Delft University of Technology

Document Version
Final published version

Licence
Dutch Copyright Act (Article 25fa)

Citation (APA)
Jafarian, M. (2026). Data-driven mechanisms of cortical memory networks in cognition. In G. Bianchin, J. Hendrickx, & R. Jungers (Eds.), *45th Benelux Meeting on Systems and Control: Book of Abstracts* (pp. 230). Katholieke Universiteit Leuven.

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

45th Benelux Meeting on Systems and Control

March 24 – 26, 2026

Center Parcs De Vossemeren, Elzen 145, 3920 Lommel, Belgium

Book of Abstracts



Institute for Information
and Communication Technologies,
Electronics and Applied Mathematics

Gianluca Bianchin, Julien Hendrickx, and Raphaël Jungers

Book of Abstracts - 45th Benelux Meeting on Systems and Control

UCLouvain
Louvain-la-Neuve, Belgium

All rights reserved. No part of the publication may be reproduced in any form by print, photo print, or by any other means without prior permission in writing from the publisher.

Part 1

Programmatic Table of Contents

Tuesday, March 24, 2026

Welcome Coffee — Fagus
Chair: — 10:15–11:00

Opening Ceremony — Fagus
Chair: Raphaël Jungers — 11:00–11:15

Plenary — Fagus
Data-Enabled Predictive Control (DeePC)
Florian Dörfler
Chair: Raphaël Jungers — 11:15–12:15

TueA01 — Fagus
Neuromorphic systems and control I
Chair: Rodolphe Sepulchre — 13:15-15:35

TueA01-1 **13.15-13.35**
Inferring Neuromodulatory Targets and Conductance Parameters From Neuronal Spike Times . 23
Julien Brandoit — University of Liège
Pierre Sacré — University of Liège
Damien Ernst — University of Liège
Julien Brandoit

TueA01-2 **13.35-13.55**
Comparing Neuromorphic Control to Sigma-delta Modulation for Nuclear Fusion Plasma Fuelling With Ice Pellets 24
Loes Jansen — DIFFER
Elena Petri — Eindhoven University of Technology
Matthijs van Berkel — DIFFER
Loes Jansen

TueA01-3 **13.55-14.15**
Modeling of Neuromorphic Circuits as Discrete Event Systems 25
Koen Scheres — KU Leuven
Rodolphe Sepulchre — KU Leuven

TueA01-4 **14.15-14.35**
Event-based Rhythmic Action Selection 26
Diego Tielens — KU Leuven
Rodolphe Sepulchre — KU Leuven

TueA01-5 **14.35-14.55**
Numerical Investigation of the Excitability of a Spiking Neuron Implemented With MOS Transistors 27
Léopold Van Brandt — UCLouvain

TueA02 — Salix
Machine Learning I
Chair: Wilm Decré — 13:15-15:35

TueA02-1 **13.15-13.35**
Constrained Extremum Seeking: Safe and Automated Engine Calibration Through In-Cylinder Pressure Shaping 28
Mats Versmissen — Eindhoven University of Technology
Maarten Vlaswinkel — Eindhoven University of Technology
Frank Willems — Eindhoven University of Technology
Mats Versmissen

TueA02-2 **13.35-13.55**
Demonstration-Free Diffusion Models for Motion Planning 29
Edward Sandra — KU Leuven
Dries Dirckx — KU Leuven
Matthias Van Eysendeyk — KU Leuven
Edward Sandra

TueA02-3 **13.55-14.15**
Randomized Spectral Node Embeddings for Graph Kernel Approximation 30
Valentin de Bassompierre — UCLouvain
Jean-Charles Delvenne — UCLouvain
Laurent Jacques — UCLouvain
Valentin de Bassompierre

TueA02-4 **14.15-14.35**
Neural Network Approximations of Discrete-Time CBF Controllers With Safety Guarantees 31
Patricia Pauli — Eindhoven University of Technology

TueA02-5 **14.35-14.55**
Gradient-Based Policy Iteration for Linear Quadratic Regulation With Finite-Time Guarantees 32
Vijayanand Digge — UCLouvain
Gianluca Bianchin — UCLouvain

TueA02-6 **14.55-15.15**
Data-Driven Multi-Step Prediction in Fed-Batch CHO Cell Cultures 33
Ali El Hadi Noura — Mons University
Alain Vande Wouwer — Mons University
Laurent Dewasme — Mons University
Ali El Hadi Noura

TueA03 — Populus
System Identification I
Chair: Philippe Dreesen — 13:15-15:35

TueA03-1 **13.15-13.35**
System Identification of a 0d Control-oriented Core Model for STEP 35
Lizan Kivits — DIFFER
Lennard Ceelen — DIFFER
Paulo Figueiredo — DIFFER
Lizan Kivits

TueA03-2 **13.35-13.55**
Model Augmentation for Thermo-mechanical Hysteretic Dynamics 36
Samarth Toolhally Eindhoven University of Technology
Siep Weiland Eindhoven University of Technology
Maarten Schoukens Eindhoven University of Technology
Samarth Toolhally

TueA03-3 **13.55-14.15**
Efficient Learning of Affine and Rationally Dependent LPV Models With Linear Fractional Representation 37
Roel Drenth Eindhoven University of Technology
Jan Hoekstra Eindhoven University of Technology
Maarten Schoukens Eindhoven University of Technology
Roel Drenth

TueA03-4 **14.15-14.35**
Screening of Impurities in Nuclear Fusion Plasmas 38
Martijn de Beij DIFFER

TueA03-5 **14.35-14.55**
Safe Extrapolation in Physics–Data Augmented Models 39
Abdullah Habboush Eindhoven University of Technology
Fahim Shakib Eindhoven University of Technology
Tom Oomen Eindhoven University of Technology
Abdullah Habboush

TueA03-6 **14.55-15.15**
Optimized Sampling Design for System Identification 40
Enrico Dozzi Eindhoven University of Technology
Tom Oomen Eindhoven University of Technology
Rodrigo Gonzalez Eindhoven University of Technology
Enrico Dozzi

TueA04 **Prunus**
Data-Driven Control I
Chair: Martina Vanelli **13:15-15:35**

TueA04-1 **13.15-13.35**
Data-driven Stabilization Using Prior Knowledge on Stabilizability and Controllability 41
Amir Shakouri University of Groningen
Henk van Waarde University of Groningen
Tren Baltussen Eindhoven University of Technology
Amir Shakouri

TueA04-2 **13.35-13.55**
Analysis of Dissipative Systems From Input-Output Data 42
Michaela Repasova University of Groningen
Henk van Waarde University of Groningen

TueA04-3 **13.55-14.15**
Safe and Efficient Iterative Controller Tuning Using Kernel Design 43
Mathias Schietecat KU Leuven
Laurens Jacobs Nikwist
Jan Swevers KU Leuven
Mathias Schietecat

TueA04-4 **14.15-14.35**
Computing Scaled Relative Graphs for MIMO LTI Systems From Frequency Response Data 44
Teun Wijfjes Eindhoven University of Technology
Timo de Groot Eindhoven University of Technology
Tom Oomen Eindhoven University of Technology
Teun Wijfjes

TueA04-5 **14.35-14.55**
Neural Approximation of Optimal Control: Sampling Strategies and Safety Guarantees 45
Andrea Giusti Flanders Make
Maxime Monsieur Flanders Make
Branimir Mrak Flanders Make
Andrea Giusti

TueA04-6 **14.55-15.15**
Distributed Data-driven Control of LTI System Using Fragmented Data 46
Surya Malladi University of Groningen
Nima Monshizadeh University of Groningen

TueA05 **Hotel Restaurant**
Optimization I
Chair: Dries Dirckx **13:15-15:35**

TueA05-1 **13.15-13.35**
System Identification and μ-synthesis for Broadband MIMO Active Noise Control in Cooling Circuits for High-tech Equipment 47
Marin Jalsic Eindhoven University of Technology
Hans Vermeulen Eindhoven University of Technology
Marc van der Wijst ASML
Marin Jalsic

TueA05-2 **13.35-13.55**
Conic Programming Duality With Applications to Lyapunov Inequalities and Linear Quadratic Control 48
Robin Bulsink University of Groningen

TueA05-3 **13.55-14.15**
Distributed Optimization of Charging Energy Hub Networks 49
Diego Fernandez-Zapico Eindhoven University of Technology
Theo Hofman Eindhoven University of Technology
Mauro Salazar Eindhoven University of Technology
Diego Fernandez-Zapico

TueA05-4 **14.15-14.35**
Hidden Convexity in Linear Neural Networks . . 50
Adeline Colson UCLouvain
Benoît Legat UCLouvain
Chenyang Yuan Toyota Research Institute
Adeline Colson

TueA05-5 **14.35-14.55**
An Efficient Metaheuristic for Optimal Sensor Placement 51
Raphaël Goetz Eindhoven University of Technology
Yuvan Dwaraga Eindhoven University of Technology
Nathan van de Wouw Eindhoven University of Technology
Raphaël Goetz

TueA05-6 **14.55-15.15**
Benign Nonconvexity in Sparse Low-rank PSD Matrix Completion 52
Sophie Lequeu UCLouvain
Chenyang Yuan MIT
Benoît Legat UCLouvain
Sophie Lequeu

TueA06 **Alnus**
Predictive Control I
Chair: Paula Chanfreut **13:15-15:35**

TueA06-1 **13.15-13.35**
Quadrotor Stabilization in Offshore Wind Conditions: Assessing the Benefit of Future Wind Information 53
Eline Wisse Eindhoven University of Technology
Bas Meere Eindhoven University of Technology
Sander Doodeman Eindhoven University of Technology
Eline Wisse

TueA06-2 **13.35-13.55**
Explainable Model Predictive Control: HVAC Case Study 54
Boulaid Boulkroune Flanders Make
Branimir Mrak Flanders Make
Bruno Depraetere Flanders Make
Boulaid Boulkroune

TueA06-3 **13.55-14.15**
A Model Predictive Track Keeping Pilot 55
Amirreza Haqshenas KU Leuven
Jan Swevers KU Leuven

TueA06-4 **14.15-14.35**
On Local Deformations of Explicit MPC Regions Under Parametric Uncertainty 56
Carlos Jose Gonzalez Rojas Eindhoven University of Technology
Valentina Breschi Eindhoven University of Technology
Leyla Özkan Eindhoven University of Technology
Carlos Jose Gonzalez Rojas

TueA06-5 **14.35-14.55**
Real-time Fleet Rebalancing: a Comparative Study of MPC and Economic MPC 57
Arthur Mélot UCLouvain
Gianluca Bianchin UCLouvain

TueA06-6 **14.55-15.15**
Dual MPC for Active Learning of Nonparametric Uncertainties 58
Tren Baltussen Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Alexander Katriniok Eindhoven University of Technology
Tren Baltussen

TueA06-7 **15.15-15.35**
Towards Cooperative Model Predictive Control in Tokamaks 59
Hari Prasad Varadarajan DIFFER

TueA07 **Betula**
Hybrid Systems I
Chair: Guillaume Berger **13:15-15:35**

TueA07-1 **13.15-13.35**
Comparison of Systems Across Time Domains . 60
Armin Pirastehzad University of Groningen
Bart Besselink University of Groningen

TueA07-2 **13.35-13.55**
On the Effect of Actuator Dynamics on the Stability of Sampled Data Systems 61
Pournami Padmalatha KU Leuven

TueA07-3 **13.55-14.15**
Bounds on Scaled Graphs of Reset Systems . . . 62
Timo de Groot Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Tom Oomen Eindhoven University of Technology
Timo de Groot

TueA07-4 **14.15-14.35**
Minimal Security Intervals for Resilience of Nonlinear Systems Under Impulsive Attacks 63
Kangping Gao Wageningen University and Research
Xiaodong Cheng Wageningen University and Research

TueA07-5 **14.35-14.55**
Incremental Input-to-state Stability for a Class of Discontinuous Piecewise Linear Systems 64
Magne Johannes Erlandsen Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Sebastiaan van den Eijnden Eindhoven University of Technology
Magne Johannes Erlandsen

TueA07-6 **14.55-15.15**
The Impact of Sensor Placement on Graph-neural-network-based Leakage Detection 65
Jarne van Gemert Eindhoven University of Technology
Valentina Breschi Eindhoven University of Technology
Mircea Lazar Eindhoven University of Technology
Jarne van Gemert

TueA08 **Fagus**
Neuromorphic systems and control II
Chair: Maurice Heemels **16:00-17:20**

TueA08-1 **16.00-16.20**
Persistent Versus Fading Memory in Serial Recall: an Analysis Through Recurrent Neural Networks 66
Lucas Stordeur University of Liège
Pierre Sacré University of Liège
Guillaume Drion University of Liège
Lucas Stordeur

TueA08-2 **16.20-16.40**
Emulation-based Spiking Control for LTI Systems 67
Elena Petri Eindhoven University of Technology
Koen Scheres
Erik Steur
Elena Petri

TueA08-3 **16.40-17.00**
Simulation of Excitable Neuromorphic Systems Across Event Scales 68
Amir Shahhosseini KU Leuven
Rodolphe Sepulchre KU Leuven

TueA08-4 **17.00-17.20**
Spiking Control for Stabilization and Tracking in LTI Systems 69
Ward Klip Eindhoven University of Technology
Elena Petri Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Ward Klip

TueA09 **Salix**
Control Theory and Systems Engineering
Chair: Arthur Fyon **16:00-17:20**

TueA09-1 **16.00-16.20**
ISS for PWA Systems 70
Radu Cimpean University of Groningen
Stephan Trenn University of Groningen
Alef Sterk University of Groningen
Radu Cimpean

TueA09-2 **16.20-16.40**
Volume-contracting Systems Are Entropy-reducing 71
Gustave Bainier University de Liège
Alessio Franci University de Liège

TueA09-3 **16.40-17.00**
Semi-physical Water Balance Model for Agricultural Polder Systems 72
Fakhira Maulana Wageningen University and Research
Martijn Wagterveld Wetsus, European Centre of Excellence for Sustainable Water Technology
Xiaodong Cheng Wageningen University and Research
Fakhira Maulana

TueA09-4 **17.00-17.20**
Structural Identifiability Determination of an Extended Biological Kinetics Model Using StrucID .
Joris Bergman Wageningen University & Research
Hans Stigter Wageningen University & Research
Karel Keesman Wageningen University & Research
Joris Bergman

TueA10 **Populus**
Energy Systems I
Chair: Lizan Kivits **16:00-17:20**

TueA10-1 **16.00-16.20**
Exploring Frequency-tunable Actuators for NTM Control in Future Nuclear Fusion Reactors . . . 74
David Becker DIFFER
Matthijs van Berkel DIFFER

TueA10-2 **16.20-16.40**
Optimal Co-design of Charging Energy Hubs for Heavy-duty Electric Transport 75
Maedeh Izadi Najafabadi Eindhoven University of Technology
Mauro Salazar, Eindhoven University of Technology
Theo Hofman Eindhoven University of Technology
Maedeh Izadi Najafabadi

TueA10-3 **16.40-17.00**
Model Predictive Control of Radiator Heating Systems With Thermostatic Valves 76
Erfan Shakhesi Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Paula Chanfreut Eindhoven University of Technology
Erfan Shakhesi

TueA11 **Prunus**
Robotics I
Chair: Jan Swevers **16:00-17:20**

TueA11-1 **16.00-16.20**
Improved Load Estimation in Hybrid Testing of Floating Wind Turbines 77
Bart Gerlagh University of Twente
Wouter Hakvoort University of Twente

TueA11-2 **16.20-16.40**
Introducing Crospi: an Open-source Constraint-based Reactive and Orchestrated Sensor-driven PIpeline for Robot Control 78
Santiago Iregui KU Leuven
Federico Ulloa Rios KU Leuven
Herman Bruyninckx KU Leuven
Santiago Iregui

TueA11-3 **16.40-17.00**
Learning UAV Navigation With Collision-Aware Predicted-Motion Embedding 79
Mohammad Javad Zallaghi Vrije Universiteit Brussel
Bryan Convens Vrije Universiteit Brussel
Kelly Merckaert Vrije Universiteit Brussel
Mohammad Javad Zallaghi

TueA11-4 **17.00-17.20**
Matrix-Free Delassus Operations: Scalable and Memory-Efficient Algorithms 80
Ajay Suresha Sathya Inria Paris
Louis Montaut Inria Paris
Yann de Mont-Marin Inria Paris
Ajay Suresha Sathya

TueA12 **Alnus**
Automated Vehicles I
Chair: Erjen Lefeber **16:00-17:20**

TueA12-1 **16.00-16.20**
Tracklet Association Using a Neural Network Model: Long Short-term Approach 81
Ruth David Eindhoven University of Technology
Jeroen Ploeg Siemens Industry Software Netherlands B.V., Helmond
Mohsen Alirezaei Siemens Digital Industries Software
Ruth David

TueA12-2 **16.20-16.40**
Velocity-through-water Estimation for Overactuated Surface Vessels Using Electric-drive Telemetry 82
Jef Billet KU Leuven
Herman Bruyninckx KU Leuven
Peter Slaets KU Leuven
Jef Billet

TueA12-3 **16.40-17.00**
Anomaly Detection in Electric Vehicles Under Dynamic Wireless Charging: Challenges and Research Gaps 83
Junhwi Mun Eindhoven University of Technology

TueA13 **Hotel Restaurant**
Optimal Control I
Chair: Flavia Sofia Acerbo **16:00-17:20**

TueA13-1 **16.00-16.20**
Design of an Optimal Positively Stabilizing State Feedback for a Spatially Discretized Heat Equation 84
Violaine Piengeon University of Namur
Joseph Winkin University of Namur

TueA13-2 **16.20-16.40**
An Optimal Power Management Policy for a Hydrogen-based Hybrid Aircraft Engine 85
Faezeh Pak Eindhoven University of Technology
Uto Perra Eindhoven University of Technology
Frank Willems Eindhoven University of Technology
Faezeh Pak

TueA13-3 **16.40-17.00**
Model-Based Fault-Tolerant Control of Multi-Engine Rocket Propulsion Systems 86
Felix Ebert TU Munich
Chiara Manfletti TU Munich

TueA13-4 **17.00-17.20**
Safe Abstraction-Based Controller Synthesis From Natural Language 87
Amir Bayat UCLouvain
Raphaël M. Jungers UCLouvain

TueA14 **Betula**
Control Applications and Discrete Event
Chair: Alexander Wierzba **16:00-17:20**

TueA14-1 **16.00-16.20**
Supervisor Synthesis for Multi-Purpose Systems . 88
Marijn MinkenbergEindhoven University of Technology

TueA14-2 **16.20-16.40**
Supervisory Control of Timed Discrete Event Systems With Partially Uncontrollable Time-interval Events 89
Mariana Guimaraes Marques Eindhoven University of Technology
Martijn Goorden Eindhoven University of Technology
Joao Carlos BasilioFederal University of Rio de Janeiro
Mariana Guimaraes Marques

TueA14-3 **16.40-17.00**
Fast Tilt-Based Electron Microscope Alignment Using A-Optimal Sensor Scheduling 90
Jilles van Hulst Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Duarte Antunes Eindhoven University of Technology
Jilles van Hulst

TueA14-4 **17.00-17.20**
Informative Experiment Design for Parameter Estimation of the Doyle-Fuller-Newman Battery Model 91
Elionai Weldeghebreal Eindhoven University of Technology

Fagus
How to Give a Scientific Talk?
Mauro Salazar
Chair: Gianluca Bianchin **17:30–18:30**

Wednesday, March 25, 2026

Plenary **Fagus**
Bits and bobs: sound synthesis with logi...
Alessandro Abate
Chair: Julien Hendrickx **08:30–09:30**

WedM01 **Fagus**
Neuromorphic systems and control III
Chair: Guillaume Drion **10:00-12:20**

WedM01-1 **10.00-10.20**
Robust Dexterous Grasping Through Virtual Hand-object Interconnection 92
Julien Vanderheyden University de Liège
Guillaume Drion University de Liège
Pierre Sacré University de Liège
Julien Vanderheyden

WedM01-2 **10.20-10.40**
Energy-based Learning in Resistor-capacitor Circuits 93
Marieke Heidema University of Groningen
Bart Besselink University of Groningen
Henk van Waarde University of Groningen
Marieke Heidema

WedM01-3 **10.40-11.00**
On Duality of Memristive Elements 94
Brayan Shali KU Leuven
Rodolphe Sepulchre KU Leuven

WedM01-4 **11.00-11.20**
A Hybrid Systems Approach for Rhythmic Neuromorphic Control of a Pendulum 95
Elena Petri Eindhoven University of Technology
Romain Postoyan University de Lorraine, CNRS, CRAN, F-54000 Nancy, France
Maurice Heemels Eindhoven University of Technology
Elena Petri

WedM01-5 **11.20-11.40**
Regulating Rhythmic Compliance in Central Pattern Generators 96
Robin Wroblowski KU Leuven
Rodolphe Sepulchre KU Leuven

WedM01-6 **11.40-12.00**
Input-Output Modeling of Neuromorphic Systems 97
Alessandro Cecconi KU Leuven
Rodolphe Sepulchre KU Leuven

WedM02 **Salix**
Machine Learning II
Chair: Fahim Shakib **10:00-12:20**

WedM02-1 **10.00-10.20**
Tempered Bayes Filtering for Improved Model-Based Estimation 98
Menno van Zutphen Eindhoven University of Technology
Domagoj Herceg Eindhoven University of Technology
Giannis Delimpaltadakis Eindhoven University of Technology
Menno van Zutphen

WedM02-2 **10.20-10.40**
Accelerating Diffusion Models for Adaptive Motion Planning in Dynamic Environments 99
Matthias Van Eysendeyk KU Leuven
Edward Sandra KU Leuven
Dries Dirckx KU Leuven
Matthias Van Eysendeyk

WedM02-3 **10.40-11.00**
MPC-Informed Residual Reinforcement Learning for Hybrid Control 100
Senne Bogaerts KU Leuven

WedM02-4 **11.00-11.20**
Can Timeseries Foundational Models Be Used for Control? 101
Lucian Cristian Iacob Flanders Make
Bruno Depraetere Flanders Make
Branimir Mrak Flanders Make
Lucian Cristian Iacob

WedM02-5 **11.20-11.40**
Path-Centric Graph Contrastive Representation Learning for Directed Acyclic Graphs 102
Jorn van Kampen Eindhoven University of Technology
Mauro Salazar Eindhoven University of Technology
Theo Hofman Eindhoven University of Technology
Jorn van Kampen

WedM02-6 **11.40-12.00**
Non-linear Model Predictive Control of Anaerobic Codigestion Using Long Short-Term Memory Networks 103
Andrés Pino-Santana Mons University
Achim Kienle Otto von Guericke Univ. / Max Planck Institute
Alain Vande Wouwer Mons University
Andrés Pino-Santana

WedM03 **Populus**
System Identification II
Chair: Rodrigo A. González **10:00-12:20**

WedM03-1 **10.00-10.20**
Improved Initialization for Port-Hamiltonian Neural Network Models 104
Gé van Otterdijk Eindhoven University of Technology
Siep Weiland Eindhoven University of Technology
Maarten Schoukens Eindhoven University of Technology
Gé van Otterdijk

WedM03-2 **10.20-10.40**
Encoder Initialisation Methods in the Model Augmentation Setting 105
Jan Hoekstra Eindhoven University of Technology
Bendeguz Györök Institute for Computer Science and Control
Roland Toth Eindhoven University of Technology
Jan Hoekstra

WedM03-3 **10.40-11.00**
Cable Hysteresis Identification and Compensation Using Frequency-domain Techniques 106
Dirk Alferink Eindhoven University of Technology
Rob Fey Eindhoven University of Technology
Nathan van de Wouw Eindhoven University of Technology
Dirk Alferink

WedM03-4 **11.00-11.20**
System Identification of Power Exhaust Dynamics in the X-Point Target Divertor at Tokamak à Configuration Variable 107
Max Winkel DIFFER
Kevin Verhaegh Eindhoven University of Technology
Matthijs van Berkel DIFFER
Max Winkel

WedM03-5 **11.20-11.40**
Least Costly Space-Filling Experiment Design for the Identification of a Nonlinear System 108
Maté Kiss Eindhoven University of Technology
Maarten Schoukens Eindhoven University of Technology
Roland Toth Eindhoven University of Technology
Maté Kiss

WedM03-6 **11.40-12.00**
A MATLAB Toolbox for OMSE-based Optimal Experimental Design in Nonlinear Bioprocess Models 109
Carlos Segura Pina University of Mons

WedM03-7 **12.00-12.20**
Contraction-preserving Unconstrained Augmentation of Lur'e-type Systems 110
Adarsh Subrahamanian Moosath Eindhoven University of Technology

WedM04 Prunus
Data-Driven Control II
Chair: Mathias Bos **10:00-12:20**

WedM04-1 **10.00-10.20**
Learning Symbolic Control Barrier Functions From Data Via Transformer-Based Generation . 111
Lufeng Xu University of Groningen
Ming Cao University of Groningen

WedM04-2 **10.20-10.40**
Feedback Stabilization of Polynomial Systems From Model-based to Data-driven Methods 112
Huayuan Huang University of Groningen
Kanat Camlibel University of Groningen
Raffaella Carloni University of Groningen
Huayuan Huang

WedM04-3 **10.40-11.00**
Scalable Nonlinear DeePC: Bridging Direct and Indirect Methods and Basis Reduction 113
Thomas de Jong Eindhoven University of Technology
Mircea Laza Eindhoven University of Technology
Siep Weilan Eindhoven University of Technology
Thomas de Jong

WedM04-4 **11.00-11.20**
Direct Vs Indirect Approach for Data-Driven Stability Analysis of Switched Linear Systems 114
Alexis Vuille UCLouvain

WedM04-5 **11.20-11.40**
Frequency-Domain Willems' Fundamental Lemma for Data-Driven Prediction With Noisy Data . . 115
Michiel Wind Eindhoven University of Technology
Tomas Meijer Lund University
Victor Dolk Eindhoven University of Technology
Michiel Wind

WedM04-6 **11.40-12.00**
*Interpolation Conditions for Data Consistency and Prediction in Noisy	Linear Systems* . . 116
Martina Vanelli UCLouvain
Nima Monshizadeh University of Groningen
Julien Hendrickx UCLouvain
Martina Vanelli

WedM04-7 **12.00-12.20**
Distributionally Robust Multi-Agent Reinforcement Learning for Intelligent Traffic Control . .
Shuwei Pei University of Groningen

WedM05 Hotel Restaurant
Optimization II
Chair: Bastiaan Vandewal **10:00-12:20**

WedM05-1 **10.00-10.20**
Contraction Estimator: a Global Optimal Solution for Robotics Optimization Problems 117
Xinghua Liu University of Groningen
Ming Cao University of Groningen

WedM05-2 **10.20-10.40**
A Two-point Random Gradient-free Method for Model-free Feedback Optimization 118
Amir Mehrnoosh UCLouvain
Gianluca Bianchin UCLouvain

WedM05-3 **10.40-11.00**
Optimization Models for Hydrogen-based Hybrid Electric Aero Engines 119
Uto Perra Eindhoven University of Technology
Faezeh Pak Eindhoven University of Technology
Theo Hofman Eindhoven University of Technology
Uto Perra

WedM05-4 **11.00-11.20**
Bandwidth Trade-offs in High-Performance, Safe Sensorless Actuators 120
Léon Borremans Vrije Universiteit Brussel

WedM05-5 **11.20-11.40**
Inertial Primal Dual Dynamics for Smooth Saddle Point Problems With Bilinear Coupling 121
Zepeng Wang University of Groningen
Juan Peypouquet University of Groningen

WedM05-6 **11.40-12.00**
Fast Optimal Control for Implicitly Constrained Dynamics 122
Louis Callens KU Leuven
Jan Swevers KU Leuven
Wilm Decré KU Leuven
Louis Callens

WedM05-7 **12.00-12.20**
Numerical Design of Online Optimization Algorithms 123
Erwan Meunier UCLouvain
Julien Hendrickx UCLouvain

WedM06 **Alnus**
Predictive Control II
Chair: Leyla Ozkan **10:00-12:20**

WedM06-1 **10.00-10.20**
Loop-Shaping for Narrow-Band Active Damping Control 124
Aditya Natu Delft University of Technology
Hassan HosseinNia Delft University of Technology

WedM06-2 **10.20-10.40**
On the Price of Constraints in Linear Tracking Model Predictive Control 125
Anastasios Vlachos University of Groningen

WedM06-3 **10.40-11.00**
Learning Cost-to-Go for Myopic Model Predictive Control Using Sobolev Training 127
Tom Minten Eindhoven University of Technology

WedM06-4 **11.00-11.20**
Exploiting Patterns for Demand Response Control With Reinforcement Learning 128
Jeroen Willems Flanders Make
Denis Steckelmacher Vrije Universiteit Brussel
Jonathan Baake Flanders Make
Jeroen Willems

WedM06-5 **11.20-11.40**
Acyclic Coupling Control: Applied to Interconnected Mechatronic Systems 129
Maarten van der Hulst Eindhoven University of Technology
Nic Dirkx ASML
Rodrigo Gonzalez Eindhoven University of Technology
Maarten van der Hulst

WedM06-6 **11.40-12.00**
Model Predictive Control of the Plasma Density Profile With Pellet Fueling in the Large Helical Device Nuclear Fusion Stellarator 130
Christopher Orrico Eindhoven University of Technology
Dinesh Krishnamoorthy Norwegian University of Science and Technology
Matthijs van Berkel DIFFER
Christopher Orrico

WedM07 **Betula**
Hybrid Systems II
Chair: Xiaodong Cheng **10:00-12:20**

WedM07-1 **10.00-10.20**
Path-complete Upper Bound on the Value Function of Switched Systems 131
Léa Ninite UCLouvain
Adrien Banse UCLouvain
Guillaume Berger UCLouvain
Léa Ninite

WedM07-2 **10.20-10.40**
Security-Aware Planning and Control of Multi-Agent Systems With LTL Tasks 132
Georgios Mitsos Eindhoven University of Technology
Dimos Dimarogonas KTH Royal Institute of Technology
Siyuan Liu Eindhoven University of Technology
Georgios Mitsos

WedM07-3 **10.40-11.00**
PCLF Extremal Lifts Are Galois Connections . . 133
Sébastien Mattenet UCLouvain
Raphaël Jungers UCLouvain

WedM07-4 **11.00-11.20**
Virtual Sensing in Urban Water Quality Monitoring Via Kalman Filtering and Machine Learning 134
Abdelrahman Habash Wageningen University and Research

WedM07-5 **11.20-11.40**
Correct-by-Design Control Synthesis of Stochastic Multi-Agent Systems: a Robust Tensor-based Solution 136
Ruohan Wang Eindhoven University of Technology

WedM07-6 **11.40-12.00**
Identification of Dynamics of Photosynthesis Regulation Under Oscillating Light Condition 137
Arviandy Aribowo Eindhoven University of Technology
Maarten Schoukens Eindhoven University of Technology
Leyla Özkan Eindhoven University of Technology
Arviandy Aribowo

WedA01 **Salix**
Neuromorphic systems and control IV
Chair: Matthijs van Berkel **13:30-15:30**

WedA01-1 **13.30-13.50**
Slow Neuronal Dynamics From Fast Activating Kinetics in Conductance-based Models 138
Arthur Fyon University of Liège
Julien Brandoit University of Liège
Kevin Jehasse University of Liège
Arthur Fyon

WedA01-2 **13.50-14.10**
Towards a Dissipativity Theory of Excitability . . 139
Guanchun Tong KU Leuven
Rodolphe Sepulchre KU Leuven

WedA01-3 **14.10-14.30**
Neuromorphic Behaviors 140
Rodolphe Sepulchre KU Leuven

WedA01-4 **14.30-14.50**
Mitigating Mismatch in Neuromorphic Circuits Using Neuromodulation Loops 141
Christian Fernandez Lorden University of Liège
Alessio Franci University of Liège

WedA01-5 **14.50-15.10**
Input-driven Timing Reliability in Mismatched Neuromorphic Neurons 142
Loris Mendolia University de Liège
Alessio Franci University de Liège

WedA02 **Populus**
Energy Systems II
Chair: Cuiling Wang **13:30-15:30**

WedA02-1 **13.30-13.50**
Joint Optimization of Hybrid Battery Sizing and Parameterized Energy Management System . . . 143
Maxime Monsieur Flanders Make
Sam Weckx Flanders Make
Ankit Surti Flanders Make
Maxime Monsieur

WedA02-2 **13.50-14.10**
Equivalent Circuit Model Parameter Identification for Lithium-ion Batteries and Adaptation to Aging 144
Ali Rahdarian University Libre de Bruxelles
Michel Kinnaert University Libre de Bruxelles

WedA02-3 **14.10-14.30**
Load-based Control Performance Testing Method for Air Conditioning Systems 145
Cuiling Wang KU Leuven
Baolong Wang Tsinghua University
Hussain Kazmi KU Leuven
Cuiling Wang

WedA02-4 **14.30-14.50**
Stochastic Model Predictive Control for Energy Management System of Grid With Hybrid Energy Storage System 146
Taranjitsingh Singh Flanders Make
Boulaid Boulkroune Flanders Make
Maxime Monsieur Flanders Make
Taranjitsingh Singh

WedA02-5 **14.50-15.10**
Control of Shared Battery Energy Storage Systems in Renewable Energy Communities: a Convex Optimization Approach 147
Alessandro Cicalese University of Groningen
Francesco Lo Iudice University of Naples Federico II

WedA02-6 **15.10-15.30**
Krasovskii Passivity-Based Control for Grid-Forming Inverters Under Matching Control . . . 148
Xinjiang Cai University of Groningen
Michele Cucuzzella University of Groningen

WedA03 **Prunus**
Robotics II
Chair: Santiago Iregui **13:30-15:30**

WedA03-1 **13.30-13.50**
Guaranteed Safe Navigation in Polytopic Free-Space: a Control Barrier Function Approach . . 149
Dries Dirckx KU Leuven
Alejandro Gonzalez-Garcia KU Leuven
Jan Swevers KU Leuven
Dries Dirckx

WedA03-2 **13.50-14.10**
Dual-arm Cooperative Control Under State and Input Constraints 150
Zemerart Asani Vrije Universiteit Brussel
Mohayad Omer Vrije Universiteit Brussel
Greet Van de Perre Vrije Universiteit Brussel
Zemerart Asani

WedA03-3 **14.10-14.30**
Slip Avoidance in Grasping Using Control Barrier Functions 151
Mohaddese Hazini University of Twente
Janset Dasdemir University of Twente
Wouter Hakvoort University of Twente
Mohaddese Hazini

WedA03-4 **14.30-14.50**
Reinforcement Learning for Navigation Skill Orchestration 152
Daan Meulendijks Eindhoven University of Technology
Jordy Senden Eindhoven University of Technology
Mehrdad Mohammadi Eindhoven University of Technology
Daan Meulendijks

WedA03-5 **14.50-15.10**
Optimization-Driven Manipulator Design For Confined Workspace 153
Chaoyue Fei Vrije Universiteit Brussel
Hafsa Nouhi Vrije Universiteit Brussel
Gaoyuan Liu Vrije Universiteit Brussel
Chaoyue Fei

WedA03-6 **15.10-15.30**
Dynamic Spline Model for Soft Robotic Grippers 154
Yulyan Hadi University of Groningen
Bayu Jayawardhana University of Groningen

WedA04 **Alnus**
Automated Vehicles II
Chair: Patrick Scheffe **13:30-15:30**

WedA04-1 **13.30-13.50**
Hardware-in-the-Loop Setup for Testing and Validation of Autonomous Driving Applications . . . 155
Shuhao Zhang KU Leuven
Jan Swevers KU Leuven
Wilm Decré KU Leuven
Shuhao Zhang

WedA04-2 **13.50-14.10**
Effects of Model Mismatch on Equivalent CACC Controller Realizations 156
Mischa Huisman Eindhoven University of Technology
Erjen Lefeber Eindhoven University of Technology
Carlos Murguia Rendon Eindhoven University of Technology
Mischa Huisman

WedA04-3 **14.10-14.30**
Control-Oriented Identification Modeling for Ship Maneuvering: A Physics-Informed and Scenario-Based Approach 157
Yan-Yun Zhang KU Leuven
Katrien Eloot Flanders Hydraulics
Jan Swevers KU Leuven
Yan-Yun Zhang

WedA04-4 **14.30-14.50**
Compensating the Effects of Hang for Agricultural Vehicles With AWS 158
Patrick Conjarts Eindhoven University of Technology
Tom van der Sande Eindhoven University of Technology
Franz Fuchshumer AGCO GmbH (Fendt)
Patrick Conjarts

WedA04-5 **14.50-15.10**
Ethical and Safety-critical Decision Making in Autonomous Driving With V2V and V2X Communication 159
Taha Bouzid Eindhoven University of Technology
Mircea Lazar Eindhoven University of Technology
Sofie Haesaert Eindhoven University of Technology
Taha Bouzid

WedA04-6 **15.10-15.30**
Criticality Measure for Scenario-based Testing of Automated Vehicles 160
Kevin Tom Kurian Eindhoven University of Technology
Erjen Lefeber Eindhoven University of Technology
Mohsen Alirezaei Eindhoven University of Technology
Kevin Tom Kurian

WedA05 **Hotel Restaurant**
Optimal Control II
Chair: Gianluca Bianchin **13:30-15:30**

WedA05-1 **13.30-13.50**
Monotone Concave Systems and Riccati Equations 161
Dongjun Wu Lund University

WedA05-2 **13.50-14.10**
Where to Place a Heavy Payload on a Multirotor UAV for Best Control Performance 162
Sander Doodeman Eindhoven University of Technology

WedA05-3 **14.10-14.30**
ASSISStaNT: Approximating Optimal Controllers for Industrial Applications 163
Andrea Giusti Flanders Make
Lucian Cristian Iacob, Flanders Make
Edward Kikken Flanders Make
Andrea Giusti

WedA05-4 **14.30-14.50**
ASSISStaNT: Retuning Approximate Optimal Controllers for Industrial Applications 164
Lucian Cristian Iacob Flanders Make
Andrea Giusti Flanders Make
Edward Kikken Flanders Make
Lucian Cristian Iacob

WedA05-5 **14.50-15.10**
Artificial Reference Time-Optimal MPC 165
Alvaro Florez KU Leuven
Jan Swevers KU Leuven
Wilm Decré KU Leuven
Alvaro Florez

WedA05-6 **15.10-15.30**
Stability of Fast Extremum Seeking Scheme for Wiener Systems Using Frequency Response Data 226
Juan Javier Palacios Roman Eindhoven University of Technology

WedA06 **Betula**
Discrete Event Systems
Chair: Koen Scheres **13:30-15:30**

WedA06-1 **13.30-13.50**
GUESS: GPU Enhanced Synthesis of Supervisory Controllers 166
Laurens van der Wal Eindhoven University of Technology

WedA06-2 **13.50-14.10**
Simultaneous Exploration of Topological and Control Space in Production Systems 167
Daniel Sarsur Eindhoven University of Technology
Patricia Pena Universidade Federal de Minas Gerais
Lucas Alves Universidade Federal de Minas Gerais
Daniel Sarsur

WedA06-3 **14.10-14.30**
A Unified Semantics-Based Framework for Timed Opacity 168
Zhe Zhang Eindhoven University of Technology
Martijn Goorden Eindhoven University of Technology
Michel Reniers Eindhoven University of Technology
Zhe Zhang

WedA06-4 **14.30-14.50**
Supervisor Synthesis for Multilevel Discrete-event Systems With Local Buses 169
Marzhan Baubekova Eindhoven University of Technology
Martijn Goorden Eindhoven University of Technology
Michel Reniers Eindhoven University of Technology
Marzhan Baubekova

WedA06-5 **14.50-15.10**
Control-oriented Modelling and Control With Stochastic Discrete Pellet Actuation for Fusion Reactors 170
Lennard Ceelen DIFFER
Christopher Orrico DIFFER
Loes Jansen DIFFER
Lennard Ceelen

WedA06-6 **15.10-15.30**
Aircraft Maintenance Scheduling Via Reinforcement Learning 171
Margot Devillers UCLouvain
Raphaël Jungers UCLouvain
Marco Saerans UCLouvain
Margot Devillers

WedA07 **Fagus**
Motion Planning and Control I
Chair: Hassan Hosseinnia **13:30-15:30**

WedA07-1 **13.30-13.50**
A Hybrid Method for Multi-agent Motion Planning on Industrial Magnetic Levitation Platforms 172
Bavo Tistaert KU Leuven

WedA07-2 **13.50-14.10**
Symbolic-Geometric Task and Motion Planning for Heterogeneous Dual-Manipulator Systems . . 173
Mehran Raisi Hatmabadi Vrije Universiteit Brussel
Gaoyuan Liu Vrije Universiteit Brussel
Mohayad Omer Vrije Universiteit Brussel
Mehran Raisi Hatmabadi

WedA07-3 **14.10-14.30**
ImpaC++t: A C++ Toolchain for Seamless Deployment of Nonlinear Model Predictive Control With Impact 174
Mathias Bos KU Leuven
Wilm Decré KU Leuven
Jan Swevers KU Leuven
Mathias Bos

WedA07-4 **14.30-14.50**
Fast Motion Planning Using Rectangular Corridor Representations of Structured Environments . . . 175
Alejandro Gonzalez-Garcia KU Leuven
Sonia De Santis KU Leuven
Jan Swevers KU Leuven
Alejandro Gonzalez-Garcia

WedA07-5 **14.50-15.10**
Co-tuning of Path Planner and Controller With Bayesian Optimization 176
Carolina Vissers Eindhoven University of Technology
René van de Molengraft Eindhoven University of Technology
Michelle Chong Eindhoven University of Technology
Carolina Vissers

WedA07-6 **15.10-15.30**
Kappa Motion Planner: A Python Framework for Analytical Motion Planning in Structured Planar Environments 177
Sonia De Santis KU Leuven
Wilm Decré KU Leuven
Jan Swevers KU Leuven
Sonia De Santis

WedA08 **Betula**
Recurrence Analysis: Methods and Applications
Chair: Martijn Boussé **16:00-18:20**

WedA08-1 **16.00-16.20**
Asymmetry and Homogeneity in β-Divergence Recurrence Plots 178
Carlos Martinez Maastricht University
Martijn Boussé Maastricht University
Philippe Dreesen Maastricht University
Carlos Martinez

WedA08-2 **16.20-16.40**
Recurrence-Based Analysis of EEG Signals for Seizure Detection Using β-Divergence 179
Filip Polakowski-Karol Maastricht University
Carlos Martinez Maastricht University
Martijn Boussé Maastricht University
Filip Polakowski-Karol

WedA08-3 **16.40-17.00**
Parameter Estimation in Dynamical Systems Based on Topology of Interacting Recurrence Networks 180
Tamas Kovacs Eötvös University

WedA08-4 **17.00-17.20**
Recurrence Analyses of Event Time-series 181
Norbert Marwan Potsdam Institute for Climate Impact Research (PIK)

WedA08-5 **17.20-17.40**
Exploring Collective Synchronisation in Football Matches Using Joint Recurrence Plots 182
Sebastian Hermann Technical University of Munich
Hendrik Meth Stuttgart Media University
Martin Lames Technical University of Munich
Sebastian Hermann

WedA09 **Fagus**
Motion Planning and Control II
Chair: Shengling Shi **16:00-18:20**

WedA09-1 **16.00-16.20**
Stability Properties of Minimal Gated Unit Neural Networks 183
Stefano De Carli University of Bergamo
Davide Previtali University of Bergamo
Mirko Mazzoleni University of Bergamo
Stefano De Carli

WedA09-2 **16.20-16.40**
Distributed Deadlock-Avoidance Control for Low-Latency Industrial Multi-Robot Systems 184
Binbin Hu University of Groningen
Brazil de Vries University of Groningen
Taco Hacquebord Philips
Binbin Hu

WedA09-3 **16.40-17.00**
Online Error Compensation of Redundant Resolvers 185
Louise Massager University Libre de Bruxelles
Michel Kinnaert University Libre de Bruxelles

WedA09-4 **17.00-17.20**
Towards Off-road Motion Planning on Elevated Terrain 186
Patrick Scheffe KU Leuven
Wilm Decré KU Leuven
Jan Swevers KU Leuven
Patrick Scheffe

WedA09-5 **17.20-17.40**
Picking Strategy for Automated Apple Harvesting 187
Marissa Jonker University of Twente
Jan de Jong University of Twente
Wesley Roozing University of Twente
Marissa Jonker

WedA10 **Salix**
Transportation Systems
Chair: Mauro Salazar **16:00-18:20**

WedA10-1 **16.00-16.20**
Energy Management and Scheduling for Regional Electric Aviation Networks 188
Finn Vehlhaber Eindhoven University of Technology
Mauro Salazar Eindhoven University of Technology

WedA10-2 **16.20-16.40**
Stability and Congestion Control of Information-Based Routing in Dynamic Transportation Networks 189
Shaya Garjani University of Groningen

WedA10-3 **16.40-17.00**
Real-time Human-in-the-loop Energy Management Policy for Electric Endurance Racing Cars . . . 190
Erik van den Eshof Eindhoven University of Technology
Jorn van Kampen Eindhoven University of Technology
Mauro Salazar Eindhoven University of Technology
Erik van den Eshof

WedA10-4 **17.00-17.20**
A Competitor-aware Race Management Policy: a Game Theoretical Approach 191
Wytze de Vries Eindhoven University of Technology
Erik van den Eshof Eindhoven University of Technology
Jorn van Kampen Eindhoven University of Technology
Wytze de Vries

WedA10-5 **17.20-17.40**
On the Engineering Trap and Responsible Innovation: Challenges and Opportunities for Systems and Control 192
Mauro Salazar Eindhoven University of Technology
Alessandro Piazza Eindhoven University of Technology

WedA10-6 **17.40-18.00**
Linear Time-Delay Control by Shaping Delay Distribution 193
Jasper Visterin KU Leuven
Wim Michiels KU Leuven

WedA10-7 **18.00-18.20**
System-Level Disturbance Control 194
Gijs Van Meerbeeck Eindhoven University of Technology
Maarten van der Hulst Eindhoven University of Technology
Nic Dirkx ASML
Gijs Van Meerbeeck

WedA11 **Populus**
Nonlinear Systems I
Chair: Dongjun Wu **16:00-18:20**

WedA11-1 **16.00-16.20**
A Density-Based Approach to the Dual Koopman Observer 195
Judicaël Mohet University of Namur
Alexandre Mauroy University of Namur
Joseph Winkin University of Namur
Judicaël Mohet

WedA11-2 **16.20-16.40**
Nonlinear Bandwidth and Bode Diagrams Based on Scaled Relative Graphs 196
Julius Krebbekx Eindhoven University of Technology
Roland Toth Eindhoven University of Technology
Amritam Das Eindhoven University of Technology
Julius Krebbekx

WedA11-3 **16.40-17.00**
Distributed Analysis of Cone-Preserving Variational Dynamics 197
Edoardo Vacchini University of Groningen
Yu Kawano Hiroshima University
Michele Cucuzzella University of Groningen
Edoardo Vacchini

WedA11-4 **17.00-17.20**
On-the-fly Surrogate Modelling With Automatic Grid-Point Allocation 198
Javier Olucha Eindhoven University of Technology
Amritam Das Eindhoven University of Technology
Roland Toth Eindhoven University of Technology
Javier Olucha

WedA11-5 **17.20-17.40**
Incremental Stability of Projection-based Control Loops 199
Riccardo Bertollo UCLouvain
Sebastiaan van den Eijnden Eindhoven University of Technology
Maurice Heemels Eindhoven University of Technology
Riccardo Bertollo

WedA11-6 **17.40-18.00**
RKHS Method for Computing Koopman-based Lyapunov Functions 200
François-Grégoire Bierwart University of Namur
Alexandre Mauroy University of Namur

WedA11-7 **18.00-18.20**
Excitable Center–Surround Recurrent Filtering for Event-Based Object–Background Segmentation 201
Fatemeh Zahed Kachaee University of Liège
Alessio Franci University of Liège

WedA12 **Prunus**
Mechanisms, Design and Control
Chair: Ajay Suresha Sathya **16:00-18:20**

WedA12-1 **16.00-16.20**
On Unwanted Nonlinearity Mitigation and Performance Analysis in Reset Control Systems 202
Ali Hosseini Delft University of Technology
Hassan HosseinNia Delft University of Technology

WedA12-2 **16.20-16.40**
Robust Performance Bounds for Component-to-assembly Uncertainty Propagation in Structural Dynamics Models 203
Rick Thijssen Eindhoven University of Technology
Rob Fey Eindhoven University of Technology
Victor Dolk Eindhoven University of Technology
Rick Thijssen

WedA12-3 **16.40-17.00**
Generalized Multilevel Coverage Using Hypergraphs 204
Ashutosh Jindal University of Groningen
Ming Cao University of Groningen

WedA12-4 **17.00-17.20**
Safe Control for Flexible Link Manipulators Using Velocity-bounding Proxy-based Sliding Mode Control 205
Hafsa Nouhi Vrije Universiteit Brussel
Chaoyue Fei Vrije Universiteit Brussel
Thierry Hubert Vrije Universiteit Brussel
Hafsa Nouhi

WedA12-5 **17.20-17.40**
A Learning Semi-Active Tuned Mass Damper Design 206
Thijs Romberg Eindhoven University of Technology
Kees Verbaan Eindhoven University of Technology
Marcel Heertjes Eindhoven University of Technology
Thijs Romberg

WedA12-6 **17.40-18.00**
Current Control Challenges for Airborne Laser Communications 207
Mauricio Sotelo Gonzalez Eindhoven University of Technology
Gert Witvoet Eindhoven University of Technology
Lex Meijer Airbus Defence and Space Netherlands
Mauricio Sotelo Gonzalez

WedA12-7 **18.00-18.20**
Evolutionary Dynamics in Continuous-time Finite-state Mean Field Games 208
Leonardo Pedroso Eindhoven University of Technology
Andrea Agazzi University of Bern
Maurice Heemels Eindhoven University of Technology
Leonardo Pedroso

WedA13 **Hotel Restaurant**
Model Reduction
Chair: Pierre Sacré **16:00-18:20**

WedA13-1 **16.00-16.20**
Abstraction-Based Model Reduction of Uncertain Systems 209
Melvyn Wildeboer University of Groningen
Jacquelien Scherpen University of Groningen
Bart Besselink University of Groningen
Melvyn Wildeboer

WedA13-2 **16.20-16.40**
Confidentiality of Linear Systems With Quadratic Output 210
Zeyad Manaa Eindhoven University of Technology
Nathan van de Wouw Eindhoven University of Technology
Michelle Chong Eindhoven University of Technology
Zeyad Manaa

WedA13-3 **16.40-17.00**
Global Synchronization Preserving Model Reduction of Nonlinear Network Systems 211
Yangming Dou University of Groningen
Xiaodong Cheng Wageningen University and Research
Jacquelien Scherpen University of Groningen
Yangming Dou

WedA13-4 **17.00-17.20**
Neural Linear Oscillator Networks and Their Application to Soft Robots Model Learning and Control 212
Jingyue Liu Delft University of Technology
Ebrahim ShahabiShalghouni Delft University of Technology
Cosimo Della Santina Delft University of Technology
Jingyue Liu

WedA13-5 **17.20-17.40**
Bimodal Atomic Force Microscopy With Nonlinear Control for High-Speed Metrology 213
Jonas Hendrikx Eindhoven University of Technology
Erik Steur Eindhoven University of Technology
Robbert Voorhoeve TNO
Jonas Hendrikx

WedA13-6 **17.40-18.00**
Moment Matching by Recursive-Based Learning . 214
Francesco Corrini University of Bergamo
Matteo Scandella University of Bergamo
Alessio Moreschini Imperial College London
Francesco Corrini

WedA13-7 **18.00-18.20**
Demographic Dependence of Vaccine Adoption Under Opinion Persuasion 215
Alessandro Casu Eindhoven University of Technology

WedA14 **Alnus**
Planning and Control under Uncertainty
Chair: Mircea Lazar **16:00-18:20**

WedA14-1 **16.00-16.20**
Optimal Controllers in Stochastic Abstractions and Sequential Games 216
Adrien Janssens UCLouvain
Raphaël Jungers UCLouvain

WedA14-2 **16.20-16.40**
Consolidating Probabilistic Optimal Control . . . 217
Ajinkya Bhole Ghent University
Mohammad Mahmoudi Filabadi Ghent University
Guillaume Crevecoeur Ghent University
Ajinkya Bhole

WedA14-3 **16.40-17.00**
Accurate Remnant Positioning Through Combined Hysteresis and Creep Modeling 218
Jurrien Keulen University of Groningen
Koen Classens ASML
Monica Acuautla University of Groningen
Jurrien Keulen

WedA14-4 **17.00-17.20**
Optimal Setpoint Adaptation Under Regime Change for Control of Mechatronic Systems . . . 219
Foeke Vanbecelaere Flanders Make

WedA14-5 **17.20-17.40**
Funnel Control for TCP Congestion Control With State-dependent Input Delays 220
Yanxin Li University of Groningen
Stephan Trenn University of Groningen

WedA14-6 **17.40-18.00**
Probabilistic Safety Under Arbitrary Disturbance Distributions Using Piecewise-Affine Control Barrier Functions 221
Matisse Teuwen KU Leuven
Mathijs Schuurmans KU Leuven
Panagiotis Patrinos KU Leuven
Matisse Teuwen

Fagus
Benelux Networking Drink
Chair: **18:30–19:30**

Thursday, March 26, 2026

Plenary **Fagus**
Mini Course: Decentralized Learning of E...
Maryam Kamgarpour
Chair: Gianluca Bianchin **08:30–10:45**

ThuM01 **Populus**
Nonlinear Systems II
Chair: Riccardo Bertollo **11:15-12:15**

ThuM01-1 **11.15-11.35**
Constrained Performance Boosting Control for Nonlinear Systems Via ADMM 222
Gianluca Giacomelli Eindhoven University of Technology
Danilo Saccani EPFL
Siep Weiland Eindhoven University of Technology
Gianluca Giacomelli

ThuM01-2 **11.35-11.55**
Compositional Design for Time-varying and Nonlinear Coordination 223
Jonas Hansson UCLouvain

ThuM01-3 **11.55-12.15**
Freq-statespace: a Python Package for Identification of (non)linear State-space Models From Periodic Input-output Data 224
Merijn Floren KU Leuven
Jan Swevers KU Leuven

ThuM02 **Hotel Restaurant**
Adaptive Systems
Chair: Alessio Franci **11:15-12:15**

ThuM02-1 **11.15-11.35**
Adaptive-Gain Control of Rational Decision Makers, Global Convergence and Comparative Performance Study 225
Rory Gavin University of Groningen
Ming Cao University of Groningen

ThuM02-2 **11.35-11.55**
Adaptive Control for Resilient Wireless Optical Communication Links 227
Arturo Morales Eindhoven University of Technology
Rodrigo Gonzalez Eindhoven University of Technology
Tom Oomen Eindhoven University of Technology
Arturo Morales

ThuM03 **Prunus**
Data-Driven Control III
Chair: Guillaume Berger **11:15-12:15**

ThuM03-1 **11.15-11.35**
Experiment Design for Continuous-time Systems Using Generalized Filtering 228
Jiwei Wang University of Groningen
Simone Baldi Southeast University
Henk van Waarde University of Groningen
Jiwei Wang

ThuM03-2 **11.35-11.55**
Data-driven Oscillatory Model of Memory Consolidation: Perspective for Control by Optogenetics 229
Mario Valencia Torres Delft University of Technology
Joshua Amato Delft University of Technology
Matin Jafarian Delft University of Technology
Mario Valencia Torres

ThuM03-3 **11.55-12.15**
Data-driven Mechanisms of Cortical Memory Networks in Cognition 230
Matin Jafarian Delft University of Technology

ThuM04 **Fagus**
Neuromorphic systems and control V
Chair: Alexander Pogromsky **11:15-12:15**

ThuM04-1 **11.15-11.35**
Spatio-temporal Excitation With Negative Threshold 231
Thomas Burger KU Leuven
Rodolphe Sepulchre KU Leuven

ThuM04-2 **11.35-11.55**
Context-dependent Manifold Learning in Dynamical Systems: A Neuromodulated Constrained Autoencoder Approach 232
Jérôme Adriaens University of Liege
Guillaume Drion University of Liege
Pierre Sacré University of Liege
Jérôme Adriaens

ThuM04-3 **11.55-12.15**
Inter-personal Variability in Muscle Synergies for Gait Control 233
Rodrigo Moreno-Morton University of Liege
Alessio Franci University of Liege

ThuM05 **Alnus**
Reinforcement Learning and applications
Chair: Elena Petri **11:15-12:15**

ThuM05-1 **11.15-11.35**
Reinforcement Learning for Diverse Locomotion of the Human Musculoskeletal System 234
Kelvin Toonen Delft University of Technology
Shengling Shi Delft University of Technology

ThuM05-2 **11.35-11.55**
Teacher–Student Reinforcement Learning for Reeds–Shepp Vehicles: Kinematically Feasible Motion Planning 235
Ibrahim Ibrahim KU Leuven
Wilm Decré KU Leuven
Jan Swevers KU Leuven
Ibrahim Ibrahim

ThuM05-3 **11.55-12.15**
AI-enhanced wastewater treatment process automation 236
Soheil Hadisi Wageningen University and Research,
Wagterveld, R.M
Mels, A.
Soheil Hadisi

ThuM06 **Betula**
Distributed Control and Estimation
Chair: Binbin Hu **11:15-12:15**

ThuM06-1 **11.15-11.35**
Overlapping Covariance Intersection: Fusion With Partial Structural Knowledge of Correlation From Multiple Sources 237
Leonardo Pedroso Eindhoven University of Technology
Pedro Batista University of Lisbon
Maurice Heemels Eindhoven University of Technology
Leonardo Pedroso
ThuM06-2 **11.35-11.55**
Biologically Inspired Sparse Event-Based Communication and Decoding of Spatiotemporal Stimuli 238
Ben Efron The University of Liège
Guillaume Drion The University of Liège
Pierre Sacré The University of Liège
Ben Efron
ThuM06-3 **11.55-12.15**
Design of Reduced Distributed Observers for Nonlinear Systems in Prime Form 239
Yaodong Li Eindhoven University of Technology
Nathan van de Wouw Eindhoven University of Technology
Erik Steur Eindhoven University of Technology
Yaodong Li

Fagus
Awards Ceremony
Chair: Gianluca Bianchin **13:15–14:15**

Fagus
Closing Ceremony
Chair: Julien Hendrickx **14:15–14:30**

Part 1: Programmatic Table of Contents 3
Overview of scientific program

Part 2: Abstracts 21

Part 3: List of Participants 241
Alphabetical list

Part 4: Organizational Comments 258
Comments, overview program, map

Part 2

Abstracts

Inferring neuromodulatory targets and conductance parameters from neuronal spike times

Julien Brandoit, Pierre Sacré, Damien Ernst, Arthur Fyon, Guillaume Drion
Department of Electrical Engineering and Computer Science, University of Liège, Belgium
{jbrandoit, p.sacre, dernst, afyon, gdrion}@uliege.be

Neuromodulators such as dopamine and histamine continuously reshape neuronal excitability by acting on ion channel densities and kinetics. This process is central to brain function; however, it remains difficult to characterize experimentally: metabotropic signaling cascades make the mapping from neuromodulator to ion channel nontrivial, and direct measurements of affected conductances are rarely feasible.

A key obstacle is neuronal degeneracy: distinct combinations of ion channels can generate similar spiking patterns, so neuromodulatory effects cannot be understood from a single "average" model. Standard approaches that attempt to fit a unique model from data therefore miss the richness of possible solutions [1].

We consider conductance-based neuron models, in which membrane dynamics are governed by ionic currents,

$$C\frac{dV}{dt} + g_{\text{leak}}(V - E_{\text{leak}}) = -\sum_{i\in\mathscr{I}} \bar{g}_i m_i^{p_i} h_i^{q_i}(V - E_i) + I_{\text{ext}},$$

where the unknown parameters are the maximal conductances $\bar{g} = [\bar{g}_1, \ldots, \bar{g}_{|\mathscr{I}|}, g_{\text{leak}}]$. Due to degeneracy, the set of solutions consistent with a given activity forms a manifold rather than a point.

We introduce a computational framework that reconstructs populations of conductance-based models directly from spike times, a widely available experimental data. The method relies on Dynamic Input Conductances (DICs), an interpretable three-dimensional mixed-feedback representation that aggregates fast, slow, and ultraslow dynamical contributions of ion channels to excitability [2]. A deep learning model maps spike trains to DICs, from which populations of degenerate conductance vectors consistent with activity x are efficiently generated [3].

Under neuromodulation, as spike patterns evolve from x_{before} to x_{after}, entire populations shift in parameter space. We identify plausible modulatory targets by minimizing parameter changes while maintaining physiological constraints. Reachability maps visualize accessible regions of DIC space and encode the cost of reaching target states, linking output changes to feasible parameter transitions (fig. 1).

By treating degeneracy as an asset rather than a limitation, the framework reveals how robustness emerges from shifting populations rather than fixed parameters, offering new insights into how neuromodulation reshapes neuronal excitability at the population level.

Acknowledgment This work is supported by the Belgian Government through the Federal Public Service Policy and Support. JB and AF are, respectively, a FRIA grantee and a Research Fellow of the Fonds de la Recherche Scientifique – FNRS.

References

[1] E. Marder and A. L. Taylor, "Multiple models to capture the variability in biological neurons and networks," *Nature Neuroscience*, vol. 14, no. 2, pp. 133–138, 2011.

[2] G. Drion et al., "Dynamic input conductances shape neuronal spiking," *eNeuro*, vol. 2, no. 1, 2015.

[3] J. Brandoit et al., "Fast reconstruction of degenerate populations of conductance-based neuron models from spike times," *arXiv preprint arXiv:2509.12783*, 2025.

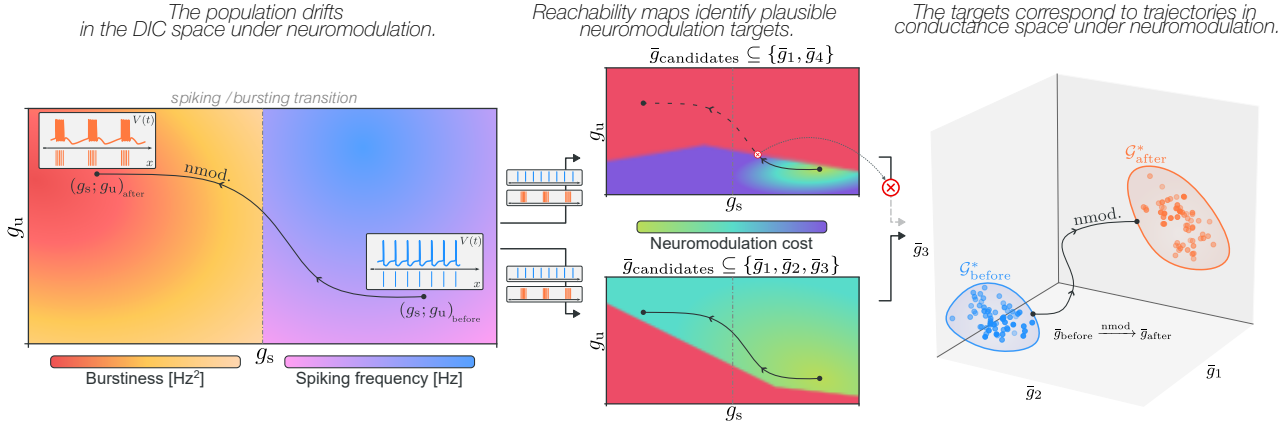

Figure 1: Population-level inference of ion channel conductances from spike times before and after neuromodulation, with automatic identification of the most plausible conductance targets. Unreachable regions are shown in red; reachable regions (green-purple) are colored by neuromodulation cost.

Comparing neuromorphic control to sigma-delta modulation for nuclear fusion plasma fuelling with ice pellets

L.L.T.C. Jansen[1,2,*], E. Petri[2], M. van Berkel[1], W.P.M.H. Heemels[2]

[1] *DIFFER - Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands.*
[2] *Eindhoven University of Technology, Department of Mechanical Engineering, Control Systems Technology section, Eindhoven, the Netherlands*
[*] *Email:* l.l.t.c.jansen@differ.nl

1 Energy optimisation in a fusion reactor

A tokamak generates energy through deuterium-tritium fusion, producing helium, neutrons, and energy. Fusion power scales with plasma density, making precise density control essential for optimising energy output. Plasma density is regulated by firing small cryogenic deuterium-tritium fuel pellets directly into the hot plasma [1].

2 Neuromorphic control vs. sigma-delta modulation

The open-loop plasma density dynamics are modelled as $\dot{n}_e = -\frac{1}{\tau} n_e(t)$, with τ a time constant. To regulate the density n_e to a prescribed reference r, a neuromorphic (NM) controller with an integrate-and-fire neuron model based on [2], adapted and analysed for pellet fuelling in [3], is compared with sigma-delta modulation (SDM) based on its implementation on the tokamak ASDEX-Upgrade (AUG) [1]. In both cases, the SISO controller integrates the error between the reference and measured density through an internal state ξ, analogous to a neuron membrane potential in NM control or the sigma variable in SDM. When ξ reaches a positive threshold Δ, the controller emits a spike corresponding to a pellet injection. Pellet injections are constrained to discrete time instants imposed by the injector hardware, potentially introducing a delay between threshold crossing and pellet delivery. When a pellet is fired, ξ is reset differently for NM control and SDM.

In the **NM controller**, the state is reset to 0 after firing

$$\xi^+ = 0. \tag{1}$$

In the **SDM**, the threshold value Δ is subtracted from ξ

$$\xi^+ = \xi - \Delta. \tag{2}$$

For both controllers we proved the density converges to the desired reference up to a quantifiable bounded error. During the transient phase, both controllers exhibit identical behaviour, injecting a pellet at every available time slot. In steady state, SDM is expected to outperform the NM controller, as it avoids a full reset of the internal state ξ and instead retains timing error information. This memory enables compensation by advancing subsequent pellet injections, since ξ reaches the threshold Δ more rapidly.

3 Results

Results from [3] show that, for the NM controller with constraints on the actuator speed, we get

$$n_e \in (r-\alpha, r+\alpha] \tag{3}$$

in steady state, where α is the jump size, i.e., the increase in density from one pellet.

SDM exhibits integrator wind-up, where two mitigation options are evaluated under the same actuator speed constraint.

1. Saturation on the input: similar to the controller implemented on AUG [1], we limit the value that goes into the SDM. In steady state, this leads to
$$n_e \in (r-2\alpha, r+\alpha]. \tag{4}$$
2. We change the reset condition for ξ after shooting a pellet to $\xi^+ = \xi - k\Delta$ with $k = \left\lfloor \frac{\xi}{\Delta} \right\rfloor$, the largest integer smaller than or equal to $\frac{\xi}{\Delta}$. In steady state, we get the same bound as the NM controller in (3).

Simulations confirm the theoretical results. SDM does not outperform the NM controller, even the modified SDM in option 2 shows no improvement over NM control. Correcting timing errors in SDM provides no additional benefit compared with fully resetting the internal state ξ in the NM controller. Moreover, the NM controller is simpler to implement and analyse.

References

[1] B. Ploeckl et al., "Targeting a Versatile Actuator for EU-DEMO: Novel Control Scheme for Multisource Pellet Injector", Fusion Science and Technology, 77(3), 2021.

[2] E. Petri et al., "Analysis of a Simple Neuromorphic Controller for Linear Systems: A Hybrid Systems Perspective", IEEE Conference on Decision and Control, 2024.

[3] L. Jansen et al., "Nuclear fusion plasma fuelling with ice pellets using a neuromorphic controller", IEEE Conference on Decision and Control, 2025.

Acknowledgement

DIFFER is part of the institutes organisation of NWO. This work was supported by the research program VIDI with project number 19695, which is (partly) financed by NWO.

Modeling of Neuromorphic Circuits as Discrete Event Systems

Koen Scheres[1], Rodolphe Sepulchre[1,2]

[1]Department of Electrical Engineering (ESAT), KU Leuven, Kasteelpark Arenberg 10, B-3001 Leuven, Belgium.

[2]Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, United Kingdom.

Email: *koen.scheres@kuleuven.be*, *rodolphe.sepulchre@kuleuven.be*

Although excitable neuromorphic circuits are physical models which produce continuous-time trajectories, these trajectories consist of sequences of discrete events. This presentation explores the possibility of extracting a discrete-event model out of the physical continuous-time model. We discuss the potential of this methodology for analysis and design of neuromorphic control systems.

Hierarchical Control

Control and estimation problems can often be separated in three categories: tracking, trajectory planning and decision-making. Although this separation is necessary in many cases (e.g., to make analysis tractable), it requires a hierarchical control structure for problems where all three are required [1]. This hierarchical structure is depicted in Figure 1.



Figure 1: By separating the control stack it is possible to do fast tracking as well as slow decision-making [1].

Consequently, several different controllers are often designed (and potentially implemented) in different domains and/or devices: the decision-making process is usually addressed in terms of automata, while the feedback control problem is typically solved via differential equations. This clear separation between the discrete and the continuous is not present in the animal world. Instead, central nervous systems orchestrate neuronal behavior, that is, *sequences of events*. At the cellular level, those events are made of "spikes." Spikes are modeled by biophysical differential equations, yet they can be counted. Hence, they combine attributes of the continuous *and* the discrete [2].

Discrete Event Systems

If we want to regard excitable systems as a unifying framework for regulation and automation, it seems crucial to combine continuous *and* discrete descriptions of their behavior. The key advantage of such an approach is that the same architecture can be used throughout the hierarchical control stack. For example, for low-level control, we can exploit *central pattern generators*, which are essential building blocks in the central nervous system for locomotion [3]. On the other hand, for high-level control we can use the *winner-take-all* principle, which was pioneered by Hopfield [4] in the classical Hopfield networks, and later shown to enable interesting computation capabilities and even universal approximation properties in some cases [5]. Many models of neurons are hybrid, but we are not aware of any methodological attempt to extract a discrete-event model from the biophysical model. However, leveraging the discrete nature allows us to endow neuromorphic circuits—electronic circuits whose operating principles are inspired by neurons—with discrete decision-making capabilities.

This background motivates our question: what is the discrete-event model of a continuous-time biophysical model of excitable neurons? The answer to this question allows us to systematically solve verification problems (which may be phrased, e.g., as temporal logic) or exploit other tools available in the discrete event system literature, such as automated synthesis of controllers, applied to neuromorphic circuits. The latter could be useful in the context of designing behaviors, see, e.g., [6]. Our approach is thus a complementary approach to the biophysical modeling as (conductance-based) differential equations.

This abstract is based on [7].

References

[1] N. Matni, A. D. Ames, and J. C. Doyle, "A quantitative framework for layered multirate control: Toward a theory of control architecture," *IEEE Control Systems*, vol. 44, no. 3, pp. 52–94, 2024.

[2] R. Sepulchre, "Spiking control systems," *Proceedings of the IEEE*, vol. 110, no. 5, pp. 577–589, 2022.

[3] E. Marder and D. Bucher, "Central pattern generators and the control of rhythmic movements," *Current Biology*, vol. 11, no. 23, pp. R986–R996, 2001.

[4] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," *Proceedings of the National Academy of Sciences*, vol. 81, no. 10, pp. 3088–3092, 1984.

[5] W. Maass, "Neural computation with winner-take-all as the only nonlinear operation," in *Advances in Neural Information Processing Systems*, vol. 12. MIT Press, 1999.

[6] Y. Huo, F. Forni, and R. Sepulchre, "A neuromorphic architecture for scalable event-based control," 2025, arxiv:2511.11924.

[7] K. Scheres and R. Sepulchre, "Discrete event system modeling of neuromorphic circuits," 2025, arXiv:2511.15594.

Event-based rhythmic action selection

Diego Tielens
KU Leuven, Leuven 3001, Belgium
diego.tielens@kuleuven.be

Rodolphe Sepulchre
KU Leuven, Leuven 3001, Belgium
rodolphe.sepulchre@kuleuven.be

1 Introduction

The action selection problem is the challenge of deciding "what to do next" to satisfy an organism's or machine's goals. Behavior possesses both continuous and discrete aspects: the world and an agent's internal processes often evolve continuously, but the agent must make discrete decisions about which action to take and when. The problems of trajectory regulation and of decision-making are often addressed separately from each other. Reactive architectures such as Brook's subsumption [2] or Braitenberg's vehicles [1] are strong at continuous reactivity, while reinforcement learning architectures are better at discrete commitment. Although many hybrid architectures try to reconcile these continuous and discrete dimensions, the field still lacks a framework where regulation and decision-making are mutually enabling rather than separate questions to be reconciled. Our research proposes a new way to address the action selection problem by building a neuromorphic architecture that generates a rhythm, which works as an internal model selecting which action to perform next and adapts to changing environments in an event-based manner [3]. We explore how the event-based nature of neuromorphic behaviors sheds new light on the classical divide between "regulation" and "automation", "feedforward" and "feedback", but also between "internal" and "external" goals. Inspired by Braitenberg [1], we adopt a bottom-up approach with simple systems reconciling simple goals.

2 A rhythmic action selection architecture

To address the action selection problem, our research studies a neuromorphic network in which only one neuron is active at a time to select the task to accomplish. It functions as a winner-take-all architecture and generates a rhythm that produces a sequence of actions with variable durations. Unlike hybrid architectures, it does not reconcile continuous and discrete behavior by layering them, but by encoding both within the same neural dynamics. Tasks are hierarchical because they can be broken down into sub-tasks. An additional advantage of a rhythm is that it naturally captures this hierarchical structure [3]. Additionally, computational complexity is reduced because the event-based framework adjusts only in response to disturbing events and adapts the rhythm up to the relevant hierarchical level based on their importance.

3 The conflict between the internal and the external

The rhythm is generated endogenously, but also adapts to external events [5]. In this way, the endogenous rhythm functions as a feedforward internal model, representing the sequence of actions the embodied machine should execute while remaining aligned with sensory feedback. Put differently, the machine acts according to what it has learned to be most effective, yet continuously reassesses and adjusts when unexpected events occur in the environment.
It is known that animal behavior depends on both internal bodily states and external environmental conditions [6]. This is fundamentally important for embodied machines: to be truly autonomous, they must be able to reconcile internal and external goals by regulating both their internal and external environments. Inspired by Braitenberg's work [1] and extensions of it [4], our research will start with simple scenarios in which an embodied machine has to accomplish externally demanded tasks, such as collecting parcels, while fullfilling internal requirements, such as recharging batteries when needed. We adress the complementarity of "interoception" and "exteroception" in embodied machines [6].

References

[1] V. Braitenberg. *Vehicles: experiments in synthetic psychology.* MIT press, Cambridge (Mass.), 1984.

[2] R. Brooks. A robust layered control system for a mobile robot. *IEEE journal of robotics and automation*, 2(1):14–23, 1986.

[3] Yongkang Huo, Fulvio Forni, and Rodolphe Sepulchre. A neuromorphic architecture for scalable event-based control. 2025. https://arxiv.org/abs/2511.11924.

[4] D. Lambrinos and C. Scheier. Building complete autonomous agents: a case study on categorization. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1:170–177, 1996.

[5] R Schmetterling, F Forni, A Franci, and Rodolphe Sepulchre. Neuromorphic control of a pendulum. *IEEE Control Systems Letters*, vol. 8:1235–1240, 2024.

[6] B. Toussaint, J. Heinzle, and K. E. Stephan. A computationally informed distinction of interoception and exteroception. *Neuroscience and biobehavioral reviews*, 159:105608–, 2024.

Numerical Investigation of the Excitability of a Spiking Neuron Implemented with MOS Transistors

Léopold Van Brandt, Grégoire Brandsteert, Jean-Charles Delvenne and Denis Flandre
ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium
leopold.vanbrandt@uclouvain.be

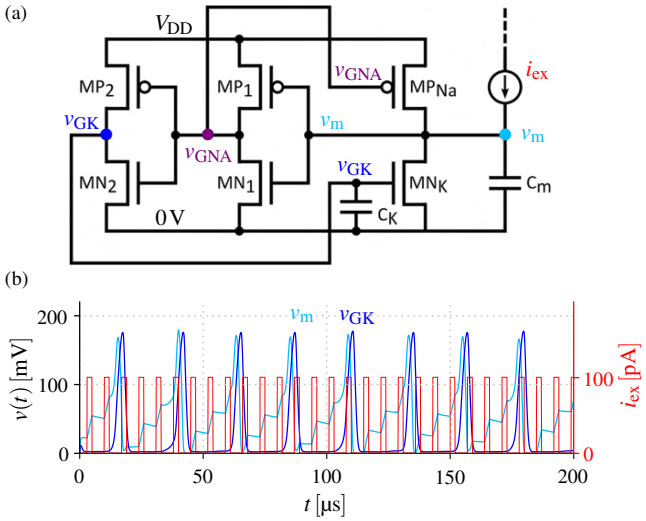

Figure 1: (a) Diagram of a CMOS analog spiking neuron [1].
(b) Action potential train (output spike train) sustained by a synaptic excitation current spike train [2].
Case study: 65 nm CMOS technology design of [1].

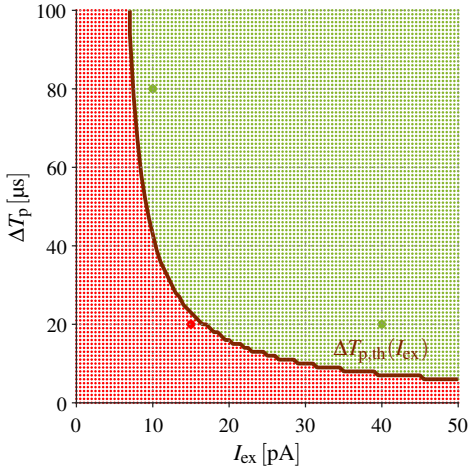

Figure 2: Threshold property of the excitable neuron of Figure 1(a), extracted from deterministic SPICE transient simulations with a sweep on excitation current pulse parameters $(I_{ex}, \Delta T_p)$, resp. amplitude and duration [2], suggested by [3, 4]. The green region corresponds to input pulses i_{ex} that lead to an action potential, i.e. a spike on output v_{GK}. Low-energy pulses located in the red region fail to trigger the spiking mechanism.

1 Introduction

Within the neuromorphic paradigm, spiking neurons are the fundamental building blocks of spiking neural networks. A large variety of neuron architectures can be found, either in CMOS technology or based on emerging devices such as volatile memristors. An ultra-low-power CMOS design proposed by [1] is shown in Figure 1(a). The work [2] investigated its *excitability* with numerical SPICE simulations, thereby complementing papers [3, 4] that worked with theoretical simplified neuron models.

2 Excitability of the Neuron

In Figure 2, we have assessed the excitability of the neuron of Figure 1(a) through SPICE noiseless transient simulations. Starting from cold initial conditions, we have injected synaptic current pulses sweeping amplitudes I_{ex} and durations ΔT_p [2]. For each case, we observe the output potential v_{GK}: the neuron has been successfully excited if one output spike is recorded (green region in Figure 2); pulses that are too low and short to trigger the spiking mechanism are located in the red region of Figure 2.

Observation suggests to determine a threshold criterion that would take into account both the amplitude *and* the duration of the injected pulse. However, our numerical investigation reveals that the critical charge $\Delta Q_{th} = I_{ex} \cdot \Delta T_{p,th}$ *does* vary with I_{ex} and hence is not a unique quantity, echoing [5, Ch. 1]. The *spiking criterion* is that the membrane potential $v_m(t)$ exceeds a certain threshold $V_{m,th}$, which is invariable and intrinsic to the neuron and its internal dynamics.

References

[1] I. Sourikopoulos, S. Hedayat, C. Loyez, F. Danneville, V. Hoel, E. Mercier, and A. Cappy, "A 4-fj/spike artificial neuron in 65 nm cmos technology," *Frontiers in neuroscience*, vol. 11, p. 123, 2017.

[2] L. Van Brandt, G. Brandsteert, and D. Flandre, "On the Excitability of Ultra-Low-Power CMOS Analog Spiking Neurons," in *IEEE LASCAS 2026, Arequipa, Peru*, 2026.

[3] R. Sepulchre, G. Drion, and A. Franci, "Excitable behaviors," in *Emerging Applications of Control and Systems Theory: A Festschrift in Honor of Mathukumalli Vidyasagar*. Springer, 2018, pp. 269–280.

[4] R. Sepulchre and G. Tong, "On the threshold of excitable systems: An energy-based perspective," *arXiv preprint arXiv:2504.02171*, 2025.

[5] E. M. Izhikevich, *Dynamical Systems in Neuroscience*. MIT press, 2007.

Constrained Extremum Seeking: Safe and Automated Engine Calibration through In-Cylinder Pressure Shaping

Mats Versmissen, Maarten Vlaswinkel and Frank Willems
Eindhoven University of Technology
Control Systems Technology
E-mail: m.versmissen@tue.nl

1 Introduction

The demand for clean and efficient combustion alternatives to the standard compression ignition (CI) engine is driven by legislation. This has led to an increased focus on advanced combustion concepts, such as dual-fuel combustion, which promises higher thermal efficiencies and lower engine-out emissions. However, finding the optimal engine control parameters is more difficult compared to their CI counterparts. Auto-calibration can accelerate the development of control strategies for advanced combustion engines and reduce reliance on expert knowledge. This work proposes a constrained Extremum Seeking Control (ESC) algorithm to optimize the fuel path actuators of a dual-fuel diesel-gasoline engine [1]. The ESC algorithm has been successfully demonstrated on the single-cylinder engine (SCE) setup in the Zero Emission Lab at the Eindhoven University of Technology.

2 In-Cylinder Pressure Shaping

Any arbitrary in-cylinder pressure trace $p_{cyl}(\theta)$ can be decomposed into $n_{\text{PC}} \in \mathbb{R}$ Principal Components (PCs) as [2]:

$$p_{cyl}(\theta,\tilde{u}) = w^{\top}(\tilde{u}) f(\theta) + f_{mp}(\theta) \tag{1}$$

where $w(\tilde{u})$ are the associated PC weights, controllable via the fuel path actuators $\tilde{u}$. This novel framework presents a direct pressure shape control approach on a cycle-to-cycle basis, independent of the underlying combustion concept.

3 In-Cylinder Pressure-Based Optimization

The presented ESC algorithm shapes the in-cylinder pressure trace towards the Ideal Thermodynamic Cycle (ITC), thereby increasing the Gross Indicated Efficiency (GIE). Using the aforementioned Principal Component Decomposition (PCD), a pressure based cost function is obtained [3]:

$$J(w(\tilde{u})) = (w(\tilde{u}) - w_{\text{ITC}}(\tilde{u}))^{\top} Z_1 (w(\tilde{u}) - w_{\text{ITC}}(\tilde{u})) \tag{2}$$

where $\tilde{u} = [\text{SOI}_{\text{DI}}, \text{BR}]$ with SOI_{DI} the start of diesel injection and BR the energy-based blend ratio. The ESC algorithm is able to comply with the requested driver demand (IMEP_{ref}) and can adhere to two combustion safety constraints (maximum pressure p_{max} and maximum pressure rise rate dp_{max}). The overall optimization problem is written as:

$$\tilde{u}^* = \min_{\tilde{u}\in\mathscr{U}} \quad J(w(\tilde{u})) \approx \text{GIE}(\tilde{u}) \tag{3a}$$

$$\text{s. t.} \quad \text{IMEP}_{\text{g}}(\tilde{u}) = \text{IMEP}_{\text{ref}} \tag{3b}$$

$$w(\tilde{u})^{\top} f(\theta) + f_{\text{mp}}(\theta) \leq p_{max,ub} \tag{3c}$$

$$w(\tilde{u})^{\top} \frac{df(\theta)}{d\theta} + \frac{df_{\text{mp}}(\theta)}{d\theta} \leq dp_{max,ub} \tag{3d}$$

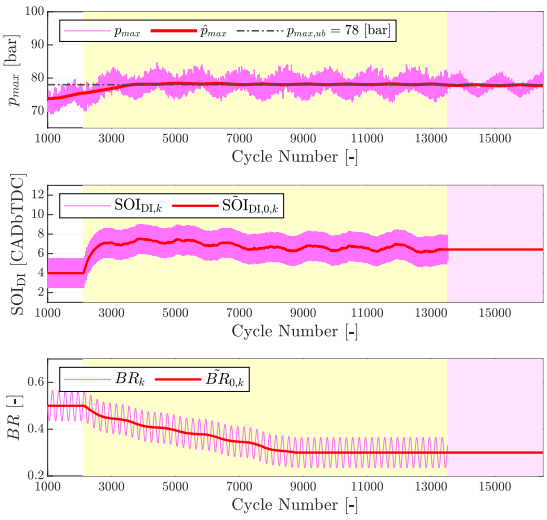

Figure 1: Experimental results of the ESC algorithm.

The ES algorithm employs gradient descent to safely determine the optimal fuel path actuators while satisfying the IMEP_{ref}-constraint via a parallel next-combustion cycle controller.

4 Experimental Results

Figure 1 displays the experimental results from the ESC algorithm (active in yellow area). By starting at: SOI_{DI} = 4 [CADbTDC] and BR = 0.5 [-] and setting an artificial maximum pressure constraint at 78 [bar], the algorithm safely converged to a stable setpoint (pink area). The new fuel path references are SOI_{DI} = 6.4 [CADbTDC] and BR = 0.3 [-], which resulted in a 3.77% increase in the GIE compared to the baseline of the predefined engine map. The driver demand is successfully controlled (4 ± 0.25 [bar]) via a next-cycle combustion controller.

5 Conclusion

The experimental results of the ESC algorithm have demonstrated that, for advanced combustion concepts, auto-calibration methods can be safely deployed on-line to accelerate engine control development.

References

[1] M. Versmissen, M. Vlaswinkel and F. Willems, "Experimental Demonstration of Safe and Automated In-Cylinder Pressure Shaping using Constrained Extremum Seeking" [Manuscript submitted for publication], 2026.

[2] M. Vlaswinkel and F. Willems, "Data-Based In-Cylinder Pressure Model with Cyclic Variations for Combustion Control: An RCCI Engine Application", In: *Energies* 17.8, 2024.

[3] M. Vlaswinkel and F. Willems, "Cylinder Pressure Feedback Control for Ideal Thermodynamic Cycle Tracking: Towards Self-Learning Engines", In: IFAC-*PapersOnLine* 56.2, 2023.

Demonstration-Free Diffusion Models for Motion Planning

Edward Sandra, Dries Dirckx, Matthias Van Eysendeyk, Jan Swevers and Wilm Decré

Department of Mechanical Engineering, KU Leuven, 3001 Heverlee, Belgium
Flanders Make@KU Leuven, Belgium
edward.sandra@kuleuven.be

1 Introduction

Robot motion planning is commonly formulated as a trajectory optimization problem that offers strong modeling fidelity but can be computationally demanding online, particularly for long-horizon, high-dimensional systems in cluttered environments. Diffusion models have recently emerged as a promising alternative for motion planning, capable of generating trajectories through iterative denoising. Most existing diffusion-based planners rely on large sets of expert demonstrations [1], which limits generalization due to the finite coverage of the demonstrated data. Others learn a policy through interaction with the environment via reinforcement learning [2], which is often sample-inefficient and requires careful reward shaping.

In this work, we introduce a demonstration-free diffusion-based motion planning framework that directly trains the reverse diffusion process using trajectory optimization objectives, eliminating the need for expert data or environment interaction. Instead of imitating demonstrations, the diffusion model leverages known cost functions, robot kinodynamics, and constraints, with the aim of transforming trajectory optimization into a fast generative model for robot motion planning.

2 Methodology

A diffusion probabilistic model [3] consists of a forward and reverse process. During the forward process, a trajectory from the data distribution $\boldsymbol{\tau}_0 \sim q(\boldsymbol{\tau}_0)$ is gradually transformed during T diffusion steps to white Gaussian noise, described as a Markov chain $q(\boldsymbol{\tau}_t|\boldsymbol{\tau}_0) = \mathcal{N}(\boldsymbol{\tau}_t|\sqrt{\overline{\alpha}_t}\boldsymbol{\tau}_0, (1-\overline{\alpha}_t)\mathbf{I})$ with variance schedule β_t, $\alpha_t = 1-\beta_t$ and $\overline{\alpha}_t = \prod_{i=1}^{t}\alpha_i$. During the reverse process, Gaussian noise is mapped to the original data distribution through the learned, parameterized distribution: $p_{\boldsymbol{\theta}}(\boldsymbol{\tau}_{t-1}|\boldsymbol{\tau}_t) = \mathcal{N}(\boldsymbol{\tau}_{t-1}|\boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{\tau}_t, \mathcal{C}, t), \boldsymbol{\Sigma}_{\boldsymbol{\theta}}(\boldsymbol{\tau}_t, \mathcal{C}, t))$ where $\mathcal{C}$ denotes additional context. Often only the added noise in the forward process is learned, and the mean is computed as $\boldsymbol{\mu}_{\boldsymbol{\theta}}(\boldsymbol{\tau}_t, \mathcal{C}, t) = \frac{1}{\sqrt{\alpha_t}}\left(\boldsymbol{\tau}_t - \frac{\beta_t}{\sqrt{1-\overline{\alpha}_t}}\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\boldsymbol{\tau}_t, \mathcal{C}, t)\right)$. The loss function then becomes:

$$\mathcal{L}_{\text{Data}}(\boldsymbol{\theta}) = \mathbb{E}_{t,\boldsymbol{\epsilon},\boldsymbol{\tau}_0,\mathcal{C}}\left[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\boldsymbol{\tau}_t, \mathcal{C}, t)\|^2\right]. \quad (1)$$

Instead of learning the reverse diffusion process from paired data $(\boldsymbol{\tau}_0, \mathcal{C})$ via (1), we adopt an objective-driven training paradigm, which does not require access to ground-truth trajectories $\boldsymbol{\tau}_0$. The reverse diffusion process is initialized from white noise $\hat{\boldsymbol{\tau}}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and rolled out for $t = T, \ldots, 1$ as

$$\hat{\boldsymbol{\tau}}_{t-1} = \boldsymbol{\mu}_{\boldsymbol{\theta}}(\hat{\boldsymbol{\tau}}_t, \mathcal{C}, t) + \sigma_t \mathbf{z}, \qquad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

yielding $\hat{\boldsymbol{\tau}}_0(\mathcal{C})$. Trajectory quality is assessed through K differentiable objective functions $\{\mathcal{J}_k(\boldsymbol{\tau}, \mathcal{C})\}_{k=1}^{K}$, which define a vector-valued objective. At each reverse diffusion step, we form the corresponding Jacobian $\mathbf{G}(\boldsymbol{\tau}, \mathcal{C}) = \begin{bmatrix}\nabla_{\boldsymbol{\tau}}\mathcal{J}_1(\boldsymbol{\tau}, \mathcal{C}) & \cdots & \nabla_{\boldsymbol{\tau}}\mathcal{J}_K(\boldsymbol{\tau}, \mathcal{C})\end{bmatrix}^\top$.

The training objective is defined as a vector-valued loss

$$\boldsymbol{\mathcal{L}}_{\text{Obj}}(\boldsymbol{\theta}) = \mathbb{E}_{\mathcal{C},\mathbf{z}}\left[\sum_{t=1}^{T} w_t \begin{bmatrix}\mathcal{J}_1(\boldsymbol{\tau}_{t-1}, \mathcal{C}) \\ \vdots \\ \mathcal{J}_K(\boldsymbol{\tau}_{t-1}, \mathcal{C})\end{bmatrix}\right] \in \mathbb{R}^K,$$

where w_t optionally reweights the contribution of each reverse diffusion step. To combine potentially conflicting objective gradients, we employ Jacobian descent with the UPGrad aggregator [4].

We plan to evaluate the proposed method on simulated robotic manipulation tasks and cluttered motion-planning benchmarks, focusing on multimodal coverage, trajectory quality, and efficiency compared to optimization-based planners and data-trained diffusion models.

References

[1] E. Sandra, L. Vanroye, D. Dirckx, R. Cartuyvels, J. Swevers, and W. Decré, "Accelerated multi-modal motion planning using context-conditioned diffusion models," 2025.

[2] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," 2024.

[3] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," 2020.

[4] P. Quinton and V. Rey, "Jacobian descent for multi-objective optimization," 2025.

Acknowledgment This work was supported by KU Leuven's Special Research Fund (BOF), grant STG/25/004, and by Flanders Make through SBO project LearnOpTra.

Randomized spectral node embeddings for graph kernel approximation

Valentin de Bassompierre
INMA/ICTEAM, UCLouvain, Belgium
valentin.debassompierre@uclouvain.be

Jean-Charles Delvenne
INMA/ICTEAM, UCLouvain, Belgium
jean-charles.delvenne@uclouvain.be

Laurent Jacques
INMA/ICTEAM, UCLouvain, Belgium
laurent.jacques@uclouvain.be

Node embeddings provide a powerful representation of networks by mapping graph vertices into low-dimensional Euclidean spaces while preserving structural information. Such embeddings are now central to a wide range of applications, including node classification, link prediction, community detection, and signal reconstruction. A common objective is to design embeddings whose inner products reflect meaningful notions of node similarity derived from the underlying graph structure.

Signal processing on graphs and spectral graph theory offer a principled framework for constructing such embeddings. By exploiting the eigenspectrum of the graph Laplacian, spectral methods define embeddings through filtering operations that encode multi-scale structural information [1, 2]. These approaches connect naturally to classical tools in harmonic analysis, while remaining well adapted to irregular network domains.

In this work, we focus on the problem of computing node embeddings with explicit kernel approximation guarantees. Graph kernels—defined as functions of the graph Laplacian—provide a flexible and theoretically grounded way to measure node similarity, but their direct computation is often prohibitively expensive for large networks. Inspired by random feature methods for kernel approximation in Euclidean spaces, we introduce randomized spectral embeddings on graphs whose dot products yield low-rank approximations of Laplacian-based graph kernels. In practice, we implement this using the efficient graph wavelet formalism applied to a series of random signals.

We present theoretical and empirical results demonstrating that our embeddings achieve more accurate kernel approximations than existing methods [3], particularly when the target kernel is localized in the spectral domain. Moreover, we highlight connections between our approach and randomized SVD, a well-established algorithm in numerical linear algebra. These findings underscore the advantages of randomized spectral constructions for scalable and principled graph representation learning.

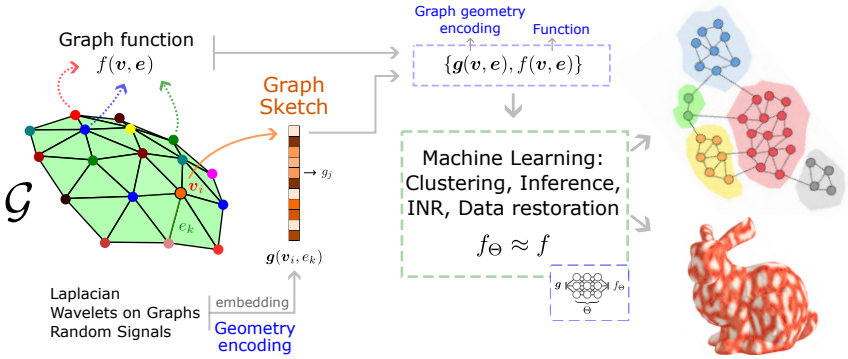

Figure 1: Schematic view of our research axes. We propose a new random graph embedding, or *graph sketch*, to encode the geometry of networks. These graph sketches provide low-rank approximations to any graph kernel, in a computationally tractable way. The overall aim is to feed machine learning techniques and enable them to integrate the properties of the graph.

References

[1] N. Tremblay and P. Borgnat. "Graph Wavelets for Multiscale Community Mining," IEEE Transactions on Signal Processing 62.20, 2014.

[2] N. Tremblay, G. Puy, R. Gribonval, and P. Vandergheynst. "Compressive Spectral Clustering," International conference on machine learning. PMLR, 2016.

[3] I. Reid, K. Choromanski, E. Berger, and A. Weller. "General Graph Random Features," 2023.

Neural Network Approximations of Discrete-Time CBF Controllers with Safety Guarantees

Patricia Pauli
Control Systems Technology, Department of Mechanical Engineering, Eindhoven, The Netherlands
Email: p.d.pauli@tue.nl

1 Introduction

Safety-critical systems in domains such as automotive engineering, robotics, and healthcare require controllers with rigorous, formally provable safety guarantees to prevent failures that could result in significant harm or damage. Control Barrier Functions (CBFs) provide a systematic framework for enforcing safety by rendering the zero-superlevel set of a CBF forward invariant while ensuring its inclusion in the safe set [1]. CBF-based controllers are well understood in continuous time, where the associated safe control law reduces to a quadratic program (QP) that can be solved efficiently online. However, the discrete-time setting lacks this favorable structure. As a result, discrete-time CBF-based controllers often involve more complex optimization problems with higher computational burden, limiting their applicability in real-time systems. To address this challenge, we propose to approximate discrete-time CBF-based controllers using neural networks (NNs) while retaining formal safety guarantees.

2 Problem Statement

Consider the nonlinear discrete-time system

$$x_{k+1} = f(x_k, u_k),$$

with state vector $x_k \in \mathbb{R}^n$, control input vector $u_k \in \mathbb{U} \subseteq \mathbb{R}^m$, both at discrete time $k \in \mathbb{N}_0$ and mapping $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$. Here, $\mathbb{U}$ denotes the control admissible set. The goal is to design a neural controller $\pi_{\text{NN}} : \mathbb{R}^n \to \mathbb{R}^m$ that is fast to evaluate and guarantees safety for the closed-loop system

$$x_{k+1} = f(x_k, \pi_{\text{NN}}(x_k)), \tag{1}$$

where safety means $x_k \in \mathscr{S}$ for all $k \in \mathbb{N}_0$ with safe set $\mathscr{S} = \{x \in \mathbb{R}^n \mid s(x) \geq 0\}$, where $s : \mathbb{R}^n \to \mathbb{R}$ is a given mapping.

3 Method

To find a safe neural controller, we first design an optimization-based robust controller $\pi : \mathbb{R}^n \to \mathbb{R}^m$ for the discrete-time system

$$x_{k+1} = f(x_k, u_k) + w_k, \tag{2}$$

with disturbance vector $w_k \in \mathbb{W} := \{w \in \mathbb{R}^n \mid \|w\| \leq w_{\max}\}$, where $\|\cdot\|$ denotes the Euclidean norm, following [2]. To guarantee robust safety for system (2), one can construct a robust controlled invariant set $\mathscr{C}$ within the safe set $\mathscr{S}$. To do so, we synthesize a safe controller for the perturbed system (2) using a robust discrete-time CBF (R-DTCBF) $h : \mathbb{R}^n \to \mathbb{R}$ with zero-superlevel set $\mathscr{C} := \{x \in \mathbb{R}^n \mid h(x) \geq 0\}$ that satisfies

$$h(f(x,u) + w) - h(x) \geq \gamma(h(x)) \quad \forall w \in \mathbb{W} \tag{3}$$

with $\gamma \in \mathscr{K}_\infty^{\leq \text{id}}$ using [2, Lemma 1]. We denote by $\mathscr{K}_\infty^{\leq \text{id}}$ the set of class-$\mathscr{K}_\infty$ functions satisfying $\gamma(r) \leq r$ for all $r \in \mathbb{R}_{\geq 0}$.

We then sample trajectories of the control law π to generate training data to approximate it with a neural controller π_{NN}. After training, we verify a uniform approximation error bound over the controlled invariant set $\mathscr{C}$: $\|e(x)\| \leq e_{\max}$ for all $x \in \mathscr{C}$.

Setting $w = f(x, \pi(x) + e(x)) - f(x, \pi(x))$, we obtain

$$f(x, \pi_{\text{NN}}(x)) = f(x, \pi(x) + e(x)) = f(x, \pi(x)) + w,$$

to describe the closed-loop dynamics (1) in terms of (2). Using this connection, we establish a condition on the approximation error e under which the neural network controller inherits the safety guarantees of the underlying R-DTCBF-based controller.

Theorem 1 *Let f be Lipschitz continuous in u on $\mathscr{C}$ and let h be Lipschitz continuous in x on $\mathscr{C}$ with given Lipschitz constants L_u, L_h, respectively. Suppose for a given pair (h, γ), there exists a control input $u \in \mathbb{U}$ such that*

$$h(f(x,u)) - h(x) + \gamma(h(x)) \geq L_h w_{max}$$

holds for each $x \in \mathscr{C}$, and assume that $\mathscr{C} \subseteq \mathscr{S}$. Let the approximation error satisfy $\|e(x)\| \leq e_{max}$ for all $x \in \mathscr{C}$. If $w_{max} \geq L_u e_{max}$, then the neural controller π_{NN} guarantees safety for the closed-loop system (1).

References

[1] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *53rd IEEE conference on decision and control*. IEEE, 2014, pp. 6271–6278.

[2] E. Shakhesi, A. Katriniok, and W. Heemels, "Counterexample-guided synthesis of robust discrete-time control barrier functions," *IEEE Control Systems Letters*, 2025.

Gradient-Based Policy Iteration for Linear Quadratic Regulation with Finite-Time Guarantees

Vijayanand Digge and Gianluca Bianchin
ICTEAM institute, UCLouvain, {vijayanand.digge, gianluca.bianchin}@uclouvain.be

1 Introduction

Reinforcement learning (RL) provides a general framework for sequential decision making in uncertain environments. In continuous-state control, RL methods have attracted significant attention, with linear quadratic regulation (LQR) serving as a canonical benchmark that enables rigorous theoretical analysis. Classical policy iteration schemes for LQR rely on least-squares policy evaluation, which requires repeated matrix inversions and becomes computationally demanding in higher dimensions. This motivates gradient-based policy iteration methods that improve scalability while retaining data efficiency and admitting explicit finite-time performance guarantees.

In this work, we propose and analyze an online implementation of policy iteration for the LQR problem. Instead of explicitly solving the least-squares policy evaluation problem, we perform a finite number of gradient descent iterations on the quadratic Bellman error and update the policy via greedy minimization of the resulting approximate Q-function. Building on the finite-time policy iteration analysis for LQR [1], we provide explicit finite-time bounds on both the number of policy improvements and the number of gradient steps required to obtain a controller whose LQR cost is within a prescribed tolerance of the optimal cost.

2 Problem Formulation

Consider the linear discrete-time system

$$x(k+1) = Ax(k) + Bu(k) \tag{1}$$

where, $x(k) \in \mathbb{R}^n$ is the state, $u(k) \in \mathbb{R}^m$ is the input to the system. The state feedback controller is given by $u(k) = -Kx(k)$. We consider LQR control problem described by the minimization of the infinite-horizon cost:

$$J(x(k),u(k)) = \sum_{k=0}^{\infty} x(k)^\top Qx(k) + u(k)^\top Ru(k) \tag{2}$$

where $Q, R \succ 0$.

Problem: Compute an optimal LQR feedback gain K^* via policy iteration, and characterize the finite-time and sample complexity of the resulting algorithm when policy evaluation is performed using gradient-based methods.

3 Methodology and Results

The proposed method follows a policy iteration structure [1] with an inexact policy evaluation step. Given a stabilizing controller K_t at policy iteration t, policy evaluation is performed by approximating the corresponding state–action value function using a linear parametrization of the Q-function. The associated Bellman error induces a convex quadratic loss whose minimizer coincides with the least-squares policy evaluation solution. The approximate Q-function obtained after T gradient descent steps is then used in a greedy policy improvement step, yielding an updated stabilizing controller K_{t+1}. After N policy improvement iterations, the algorithm returns the controller K_N.

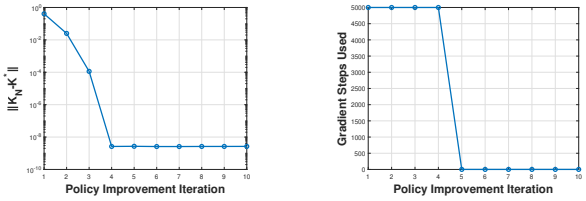

(a) Policy convergence (b) Adaptive gradient steps

Figure 1: This figure illustrates the simulation of gradient-based policy iteration. Figure 1a depicts convergence of the estimated policy K_N toward the optimal LQR gain K^*, illustrated by the evolution of the error norm $\|K_N - K^*\|$ over policy iteration steps N. Figure 1b represents adaptive gradient steps over policy iteration steps N.

Main Result: Let K_0 be an initial stabilizing controller. For any accuracy $\delta > 0$, there exist integers T and N, scaling logarithmically with $1/\delta$, such that if at least T gradient descent steps are performed at each policy evaluation, then the controller K_N obtained after N policy improvement iterations satisfies

$$\|K_N - K^*\| \leq \delta. \tag{3}$$

The bounds depend explicitly on the cost matrices, the initial controller, and the condition number of the sampled data covariance.

Simulation Results: We consider the example described in [2]. Figure 1a illustrates the convergence behavior of the policy across successive iterations. As observed, the policy quickly converges to the optimal LQR gain within a few iterations. Figure 1b shows the number of gradient descent steps required per policy iteration.

Future work will consider measurement noise in the data, leading to high-probability finite-sample guarantees.

References

[1] K. Krauth, S. Tu, and B. Recht, "Finite-time analysis of approximate policy iteration for the linear quadratic regulator," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[2] D. Lee, "Convergence of dynamic programming on the semidefinite cone for discrete-time infinite-horizon lqr," *IEEE Transactions on Automatic Control*, vol. 67, no. 10, pp. 5661–5668, 2022.

Data-Driven Multi-Step Prediction in Fed-Batch CHO Cell Cultures

Ali El Hadi Noura[a,b] ali.noura@ucb.com
Bassem Ben Yahia[b] Bassem.BenYahia@ucb.com
Laurent Dewasme[a] laurent.dewasme@umons.ac.be
Alain Vande Wouwer[a] alain.vandewouwer@umons.ac.be

[a]Systems, Estimation, Control and Optimization (SECO), University of Mons, Belgium
[b]UCB Pharma S.A., Avenue de l'Industrie, B-1420, Braine l'Alleud, Belgium

1 Introduction

Mathematical models are essential for the analysis, monitoring, and control of bioprocesses, enabling process understanding, optimization, and decision making. In fed-batch CHO cell cultures, modeling is challenging due to strong nonlinearities, time-varying dynamics, and significant process variability.

Mechanistic models based on mass balances and kinetic expressions are widely used due to their interpretability and strong link to process knowledge, but their development and maintenance often require substantial process knowledge and iterative model refinement, while their predictive performance can still be limited by structural mismatch, parameter uncertainty, and identifiability issues [1]. Data-driven approaches provide an alternative by learning process dynamics directly from data, yet face challenges related to limited datasets, heterogeneous sampling, non-stationarity, and robustness outside the training domain [2].

2 Methodology

A data-driven modeling framework based on recurrent neural networks is investigated for multi-step state prediction and soft sensing in fed-batch CHO cell cultures. Long Short-Term Memory (LSTM) networks are considered due to their ability to capture nonlinear and long-range temporal dependencies in time-series data.

Several forecasting strategies are evaluated, including one-step-ahead recursive prediction and direct multi-output formulations. Particular emphasis is placed on the Multiple-Input Multiple-Output (MIMO) strategy, which predicts a future horizon simultaneously and mitigates error accumulation in long-term forecasting [2]. A comparative study is conducted to identify the most suitable forecasting strategy under different operating scenarios.

Feature selection and feature engineering are investigated to improve model informativeness and robustness. Raw process measurements are transformed and combined to obtain more informative input features, while reducing redundancy and sensitivity to noise. Hyperparameter optimization is performed to balance model complexity and generalization performance.

Prior to training, the available datasets are analyzed across the experimental design space to ensure data consistency, coverage, and relevance. This step aims to limit overfitting and improve generalization across operating conditions. In addition, a calibrated mechanistic model is used not only as a performance baseline, but also explored as a tool for data augmentation, enabling the generation of additional training scenarios when historical batch data are insufficient.

3 Results and Discussion

The results show that the performance of data-driven models for multi-step prediction in fed-batch bioprocesses strongly depends on the selected forecasting strategy, feature representation, and dataset preparation. Direct multi-output formulations based on the MIMO strategy provide more stable long-horizon predictions than recursive approaches, particularly in the presence of sparse and heterogeneous data.

These findings underline that data-driven models can be effective tools for bioprocess monitoring when their limitations are explicitly addressed, and that methodological choices play a critical role in achieving robust and generalizable performance.

References

[1] B. Ben Yahia, L. Malphettes, and E. Heinzle, "Predictive macroscopic modeling of cell growth, metabolism and monoclonal antibody production: Case study of a CHO fed-batch production," *Metabolic Engineering*, vol. 66, pp. 204–216, 2021.

[2] S.-Y. Park, S.-J. Kim, C.-H. Park, J. Kim, and D.-Y. Lee, "Data-driven prediction models for forecasting multi-step ahead profiles of mammalian cell culture toward bioprocess digital twins," *Biotechnology and Bioengineering*, vol. 120, pp. 2494–2508, 2023.

Distribution-Aware Sharpness via Derandomized PAC-Bayes

Domagoj Herceg
Control Systems Technology Section
Eindhoven University of Technology
Email: d.herceg@tue.nl

1 Introduction and Motivation

Modern deep networks are routinely overparameterized, i.e. they have many more parameters than data. Consequently, there are typically multiple *global* optima, but not all of them generalize well on unseen data. Sharpness, a measure of the sensitivity of the loss landscape to parameter perturbations, is a powerful heuristic for selecting solutions that generalize well in deep learning. However, despite its empirical success, sharpness often lacks a clear theoretical motivation. In contrast, PAC-Bayes generalization theory [1] provides some of the tightest known generalization guarantees for over-parameterized models, but the role of sharpness in these bounds is implicit and typically interpreted post-hoc. In this work, we aim to reconcile these two perspectives. By derandomizing the PAC-Bayesian expectation of the empirical loss, we show that a novel notion of sharpness emerges naturally. The resulting sharpness term relates to several existing sharpness definitions as special cases and, importantly, becomes directly usable for optimization.

2 Background and Related Work

We briefly cover two concepts central to our work: sharpness and PAC-Bayes generalization theory. Sharpness is *often* defined (see [2]) as the maximum local increase of the loss $L(w)$ under bounded parameter perturbations,

$$\text{sharpness}(w) = \max_{\|u\|\leq\rho} L(w+u) - L(w), \tag{1}$$

where $w \in \mathbb{R}^d$ is the parameter vector, and $u \in \mathbb{R}^d$ is a local perturbation. A prominent example of the operational use of the above is Sharpness-Aware Minimization (SAM) [2]. Informally, SAM seeks parameters that are insensitive to local perturbations leading to

$$\min_{w\in\mathbb{R}^d} \max_{\|u\|\leq\rho} L(w+u) + \lambda\|w\|^2, \tag{2}$$

where the last term is the regularizer. On the other hand, a typical PAC-Bayes bound takes the form

$$\mathbb{E}_{w\sim Q}[L(w)] \;\leq\; \mathbb{E}_{w\sim Q}[\widehat{L}_S(w)] \;+\; \text{cmplx}(Q,P,n,\delta) \tag{3}$$

where $\widehat{L}_S(w)$ denotes the empirical loss on a dataset S, $L(w)$ the population loss, P a data-independent prior, and Q a posterior. The complexity term penalizes the "distance" between Q and P via the Kullback-Leibler divergence and is inversely proportional to the number of samples n. If we further assume the data are drawn i.i.d. from a (unknown) distribution $\mathscr{D}$, then the statement holds with confidence $1-\delta$.

3 De-randomization and Emergent Sharpness

PAC-Bayes bounds are traditionally formulated for randomized predictors and expectations. Using a local matrix smoothness assumption on the empirical loss, we show that the randomized empirical loss term $\mathbb{E}_{w\sim Q}[\widehat{L}_S(w)]$ can be upper-bounded deterministically. Specifically, we assume the existence of a locally valid Hessian proxy $M \succeq 0$ such that, for all w in a neighborhood of a reference point $\bar{w}$, with $\Delta w = w - \bar{w}$, the following holds:

$$\widehat{L}_S(w) \leq \widehat{L}_S(\bar{w}) + \langle \nabla\widehat{L}_S(\bar{w}), \Delta w\rangle + \frac{1}{2}\Delta w^\top M \Delta w. \tag{4}$$

Under the standard Gaussian restrictions of Q and P, we get:

$$\mathbb{E}_{w\sim Q}[\widehat{L}_S(w)] \leq \widehat{L}_S(\bar{w}) + \frac{1}{2}\text{tr}(M\Sigma), \tag{5}$$

where Σ denotes the posterior covariance. We interpret the second term as a sharpness measure given by a covariance-curvature interaction. Importantly, replacing the randomized expectation with this deterministic upper bound preserves the validity of the PAC-Bayes guarantee. The resulting notion of sharpness includes SAM and related approaches as special cases while avoiding known issues such as sensitivity to reparameterization. Moreover, the sharpness term is anisotropic and directly differentiable, making it suitable for optimization. Intuitively, sharp directions of the loss landscape are penalized only if the posterior assigns significant probability mass to them, reflecting the fact that generalization depends on predictions rather than worst-case perturbations. We conclude with numerical examples, starting from a simple toy problem and progressing to standard benchmarks on CIFAR-10 and CIFAR-100.

References

[1] O. Catoni. *PAC-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*. Institute of Mathematical Statistics, 2007.

[2] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021.

System identification of a control-oriented core model for STEP

Lizan Kivits[†‡*], Lennard Ceelen[†‡], Paulo Figueiredo[†‡], Matthijs van Berkel[†‡], and Marco de Baar[†‡]
[†]DIFFER - Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands
[‡]Eindhoven University of Technology, Eindhoven, the Netherlands
[*]Email: e.m.m.kivits@differ.nl

1 Nuclear fusion

Nuclear fusion is a conceivable future energy production method that is considered to be clean, safe, and efficient. In nuclear fusion reactors, atomic nuclei (typically of the hydrogen isotopes deuterium and tritium) are merged together, for which very high temperatures are required. The heating leads to ionisation of the hydrogen gas into plasma: a mixture of freely moving atomic nuclei and electrons.

The plasma interacts with reactor components and contains many internal processes. Due to the complexity of the plasma's behaviour, it needs to be controlled to guarantee safe operation and maximum performance of the reactor [1]. Next to the plasma control system, an advanced supervisory control system is required to handle reactor constraints, hardware failure, and plasma limits [2]. To design these controllers, a model of the plasma is needed.

2 Control-oriented core model

The overall objective of this research is to construct a plasma model for the STEP tokamak. In this particular work, we aim at a control-oriented model of the core, which can be used for controller and observer design and for prediction of the profile evolution. This model is identified from simulated data that is generated with JINTRAC: a system of interfaced Tokamak-physics codes [3].

First, we have to decide which profiles and actuator signals are relevant for this core model. During (nominal) flat-top operation of the tokamak, the plasma in the core acts like a self-organizing system, as in Figure 1. Based on this scheme, the selected actuators are the pellet injector (for fuelling) and the Electron Cyclotron Resonance Heating (ECRH) (for heating and current drive). The selected profiles are the electron temperature, electron density, plasma current, and the fusion power.

3 Simulation and identification

Second, the STEP tokamak is simulated with JINTRAC and the data of the selected actuations and profiles is obtained. We will employ the Local Polynomial Method (LPM) [4, 5] to identify a frequency-domain model of the core. We select the LPM, because transients can be corrected for. From the identified core-model, the relations between the plasma profiles can be extracted, which give further insight into the profiles that are relevant for controller and observer design.

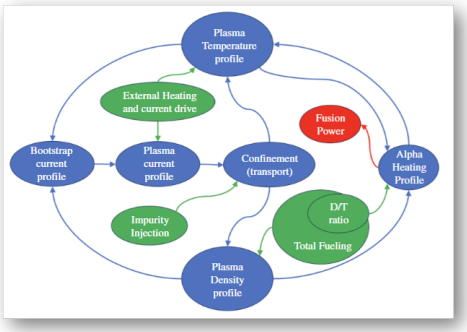

Figure 1: Plasma self-organizing system [1].

4 Future steps

In an actual tokamak, it will not be possible to directly measure the plasma profiles; they will be detected by a diagnostic system instead. To make the mapping from measured diagnostic signals to the plasma profiles, a model of the diagnostic system has to be available. Similar to the core model, the diagnostic model can be identified from JINTRAC simulation data. This will be one of the next steps in this research.

Acknowledgements

DIFFER is part of the institutes organization of NWO. This work has been funded by the Spherical Tokamak for Energy Production (STEP), a UKAEA programme to design and build a prototype fusion energy plant and a path to commercial fusion.

References

[1] M. Lennholm, et al., Phil. Trans. R. Soc. A. 382 (2024) 20230403.

[2] P. C. de Vries, et al., Fusion Eng. Des. 204 (2024) 114464.

[3] M. Romanelli, et al., Plasma Fusion Res. 9 (2014) 3403023.

[4] J. Schoukens, et al., Eur. J. Control 15 (2009) 260–274.

[5] M. van Berkel, et al., Plasma Phys. Control. Fusion 62 (2020) 094001.

Model augmentation for thermo-mechanical hysteretic dynamics [1]

Samarth Toolhally[1], Ozan Celik[2], Daniel Deenen[2], Maarten Schoukens[1], and Siep Weiland[1]

[1]Control Systems Group, Electrical Engineering, Eindhoven University of Technology
[2]ASML, Veldhoven, The Netherlands
Emails: s.toolhally@tue.nl

1 Introduction

Accurate real-time prediction of thermo-mechanical wafer deformation is crucial for minimizing overlay errors in advanced lithography. The wafer, supported by burls, forms a high-dimensional system with nonlinear hysteretic behavior. Real-time predictions require models that are both precise and computationally efficient. This work proposes model augmentation, combining a physics-based baseline model with system data to create fast, interpretable surrogate models. Preliminary results from a scaled-down model demonstrate the feasibility of the approach.

2 System description

The dynamics of the true system is of the form:

$$x_{k+1} = f(x_k, u_k), \quad y_k = h(x_k, u_k). \tag{1}$$

$x_k \in \mathbb{R}^{n_x+n_s}$ is temperature and slip states for each burl, with n_x burls and n_s slip states. $u_k \in \mathbb{R}^{n_u}$ is the heat input, and $y_k \in \mathbb{R}^{n_y}$ is the X-Y displacement of the wafer over each burl. The first-principles linear baseline model, which does not capture the hysteretic behaviour is given by:

$$\tilde{x}_{k+1} = A\tilde{x}_k + Bu_k, \quad \tilde{y}_k = C\tilde{x}_k + Du_k, \tag{2}$$

where $\tilde{x}_k \in \mathbb{R}^{n_{\tilde{x}_k}}$ are the temperature states, $\tilde{y}_k \in \mathbb{R}^{n_{\tilde{y}_k}}$ are the burl displacements for the linear model.

3 Methodology

An LFR-based model augmentation approach, as developed in [1], is used to identify missing system dynamics, utilizing an ANN parameterisation for flexible model structure accommodation.

$$\begin{bmatrix}\hat{x}_{k+1} & \hat{y}_k & z_{1,k} & z_{2,k}\end{bmatrix}^\top = \boldsymbol{S}\begin{bmatrix}\hat{x}_k & u_k & w_{1,k} & w_{2,k}\end{bmatrix}^\top \tag{3a}$$

$$w_{1,k} = \phi_{\text{base}}(z_{1,k}), \quad w_{2,k} = \phi_{\text{aug}}(z_{2,k}) \tag{3b}$$

where $\hat{y}_k$ is the predicted output, $\boldsymbol{S}$ is a selection matrix and $\hat{x}_k = [\tilde{x}_k\ \bar{x}_k]^\top$, where $\bar{x}_k$ are the additional states for the dynamic learning component, resulting in two possible structure as in Fig.1. The Learning component uses a neural network to capture unmodeled dynamics and non-linearities from the system data using *Sub-space Encoder Network* (SUBNET) method [2]. It employs a truncated objective function that supports batch optimization, commonly used in machine learning. The objective is formed by dividing the data $D_N = (y_k, u_k)_{k=1}^N$ into segments of length T.

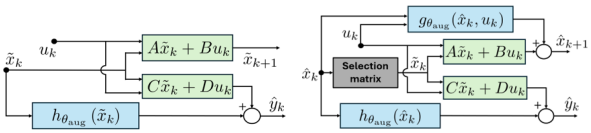

(a) Static augmentation (b) Dynamic augmentation

Figure 1: Augmented structures

4 Results

The system is excited using a multisine signal with 10 frequencies, a base frequency of 0.1 Hz, and random phase shifts. The signal is sampled at 10 Hz for a total of 100 seconds. A feedforward neural network with 2 hidden layers of 128 neurons each was trained for 350 epochs using data from 50 experiments for training, 9 for validation, and 1 for testing. Dynamic augmentation with 10 additional states achieved the best performance, with an NRMSE of 0.1120, compared to static augmentation with 0.1616 and the baseline model at 0.5996.

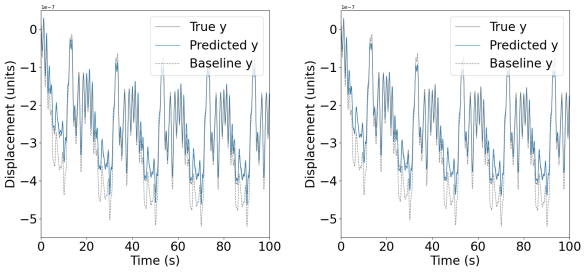

(a) Static augmentation (b) Dynamic augmentation

Figure 2: Comparison of the augmented model and the baseline model with the true system

References

[1] J. H. Hoekstra, C. Verhoek, R. Tóth, and M. Schoukens, "Learning-based model augmentation with lfrs," *European Journal of Control*, vol. 86, p. 101304, Nov. 2025.

[2] G. I. Beintema, M. Schoukens, and R. Tóth, "Deep subspace encoders for nonlinear system identification," *Automatica*, vol. 156, 10 2023.

[1]This project is co-financed by Holland High Tech, top sector High-Tech Systems and Materials, with a PPP innovation subsidy for public-private partnerships for research and development.

Efficient Learning of Affine and Rationally Dependent LPV Models With Linear Fractional Representation

Roel Drenth[1], Jan H. Hoekstra[1], Maarten Schoukens[1], and Roland Tóth[1,2]

[1] Control Systems Group, Eindhoven University of Technology, The Netherlands

[2] Systems and Control Laboratory, Institute for Computer Science and Control, Budapest, Hungary

r.drenth@tue.nl, j.h.hoekstra@tue.nl, m.schoukens@tue.nl, r.toth@tue.nl

1 Introduction

Identifying control-friendly models of nonlinear systems remains one of the major challenges at the intersection of system identification and control. The *Linear Parameter-Varying* (LPV) framework offers a promising solution, but existing identification methods often rely on model structures with affine scheduling dependency. Instead, this work proposes the use of LPV models with *Linear Fractional Representation* (LFR) admitting a rational scheduling-dependency, capable of modelling complex nonlinear systems with less scheduling variables compared to affine models. This can reduce conservatism (called overbounding) in controller synthesis, for which the rationally dependent LPV-LFR model structure is still suited. This work introduces a parameterization of rational LPV-LFR models that ensures well-posedness in a constraint free form. Furthermore, an identification approach is proposed that by joint-estimation of such LPV models, the scheduling map and a noise model, using only input-output data, is capable of modelling nonlinear systems under complex noise scenarios.

2 Method

We aim to identify discrete-time models of plant and noise processes using the identification setting proposed in [1], minimizing the error of the one-step ahead prediction:

$$\hat{y}(k|k-1) = G_\theta u(k) + (I - H_\theta^{-1})(y(k) - G_\theta u(k)), \quad (1)$$

where $k \in \mathbb{Z}$ is the discrete time, $u(k) \in \mathbb{R}^{n_\mathrm{u}}$ is the input, $\hat{y}(k) \in \mathbb{R}^{n_\mathrm{y}}$ is the model output, G_θ is the plant model and H_θ^{-1} is the model of the inverse noise process. The plant is parametrized in LFR form:

$$\begin{aligned} x(k+1) &= A_x x(k) + B_w w(k) + B_u u(k), && \text{(2a)}\\ z(k) &= C_z x(k) + D_{zw} w(k) + D_{zu} u(k), && \text{(2b)}\\ y(k) &= C_y x(k) + D_{yw} w(k) + D_{yu} u(k), && \text{(2c)}\\ w(k) &= \Delta(p(k)) z(k), && \text{(2d)} \end{aligned}$$

where $x(k) \in \mathbb{R}^{n_\mathrm{x}}$ is the state, $w(k), z(k) \in \mathbb{R}^{n_\mathrm{w}}$ are the latent variables, $A_\mathrm{x}, \dots D_{yu}$ are real matrices of appropriate dimensions, and $p(k) \in \mathbb{P} \subseteq \mathbb{R}^{n_\mathrm{p}}$ is the scheduling variable. The matrix function $\Delta(p(k)) : \mathbb{P} \to \mathbb{R}^{n_\mathrm{w} \times n_\mathrm{w}}$ is a diagonal matrix with repeated elements of $p(k)$. The inverse noise model in (1) is parameterized with a similar LFR structure.

This work is funded by the European Union (ERC, COMPLETE, 101075836), by the Air Force Office of Scientific Research under award number FA8655-23-1-7061, and the Mathworks Inc.

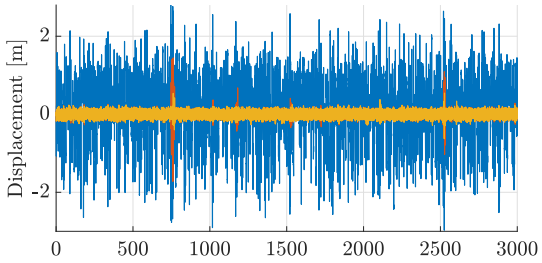

Figure 1: True output (—), affine LPV-LFR-BJ ($n_\mathrm{p}=1$) error (—) and rational LPV-LFR-BJ ($n_\mathrm{p}=1$) error (—) on the test data of the NL-MSD.

To represent nonlinear behaviour we identify a scheduling map $\phi : \mathbb{R}^{n_\mathrm{x}} \times \mathbb{R}^{n_\mathrm{u}} \to \mathbb{P}$ which maps inputs and states of the plant model to a scheduling variable shared by the plant and noise processes. The scheduling map is parametrized as a feedforward neural network.

Free parameterization of D_{zw} in (2) can introduce algebraic loops between $w(k)$ and $z(k)$, making the model ill-posed. To avoid this, D_{zw} is typically forced to be a zero-matrix, resulting in an affine scheduling dependency. Instead, we propose a well-posed parameterization inspired by [2], which admits rational scheduling dependency. To achieve this, we construct D_{zw} by a transformation of free variables, using the matrix exponential. This can, under mild conditions, guarantee well-posedness of the LFR.

3 Example

We consider a *Nonlinear Mass-Spring-Damper* (NL-MSD) system with nonlinear spring and nonlinear input-gain. In this example, output data is contaminated with gaussian noise, filtered by a nonlinear noise process. We show that rational-dependency-based model structures using our well-posed parameterization result in more accurate model estimates than affine-dependency-based models when the same scheduling order ($n_\mathrm{p} = 1$) is used. This is shown in Figure 1. Moreover, by using a *Box-Jenkins* (BJ) noise structure, the surrogate models are capable of accurately capturing the plant and noise processes of the data-generating system.

References

[1] Bemporad, A. and Tóth, R. (2025). Efficient identification of linear, parameter-varying, and nonlinear systems with noise models. arXiv preprint arXiv:2504.11982.

[2] Revay, M. et al. (2020). Lipschitz Bounded Equilibrium Networks. arXiv preprint arXiv:2010.01732.

Screening of impurities in nuclear fusion plasmas

M.P. de Beij[1,2,†], M. van Berkel[1,2], A. Das[2], K.J. Brunner[3]

[1] DIFFER - Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands [1] – [†] m.p.debeij@differ.nl

[2] Eindhoven University of Technology, Control Systems, Eindhoven, the Netherlands

[3] Max Planck Institute for Plasma Physics, IPP Greifswald, Germany

1 Impurities

The nuclear fusion reaction between hydrogen isotopes is a promising future energy source with potentially unlimited fuel. Although fusion reactors would be safe, conditions inside the reactors required for energy production are extreme. The fuel, a plasma is contained by a magnetic field but still interacts with the reactor wall. In the fusion reactor Wendelstein 7-X (W7-X), carbon particles from the wall ionize, enter the plasma and collect in the plasma. This dilutes the plasma, radiates away energy and significantly hampers the performance. This project focuses on controlling impurity seeding at W-7X by influencing the plasma temperature and density [2].

The characteristics of fusion reactors allows reducing the spatial domain to a single dimension; spanning from the magnetic axis of the plasma to the plasma edge. On this domain the spatially varying parameters are considered *profiles*. Contemporary diagnostic systems cannot enable effective profile control, requiring observers for real-time state estimation. The effectiveness of the observer and control strategy highly depend on the accuracy of the underlying transport model. Between fusion reactors, dynamics could vary significantly, requiring different approaches. For this reason, several models will be investigated, which will be compared on the basis of state estimation accuracy and controllability of the plasma parameters of W-7X.

2 Transport model identification

Identification of the control-oriented model begins with a hyperbolic cylindrical diffusion-drift-transport PDE (1) which is considered based on current knowledge;

$$\frac{\partial z}{\partial t} = \left(\frac{\partial V}{\partial \rho}\right)^{-1} \frac{\partial}{\partial \rho}\left(\frac{\partial V}{\partial \rho}\left(\frac{d\rho}{dr}\right)^2 \kappa \frac{\partial z}{\partial \rho} + \frac{\partial V}{\partial \rho}\frac{d\rho}{dr} v z\right) + Pu, \tag{1}$$

with heat or density z, magnetic coordinate ρ, plasma radius r and volume V; first- and second-order spatial derivative transport parameters κ and v, respectively; source term P and input u. Variations of this PDE will be fitted to experimental data. This includes nonlinear, linear and time-varying transport parameters.

[1] DIFFER is part of the institutes organisation of NWO.

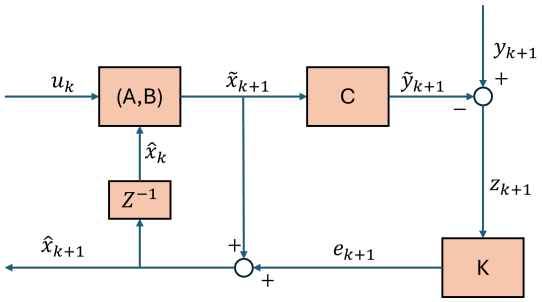

Figure 1: Diagram of a discrete-time KF-based observer with input u, state x, measurement y, residual z, dynamic model (A,B), synthetic diagnostic model C and observer gain K.

3 Observing the plasma

Both the identified model together with the available diagnostic signals are used for real-time profiles estimation for plasma temperature and density. A linear Kalman filter based observer has been developed for this application based on the work in [1].

The observer relies on synthetic models of the diagnostic systems capable of measuring plasma temperature and density and tuning of covariances. An effort will be made to simplify tuning the observer by reducing the number of hyperparameters related to the covariances [3].

4 Control

The control strategy is highly dependent on the actuator used for controlling the plasma profiles. This includes PID-controllers or more advanced MPC, dor instance. Other control systems will be active during operation of the fusion reactor as well, focusing on other aspects such as plasma disruptions. This means future control systems will include some form of supervisory control, incorporating the impurity screening controller. This is taken into account as well.

References

[1] Blanken, T.C., et al., Fusion Eng. Des. **126**, (2018)

[2] Wade, M.R., et al., Phys. Rev. Lett. **84**, 282 (2000)

[3] Beelen, H.P.G.J., et al., IEEE Trans. on Ctrl. Sys. Tech. **29**(3), 1087

Safe Extrapolation in Physics–Data Augmented Models*

Abdullah Habboush, Fahim Shakib, Tom Oomen, and Nathan van de Wouw
Eindhoven University of Technology, The Netherlands
Emails: {a.habboush, m.f.shakib, t.a.e.oomen, n.v.d.wouw}@tue.nl

1 Introduction

Hybrid models combine physics-based models with data-driven components to achieve high accuracy on a given dataset. However, their performance can degrade when extrapolating beyond the training data, often producing predictions inconsistent with the underlying physics and potentially compromising stability. In this work, we address this challenge by enforcing consistency with the physics-based model away from the dataset while maintaining accuracy on it. We propose a model modification technique that employs a transition mapping to ensure stability-preserving extrapolation based on the known stability properties of the underlying physics-based model. The method is applicable to any nonlinear hybrid model, offering a practical and scalable modification particularly suited for high-order systems.

2 Problem Statement

We consider a data-generating system of the form

$$\dot{x}^*(t) = f^*(x^*(t)), \tag{1}$$

with a measurable state $x^*(t) \in \mathbb{R}^n$, and an unknown vector field $f^*(\cdot) : \mathbb{R}^n \to \mathbb{R}^n$. A hybrid model of (1) can be constructed by integrating a physics-based model with a data-driven component as

$$\dot{x}_h(t) = f_p(x_h(t)) + f_d(x_h(t)) \tag{2}$$

where the known physics-based vector field $f_p(\cdot) : \mathbb{R}^n \to \mathbb{R}^n$ (assumed to have a globally asymptotically stable origin) is augmented with a data-driven mapping $f_d(\cdot) : \mathbb{R}^n \to \mathbb{R}^n$, trained on samples from true system trajectories to compensate for modeling errors. While this typically yields satisfactory accuracy near the training data, the hybrid model (2) becomes unreliable away from the data and can produce physically inconsistent or even unstable trajectories.

3 Methodology

In line with [1], to ensure physical consistency outside the training domain while preserving accuracy within it, we modify the hybrid model (2) as

$$\dot{x}(t) = f_p(x(t)) + f_d(x(t))\Phi(x(t)), \tag{3}$$

where $\Phi(\cdot) : \mathbb{R}^n \to [0,1]$ is a validity function defined as

$$\Phi(x) = \begin{cases} 1, & \|x\| < r_1, \\ \phi(x), & r_1 \le \|x\| \le r_2, \\ 0, & \|x\| > r_2, \end{cases} \tag{4}$$

*This work is co-financed by Holland High Tech, top sector High-Tech Systems and Materials, with a PPP innovation subsidy for public-private partnerships for research and development.

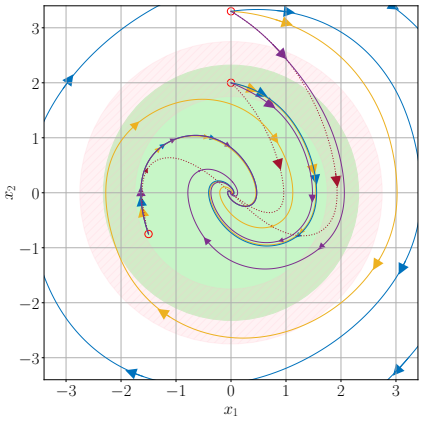

Figure 1: Trajectories of the true system (——), the physics-based model (····), the hybrid model (——), and the *modified* hybrid model (——), starting from multiple initial conditions. The training domain is highlighted by (——) and the transition region by (——).

for some design parameters $0 < r_1 < r_2$, which are chosen to preserve the accuracy of $f_d(\cdot)$ near the training data and suppress its contribution away from the data, thereby transitioning back to the physics-based vector field $f_p(\cdot)$ in regions where $f_d(\cdot)$ is deemed unreliable. This transition is achieved via a smooth *transition mapping* $\phi(\cdot) : \{x \mid r_1 \le \|x\| \le r_2\} \to [0,1]$, which is designed to ensure two key properties of the modified hybrid model: 1) well-posedness and 2) stability-preserving extrapolation. The former is achieved by ensuring that $\Phi(\cdot)$ is Lipschitz continuous on $\mathbb{R}^n$, whereas the latter is achieved based on the known stability properties of the physics-based model. Consequently, we define the transition mapping as

$$\phi(x) = \left(\frac{\|x\| - r_2}{r_1 - r_2}\right)^p \frac{W_p(x)}{\max\left(\frac{\partial V_p}{\partial x} f_d(x), W_p(x)\right)} \tag{5}$$

where $p \ge 1$ controls the sharpness of the transition, and $V_p(x)$ is a known Lyapunov function of the physics-based model, satisfying $\frac{\partial V_p}{\partial x} f_p(x) \le -W_p(x)$ for some $W_p(x) > 0$.

4 Results & Conclusions

Figure 1 presents a simulation case study on a second-order nonlinear system, showing that the proposed modified hybrid model outperforms the original one as it preserves accuracy near the data, aligns with the physics-based model farther away, and provides stability guarantees. Future work will extend the approach to systems with external inputs.

References

[1] A. Habboush, F. Shakib, T. Oomen, and N. van de Wouw, "Safe extrapolation of autonomous data-driven augmented models," in *2025 IEEE 64th Conference on Decision and Control (CDC)*. IEEE, 2025.

Optimized Sampling Design for System Identification

Enrico Dozzi
TU/e
e.dozzi@tue.nl

Tom Oomen
TU/e
t.a.e.oomen@tue.nl

Rodrigo A. González
TU/e
r.a.gonzalez@tue.nl

1 Background

Accurate identification for mechatronic systems traditionally relies on periodic excitation signals and uniform sampling. However, periodic excitations may be impractical to generate, while modern data-acquisition increasingly employs irregular or event-triggered sampling [1]. Under such experimental conditions, standard identification methods can exhibit reduced data informativeness and degraded accuracy due to spectral leakage.

2 Problem Formulation

We consider the identification of a continuous-time, single-input single-output (SISO) linear time-invariant (LTI) system $G(\mathrm{p})$ from nonuniformly sampled input-output data $\{u(t_k), y(t_k)\}_{k=1}^{N}$. The system is excited over the observation window $[0, T]$ by a possibly nonperiodic multisine

$$u(t) = a_0 + \sum_{m=1}^{M} a_m \cos(\omega_m t + \phi_m), \tag{1}$$

where the excitation frequencies ω_m are not constrained to be integer multiples of $2\pi/T$. The objective of this work is to design irregular sampling schemes that enhance the statistical accuracy of Frequency Response Function (FRF) estimation under nonperiodic multisine excitation by improving data informativeness and mitigating spectral leakage. The proposed methodology and its theoretical foundations are detailed in [2].

3 Approach

Leveraging the sparse spectral structure of the nonperiodic multisine, we formulate a least-squares estimator for the FRF at the excited frequencies that is applicable to nonuniformly sampled data. By modeling the sampling instants as random variables drawn from a tunable distribution, we derive finite-sample statistical properties of the estimator and reveal the impact of the sampling distribution. Minimizing spectral leakage in expectation is shown to be equivalent to maximizing the determinant of the expected Fisher Information Matrix (FIM) associated with the estimator, motivating a D-optimal sampling design

$$\begin{aligned} \min_{p_T} \quad & -\log\det\big(R_0(p_T)\big) \\ \text{s.t.} \quad & p_T(t) \geq 0 \quad \text{for } t \in [0,T], \quad \int_0^T p_T(t)\,dt = 1, \end{aligned} \tag{2}$$

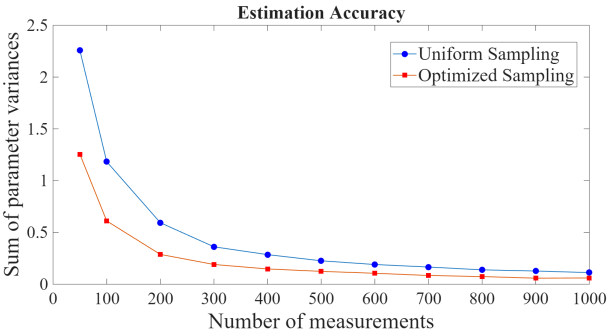

Figure 1: Accuracy of the FRF estimator for uniform sampling and for the D-optimal sampling sequence.

where R_0 is the expected FIM under the sampling distribution p_T. Finally, deterministic sampling sequences approximating the optimized distributions are generated via an iterative algorithm, preserving the statistical benefits predicted by the analysis.

4 Results and Outlook

We illustrate the methodology on the system $G(\mathrm{p}) = 2/(\mathrm{p}^2+2\mathrm{p}+1)$, excited by a nonperiodic multisine with frequencies $\Omega = \{0, \pm 3\sqrt{2}, \pm 3\pi, \pm 12, \pm 12.1\}$ over an observation window $[0, 10]$. For a range of sample sizes, the proposed sampling strategy consistently achieves higher FRF estimation accuracy than uniform sampling (Fig. 1), demonstrating the benefit of shaping the sampling pattern when periodic excitation is not available.

This approach can be viewed as a first step toward event-based system identification, where sampling instants are scheduled based on the information content of the identification problem rather than a fixed clock. This perspective opens the way to more efficient data-acquisition strategies for system identification in scenarios where experiment time, sensing resources, or excitation design are constrained.

References

[1] M. Miskowicz, *Event-based control and signal processing*. CRC press, 2018.

[2] E. Dozzi, T. Oomen, and R. A. González, "D-optimized sampling design for system identification." Manuscript submitted for conference, 2026.

Data-Driven Stabilization Using Prior Knowledge on Stabilizability and Controllability

Amir Shakouri*, Henk J. van Waarde*, Tren M.J.T. Baltussen†, W.P.M.H. (Maurice) Heemels†
* Bernoulli Institute, University of Groningen, The Netherlands
† Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands
a.shakouri@rug.nl h.j.van.waarde@rug.nl t.m.j.t.baltussen@tue.nl m.heemels@tue.nl

1 Introduction

The majority of the literature on data-driven control work in the setting where the parameters of the system are completely unknown. However, for most physical systems, this is a rather conservative modeling framework since we often have access to some *prior knowledge* on the system parameters. For such cases, using both prior knowledge and data can lead to design methods that are less conservative than using data alone. In this work [1], we extend the data informativity framework of [2] to include prior knowledge by requiring the existence of a controller that stabilizes all systems consistent with the data *and* the prior knowledge. In particular, here, we study prior knowledge on stabilizability and controllability of the true system.

2 Problem Formulation

Let $n, m \in \mathbb{N}$. Consider the LTI system

$$x(t+1) = A_{\text{true}}x(t) + B_{\text{true}}u(t), \tag{1}$$

referred to as the *true system*, where $t \in \mathbb{Z}_+$ denotes time, $x(t) \in \mathbb{R}^n$ is the state, and $u(t) \in \mathbb{R}^m$ is the input. The system matrices $(A_{\text{true}}, B_{\text{true}}) \in \mathscr{M} := \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times m}$ are assumed to be unknown. However, we have access to input-state data of the form $\mathscr{D} := \big(\begin{bmatrix}u(0) & \cdots & u(T-1)\end{bmatrix}, \begin{bmatrix}x(0) & \cdots & x(T)\end{bmatrix}\big)$ collected from (1) within the time horizon $T \in \mathbb{N}$. Given $\mathscr{D}$, we define the matrices $U_- := \begin{bmatrix}u(0) & \cdots & u(T-1)\end{bmatrix}$, $X_- := \begin{bmatrix}x(0) & \cdots & x(T-1)\end{bmatrix}$, and $X_+ := \begin{bmatrix}x(1) & \cdots & x(T)\end{bmatrix}$.

We define the set of *data-consistent systems* as

$$\Sigma_{\mathscr{D}} := \{(A,B) \in \mathscr{M} \mid X_+ = AX_- + BU_-\}. \tag{2}$$

The goal of data-driven stabilization is to find a single controller that stabilizes all data-consistent systems, leading to the following definition.

Definition 1 ([2, Def. 12]) *We say $\mathscr{D}$ is* informative for stabilization *if there exists a $K \in \mathbb{R}^{m\times n}$ such that $A+BK$ is Schur for all $(A,B) \in \Sigma_{\mathscr{D}}$.*

Informativity for stabilization has been fully characterized as follows.

Proposition 1 ([2, Thm. 17]) *The data $\mathscr{D}$ is informative for stabilization* if and only if *there exists $\Theta \in \mathbb{R}^{T\times n}$ such that*

$$X_-\Theta = \Theta^\top X_-^\top \text{ and } \begin{bmatrix} X_-\Theta & X_+\Theta \\ \Theta^\top X_+^\top & X_-\Theta \end{bmatrix} > 0. \tag{3}$$

Moreover, $A+BK$ is Schur for all $(A,B) \in \Sigma_{\mathscr{D}}$ if and only if *$K = U_-\Theta(X_-\Theta)^{-1}$ for some Θ satisfying* (3).

Now, let $\Sigma_{\text{pk}} \subseteq \mathscr{M}$ be a set capturing our prior knowledge on the true system, i.e., $(A_{\text{true}}, B_{\text{true}}) \in \Sigma_{\text{pk}}$. Now, we extend the informativity notion in Definition 1 to the following that takes the prior knowledge into account.

Definition 2 *We say $\mathscr{D}$ is* Σ_{pk}–informative for stabilization *if there exists a $K \in \mathbb{R}^{m\times n}$ such that $A+BK$ is Schur for all $(A,B) \in \Sigma_{\mathscr{D}} \cap \Sigma_{\text{pk}}$.*

In this work, we are interested in two important sets of prior knowledge that capture the *stabilizability* and *controllability* of the true system. Denote the sets of controllable and stabilizable systems in $\mathscr{M}$, respectively, by Σ_{cont} and Σ_{stab}. Formally, we thus consider the following problem.

Find necessary and sufficient conditions under which $\mathscr{D}$ is
(i) Σ_{cont}–informative for stabilization;
(ii) Σ_{stab}–informative for stabilization.

3 Summary of Results

Our main results are twofold. First, we show that data-driven stabilization using controllability as prior knowledge is equivalent to data-driven stabilization without prior knowledge. Next, we show that stabilizability as prior knowledge leads to necessary and sufficient conditions for data-driven stabilization that are weaker than those without prior knowledge. In addition, we provide a tractable method to compute a stabilizing feedback gain from data that is Σ_{stab}–informative for stabilization.

References

[1] A. Shakouri, H. J. van Waarde, T. M. J. T. Baltussen, W. P. M. H. Heemels, "Data-driven stabilization using prior knowledge on stabilizability and controllability," *arXiv preprint arXiv:2510.25452*, 2025.

[2] H. J. van Waarde, J. Eising, H. L. Trentelman, M. K. Camlibel, "Data informativity: A new perspective on data-driven analysis and control," *IEEE Transactions on Automatic Control*, vol. 65, no. 11, pp. 4753–4768, 2020.

Analysis of Dissipative Systems from Input-Output Data

Michaela Repášová and Henk van Waarde
Bernoulli Institute, University of Groningen, 9747 AG Groningen, The Netherlands,
m.repasova@rug.nl, h.j.van.waarde@rug.nl

1 Introduction

Dissipativity, first developed in the 1970s by Jan Willems [1], is one of the most significant concepts in systems theory. A system is said to be dissipative if at any time the amount of stored energy can not exceed the amount of energy that has been supplied to it in the past.

We study dissipativity from a behavioral approach, where dynamical systems are defined by the set of all admissible trajectories. In this setting, dissipativity is formalized using so-called quadratic difference forms.

A *quadratic difference form* (in short, QDF) Q_Ψ, associated with the symmetric matrix $\Psi \in \mathbb{R}^{(N+1)q\times(N+1)q}$, is an operator that maps functions $w\colon \mathbb{Z}_+ \to \mathbb{R}^q$ to functions $Q_\Psi(w)\colon \mathbb{Z}_+ \to \mathbb{R}$, given by

$$(Q_\Psi(w))(t) = \sum_{k,l=0}^{N} w(t+k)^\top \Psi_{k,l} w(t+l).$$

Here $\Psi_{k,l} \in \mathbb{R}^{q\times q}$ denotes the block entry of Ψ in the k-th row and j-th column block.

We say that a behavior $\mathscr{B}$ is *dissipative* with respect to the QDF Q_Φ if there exists a QDF $Q_\Psi\colon (\mathbb{R}^q)^{\mathbb{Z}_+} \to \mathbb{R}^{\mathbb{Z}_+}$ such that $Q_\Psi(w)(t) \geq 0$ for all $w \in \mathscr{B}$ and all $t \in \mathbb{Z}_+$, and the dissipation inequality

$$Q_\Psi(w)(t+1) - Q_\Psi(w)(t) \leq Q_\Phi(w)(t),$$

holds for all $w \in \mathscr{B}$ and all $t \in \mathbb{Z}_+$. The QDF Q_Φ is called the *supply rate* and the QDF Q_Ψ is called the *storage function*.

We focus on the analysis of dissipativity in a data-driven context where we study a model class of autoregressive (AR) systems through input-output data collected from an unknown true system.

2 Problem Formulation

We consider the model class of systems of known order L, input dimension m, and output dimension p such that

$$R(\sigma)w = v, \tag{1}$$

where $w = \mathrm{col}(u,y) \in \mathbb{R}^q$, $q = p+m$, $v(t)$ is an additive noise, and σ is the shift operator. Here $R(\xi) \in \mathbb{R}^{q\times p}[\xi]$ is a polynomial matrix of degree L such that

$$R(\xi) = \begin{bmatrix} 0 & I \end{bmatrix} \xi^L + R_{L-1}\xi^{L-1} + \cdots + R_1\xi + R_0.$$

We identify system (1) by the matrix $\tilde{R} \in \mathbb{R}^{p\times qL}$ containing the coefficients $R_0, \ldots, R_{L-1}$ of R.

We collect data $\{w(t)\}_{t=0}^T$ from an unknown true system

$$R_{\text{true}}(\sigma)w = v$$

within the model class. During the experiment, the system is influenced by the noise samples $v(0), \ldots, v(T-L)$ which are not measured. However, it is assumed that the noise samples satisfy the quadratic matrix inequality

$$\begin{bmatrix} I \\ V^\top \end{bmatrix}^\top \Pi \begin{bmatrix} I \\ V^\top \end{bmatrix} \geq 0, \quad V = \begin{bmatrix} v(0) & v(1) & \ldots & v(T-L) \end{bmatrix} \tag{2}$$

for some $\Pi = \begin{bmatrix} \Pi_{11} & \Pi_{12} \\ \Pi_{21} & \Pi_{22} \end{bmatrix} \in \mathbb{S}^{q+r}$ such that $\Pi_{22} \leq 0$, $\Pi \mid \Pi_{22} \geq 0$, and $\ker \Pi_{22} \subseteq \ker \Pi_{12}$. This assumption on the noise captures different bounds on the noise, such as energy noise bounds, individual noise sample bounds, and sample covariance bounds [3].

Define H_1 as the depth-L Hankel matrix of $\{w(t)\}_{t=0}^{T-1}$, and define

$$H_2 = \begin{bmatrix} y(L) & y(L+1) & \ldots & y(T) \end{bmatrix}.$$

We define the *set of systems consistent with the data* $\{w(t)\}_{t=0}^T$ as

$$\Sigma = \left\{ \tilde{R} \in \mathbb{R}^{p\times qL} \;\middle|\; \begin{bmatrix} \tilde{R} & I \end{bmatrix} \begin{bmatrix} H_1 \\ H_2 \end{bmatrix} = V \text{ with } V \text{ satisfying (2)} \right\}.$$

The coefficient matrix $\tilde{R}_{true}$ of the true system is a member of Σ, but in general Σ contains many other systems. As our main result, we provide conditions under which all data-consistent systems $\tilde{R} \in \Sigma$ are dissipative with respect to a given supply rate.

References

[1] J.C. Willems, "Dissipative dynamical systems Part I: General theorem," Archive for Rational Mechanics and Analysis, vol. 45, pp. 321-351, Jan. 1972

[2] J. C. Willems and H. L. Trentelman, "On quadratic differential forms," *SIAM Journal on Control and Optimization*, vol. 36, no. 5, pp. 1703–1749, Sept. 1998.

[3] H. J. van Waarde, K. M. Camlibel, J. Eising, and H. L. Trentelman, "Quadratic matrix inequalities with applications to data-based control," *SIAM Journal on Control and Optimization*, vol. 61, no. 4, pp. 2251–2281, 2023.

Safe and Efficient Iterative Controller Tuning Using Kernel Design

Mathias Schietecat[1,2], Laurens Jacobs[3], Jan Swevers[1,2]
[1]MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
[2]Flanders Make@KU Leuven, Belgium, [3]Nikwist, Belgium
E-mail: mathias.schietecat@kuleuven.be

1 Introduction

Industry relies heavily on linear controllers within a wide range of applications. However, more often than not, these controllers are (i) poorly tuned which results in poor performance, or (ii) tuned very conservatively to ensure robustness in case the controlled process behaves different than expected. While these controllers are functioning, a lot of performance is often left unexplored. In the last few years, Bayesian optimization (BO) has become a popular method for data-based tuning of controller parameters, and has been validated on a wide range of practical problems.

Last year, an extension to these BO methodologies was presented, in which some prior model of the system is assumed to be known [1]. Using this prior information, case specific kernels can be constructed. These kernels form the basis of Gaussian Processes (GP), on which BO is based upon. Proper kernel design can significantly increase the convergence rate of the optimization algorithm. However, the methodology was not able to cope with safety or stability constraint, meaning that the optimization process could result in unstable conditions.

Therefore, GoOSE[2] is selected as a framework to be built around the BO. GoOSE is a Goal Oriented Safe Exploration algorithm which wraps around an existing optimization algorithm and allows the user to apply constraints which should be guaranteed throughout the optimization process, making it safe.

2 Safe Bayesian optimization - GoOSE

Assume a controller $\mathscr{K}(\boldsymbol{x})$ parameterized by the vector $\boldsymbol{x}$ which is constrained within the finite decision space $\mathcal{D} \subseteq \mathbb{R}^d$. The aim is to solve following optimization problem based purely on input-output data of a closed-loop system:

$$\begin{aligned} \min_{\boldsymbol{x}} \quad & f(\boldsymbol{x}) \\ \text{s.t.} \quad & q(\boldsymbol{x}) \geq 0 \end{aligned} \tag{1}$$

In order to solve this problem, BO is used in combination with GoOSE. Therefore both $f(\boldsymbol{x})$ and $q(\boldsymbol{x})$ are approximated by GPs:

$$\begin{aligned} f(\boldsymbol{x}) &\sim \mathcal{GP}\big(\mu_f(\boldsymbol{x}), k_f(\boldsymbol{x},\boldsymbol{x}')\big) \\ q(\boldsymbol{x}) &\sim \mathcal{GP}\big(\mu_q(\boldsymbol{x}), k_q(\boldsymbol{x},\boldsymbol{x}')\big) \end{aligned} \tag{2}$$

At every iteration i, BO uses the current mean $\mu_{f,i}(\boldsymbol{x})$ and standard deviation $\sigma_{f,i}(\boldsymbol{x})$ to propose the next control parameters $\boldsymbol{x}_i^*$ to be tested. In order to make sure that the the parameters $\boldsymbol{x}_i^*$ are safe to evaluate, GoOSE utilizes the current estimate of $q(\boldsymbol{x})$. Depending on $\mu_{q,i}(\boldsymbol{x})$ and its according uncertainty $\sigma_{q,i}(\boldsymbol{x})$, some regions within $\mathcal{D}$ are deemed unsafe, *optimistically* safe or *pessimistically* safe depending on the zero crossing of the upper and lower confidence bound. The described estimated functions and safe or unsafe sets are illustrated in Figure 1.

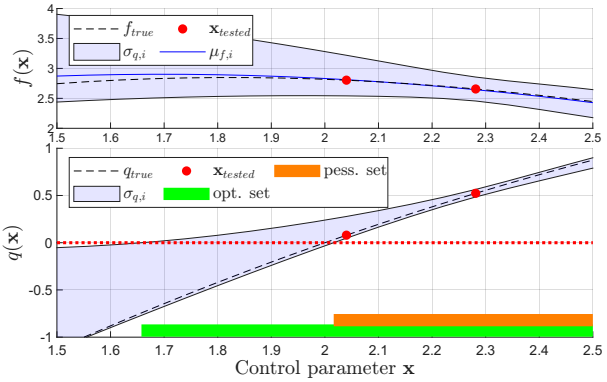

Figure 1: Estimates of functions $f(\boldsymbol{x})$ and $q(\boldsymbol{x})$ at time t

3 Problem Specific Kernel design

The proposed method aims to speed up the combined BO and GoOSE approach by replacing the usually very general kernels, eg. the squared exponential (SE) kernel $k_{SE}(\boldsymbol{x},\boldsymbol{x}')$ by a problem specific one. This kernel is created by realizing that the sum of two Gaussian distributions is again a Gaussian distribution. An identical property is true for GPs. Therefore, the more general SE kernel can be combined with a problem specific kernel $k_{PS}(\boldsymbol{x},\boldsymbol{x}')$ to increase the accuracy of the approximation of $f(\boldsymbol{x})$ and therefore result in significantly faster convergance towards optimality.

$$f(\boldsymbol{x}) \sim \mathcal{GP}\big(\mu_f(\boldsymbol{x}), k_{SE}(\boldsymbol{x},\boldsymbol{x}) + k_{PS}(\boldsymbol{x},\boldsymbol{x}')\big) \tag{3}$$

References

[1] M. Schietecat, "Kernel Design for Efficient Bayesian Optimization in an LPV Framework," in *Book of Abstracts of the 44th Benelux Meeting on Systems and Control*, 2025

[2] M. Turchetta, F. Berkenkamp, and A. Krause, "Safe exploration for interactive machine learning," in *Advances in Neural Information Processing Systems*, 2019

Acknowledgment This work was funded by Flanders Make through SBO project 'ADCEE: Adaptive Dual Control of Actuators for Exploration and Exploitation'.

Computing Scaled Relative Graphs for MIMO LTI Systems From Frequency Response Data

T. Wijfjes*, T. de Groot*, T. Oomen*, S.J.A.M. van den Eijnden*

*Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands

t.wijfjes@student.tue.nl { t.d.groot2, t.a.e.oomen, s.j.a.m.v.d.eijnden}@tue.nl

1 Introduction

Classical control design is based on *linear time-invariant* (LTI) models, for which stability and robustness can be assessed using intuitive frequency-domain tools. A key advantage of these methods is their direct applicability to measured *frequency response functions* (FRFs), enabling nonparametric-based controller design. However, fundamental performance limitations of linear control motivate the use of nonlinear controllers, for which no universal design framework exists. The *scaled relative graph* (SRG) [1], provides a graphical framework for analyzing (possibly nonlinear) operators, enabling direct stability verification of feedback interconnections via separation of SRGs.

Existing methods for computing the SRG assume the availability of an explicit parametric system model. In this work, we propose method for constructing an over-approximation of the SRG of a multi-variable, bounded LTI system directly from measured data, aiming to make SRG-based stability analysis compatible with established nonparametric frequency-domain design tools.

2 Approach

Our approach builds on the connection between the SRG and the numerical range of a linear operator. In [2] it is shown that the SRG of a linear system can be obtained from the numerical range of a transformed system. This relation allows SRG computation to be formulated as a numerical range approximation problem.

Starting from measured frequency response data, we exploit the fact that evaluating the FRF at a single frequency yields a complex matrix. For each frequency, this matrix is first mapped via the Beltrami–Klein transformation, after which the numerical range of the transformed matrix is computed independently using an efficient boundary tracing method. The resulting frequency-wise numerical ranges are combined by taking the union over frequency. Applying the inverse Beltrami–Klein transformation then yields an SRG that is valid over the measured frequency band.

To account for unmeasured low- and high-frequency regions, we conservatively over-bound the SRG using the worst-case $\mathscr{L}_2$-gain outside the measurement band. Under assumptions of properness and boundedness of the underlying system, this construction provides an SRG over-approximation that is valid over the entire frequency axis.

Accurate SRG construction requires a fine frequency grid to capture sharp resonant dynamics, which is further amplified by the geometric warping induced by the Beltrami–Klein mapping. Direct measurement at such high resolution is not feasible. Instead, we employ the Local Rational Method (LRM), which fits locally valid rational models over narrow frequency windows [3]. These local models are evaluated on a desired fine frequency grid and combined via weighted averaging to reconstruct a finely sampled frequency response, which is subsequently used for SRG computation.

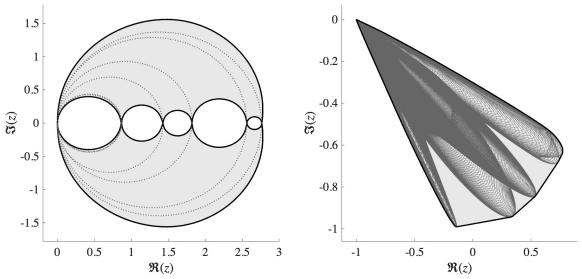

Figure 1: SRG and the systems characteristic loci (left) and Numerical range and frequency-wise numerical ranges (right).

3 Example

Figure 1 demonstrates the proposed SRG construction via the numerical range for a MIMO system, showing how the SRG boundary is determined by the numerical range. Furthermore, the principal shape of the SRG reflects the characteristic loci of the underlying system.

References

[1] T. Chaffey, F. Forni, and R. Sepulchre, "Graphical Nonlinear System Analysis," *IEEE Transactions on Automatic Control*, vol. 68, no. 10, 2023.

[2] R. Pates, "The Scaled Relative Graph of a Linear Operator," arXiv:2106.05650, 2021.

[3] R. J. Voorhoeve, A. van der Maas, and T. Oomen, "Non-parametric MIMO system identification using local rational models," *Mechanical Systems and Signal Processing*, vol. 105, 2018.

Neural Approximation of Optimal Control: Sampling Strategies and Safety Guarantees

Andrea Giusti, Maxime Monsieur and Branimir Mrak
MotionS, Flanders Make VzW
3920 Lommel, Belgium
Email: branimir.mrak@flandersmake.be

1 Motivation

Model Predictive Control (MPC) is a powerful framework for the optimal control of constrained, nonlinear, and multi-degree-of-freedom systems. Despite its effectiveness, the real-time deployment of MPC is often limited by its computational complexity, particularly on resource-constrained platforms or in high-frequency control loops.
Neural network-based approximations of MPC (NN-MPC) provide a promising alternative, offering rapid inference while closely emulating the optimal control policy [1]. However, this approach introduces new challenges, most notably: (i) the efficient generation of informative training data, and (ii) the guarantee of safe and reliable closed-loop behavior.

2 Approach

Our work leverages an iterative imitation learning framework to efficiently train NN-MPC controllers while systematically validating their robustness across the task space. Rather than relying on large, uniformly sampled datasets, the method selectively queries the MPC expert only in regions where the NN-MPC exhibits performance degradation. We introduce the Sampling Trajectories for Robust Imitation and Dynamic Evaluation (STRIDE) algorithm, which extends the DAGGER approach [2]. STRIDE integrates two complementary mechanisms, as shown in Algorithm 1: one *focuses* on the training MPC trajectories to ensure closed-loop robustness, and the other *explores* new test conditions. This cycle continues until the NN-MPC demonstrates satisfactory performance across the entire task space. Finally, safety of the neural network controller and the closed-loop system is analyzed through set propagation, ensuring that neither state nor input constraints are violated.

3 Results

The approach is validated on a parallel SCARA robot executing energy-optimal point-to-point motions under a 1 kHz control loop (see Figure 1; video available at https://youtu.be/D_TQ-tqjhYQ). The resulting NN-MPC, implemented as a compact neural network, achieves near-optimal performance while dramatically reducing computational overhead. The iterative training strategy converges rapidly, minimizing the number of MPC queries required and ensuring robust closed-loop behavior throughout the task space.

Algorithm 1 STRIDE

- Generate initial MPC dataset (grid of initial/final conditions)
- Explore
  - Focus
    - Train NNMPC
    - Test NNMPC on training trajectories
    - Select point where NNMPC trajectory diverges from dataset
    - Generate new MPC from that point
  - Generate new initial conditions outside of existing dataset
  - Test NNMPC from new initial conditions
  - Select trajectories where NNMPC failed and generate new MPC

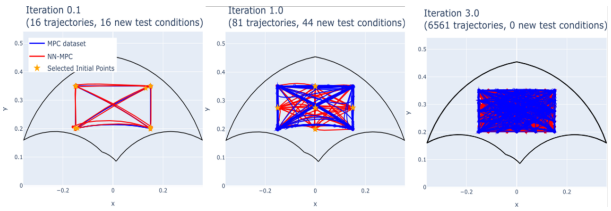

Figure 1: Trajectories generated by MPC and NN-MPC controllers in the workspace of the SCARA robot.

References

[1] Edward Kikken et al. "Approximating MPC solutions using Deep Neural Networks: Towards Application in Mechatronic Systems". en. In: *Proceedings of the 22nd International Conference on Informatics in Control, Automation and Robotics*. Vol. 1. 2025, pp. 71–81. DOI: 10.5220/0013712900003982.

[2] Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. "A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning". In: *Proceedings of Machine Learning Research* (Mar. 2011).

Acknowledgment This research was supported by Flanders Make, the strategic research center for the manufacturing industry, in the framework of the LearnOpTra_SBO project.

Distributed data-driven LQR of LTI system[1]

Surya Malladi
ENTEG, University of Groningen
v.s.p.malladi@rug.nl

Nima Monshizadeh
ENTEG, University of Groningen
n.monshizadeh@rug.nl

1 Introduction

Data-driven control has emerged as a powerful alternative to classical model-based approaches, offering the ability to synthesize controllers directly from collected data. A central assumption of these works is the fact that data is available centrally. Motivated by the fragmented availability of data due to architectural and privacy constraints, we design distributed data-driven algorithm to compute the Linear Quadratic Regulator (LQR) of linear time-invariant systems.

2 Problem Statement

Consider a linear system with input $\dot{x} = Ax + Bu$, where (A,B) is stabilizable. We suppose that N computing agents are available, where each agent i has access only to a single data sample $\{u(t_i), x(t_i), r(t_i)\}$, with $i \in \mathscr{I} := \{1,2,\dots,N\}$, where $u(t_i) \in \mathbb{R}^m$, $x(t_i) \in \mathbb{R}^n$ and $r(t_i)$ respectively denote the input data, measured state data and rate of change of state data of the LTI system at time t_i. We aim to compute the LQR controller, which is obtained using the unique positive definite solution P^* of the following algebraic Riccati equation (ARE)

$$A^T P + PA + Q - PBR^{-1}B^T P = 0, \tag{1}$$

where $Q, R > 0$. It is well-known that the controller gain given by $K^* = -R^{-1}B^T P^*$, minimizes the infinite horizon linear quadratic cost

$$J(K) = \int_0^\infty x^T(t)(Q + K^T R K)x(t)dt.$$

We assume that the input matrix B is publicly known, whereas the system matrix A is unknown.

Problem 1 *Given the matrices B, Q, R and the data $\{x(t_i), r(t_i), u(t_i)\}$ available to each agent $i \in \mathscr{I}$, design a distributed data-driven algorithm to compute the optimal control gain K^*.*

We impose the following assumption on the data: it holds that $\operatorname{rank}(X_0) = n$, where $X_0 := \begin{bmatrix} x(t_1) & x(t_2) & \cdots & x(t_N) \end{bmatrix}$ denotes the combined state data.

3 Main result

As the first step, we devise a scheme which splits the unknown matrix A into N (low-rank) shares A_i, such that $A = \sum_{i=1}^N A_i$ and each A_i is available to agent i. The data samples satisfy $r(t_i) = Ax(t_i) + Bu(t_i)$, for each $i \in \mathscr{I}$. Defining $\hat{r}(t_i) := r(t_i) - Bu(t_i)$ for each $i \in \mathscr{I}$, observe that the system matrix A satisfies

$$A = \sum_{i=1}^N \hat{r}(t_i) y^T(i), \quad \text{for } \{y(i)\}_{i=1}^N \text{ s.t. } I_n = \sum_{i=1}^N x(t_i) y^T(i) \tag{2}$$

By treating the constraint (2) as a distributed resource allocation problem, the agents, relying only on neighboring communication, can compute their respective variables $y(i)$, such that the resulting set of computed vectors satisfy (2). This enables each agent i to compute the term $A_i := \hat{r}(t_i) y^T(i)$, which precisely coincides with one of the N terms appearing in the first summation in (2).

Theorem 1 *Let each agent $i \in \mathscr{I}$ has access to rank one matrix A_i above, and employ the distributed dynamics*

$$\begin{aligned} \dot{P}_i =& NA_i^T P_i + NP_i A_i + Q - P_i BR^{-1}B^T P_i \\ &+ \gamma \sum_{j \in \mathscr{N}_i} (P_j - P_i) + \gamma \sum_{j \in \mathscr{N}_i} (Y_j - Y_i) \end{aligned} \tag{3a}$$

$$\dot{Y}_i = -\gamma \sum_{j \in \mathscr{N}_i} (P_j - P_i) \tag{3b}$$

with $P_i(0) = \mathbf{0}$, $Y_i(0) = \mathbf{0}, \forall\ i \in \mathscr{I}$. Then, there exists a coupling gain $\gamma^* > 0$ such that for any $\gamma > \gamma^*$, the solutions of the distributed dynamical algorithm (3) *satisfy $\lim_{t\to\infty} \|P_i(t) - P^*\| = 0$, where P^* is the unique stabilizing solution to* (1).

The following plot shows the result of the application of the dynamics (3) to a linearized model of the helicopter hover dynamics [1] with $n = 8$, $m = 4$, $N = 16$, and $\gamma = 10000$.

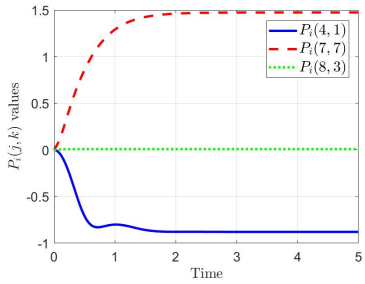

Figure 1: Asymptotic convergence to the ARE's solution

References

[1] G. D. Padfield, *Helicopter flight dynamics: the theory and application of flying qualities and simulation modelling*. John Wiley & Sons, 2008.

[1] The work of authors is supported by NWO grant OCENW.KLEIN.257.

System Identification and μ-Synthesis for Broadband MIMO Active Noise Control in Cooling Circuits for High-Tech Equipment

Marin Jalšić[1], Hans Vermeulen[1,2], Marc van der Wijst[2] and Claire Bourquard[3]
[1] Control Systems Technology, Eindhoven University of Technology, Eindhoven, The Netherlands.
[2] ASML, Veldhoven, The Netherlands.
[3] Dynamics and Control, Eindhoven University of Technology, Eindhoven, The Netherlands.
m.j.jalsic@tue.nl

1 Introduction

Active Noise Control (ANC) in heating, ventilation and air conditioning (HVAC) systems is commonly investigated in an effort to reduce noise in workplace environments generated, for example, by fans [1]. However, a conceptually similar problem arises in precision engineering applications, namely, in photolithography. Due to the strict thermal requirements, sensitive elements may require water cooling, and because a cooling medium is introduced, disturbances propagate through the cooling system, negatively affecting the precise positioning of the elements. These disturbances may also have dominant low-frequency content, so passive mitigation measures, such as Helmholtz resonators, which become bulky when targeting lower frequency bands [2], have limited effectiveness. With this in mind, ANC becomes an attractive approach to reduce the impact of such phenomena on machine performance. A major difference between HVAC systems and cooling circuits for photolithography machines is that in waterlines dominant disturbances may originate anywhere where there is a pipe discontinuity: both up- and downstream from the location where they are to be minimized. Because of this, considering multi-input-multi-output (MIMO) models and control systems may be necessary. In this work, experimental results are presented, in which a controller designed using μ-synthesis for an air-filled pipe system mitigates the influence of disturbances that originate both up- and downstream of the region where they must be minimized.

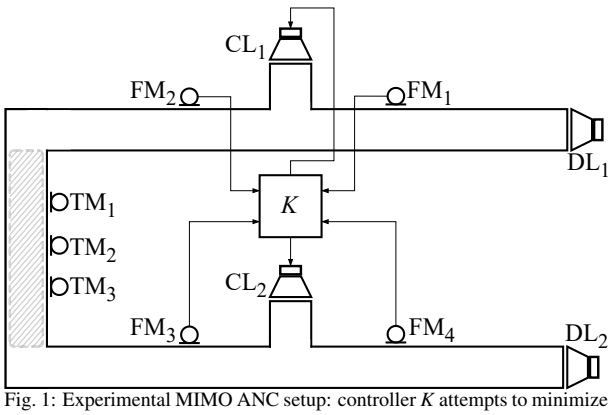

Fig. 1: Experimental MIMO ANC setup: controller K attempts to minimize the effect of disturbance loudspeakers DL_1 and DL_2 on target microphones TM_1, TM_2 and TM_3 (gray hatched area) by controlling the input of control loudspeakers CL_1 and CL_2, given measurements of feedback microphones FM_1, FM_2, FM_3 and FM_4

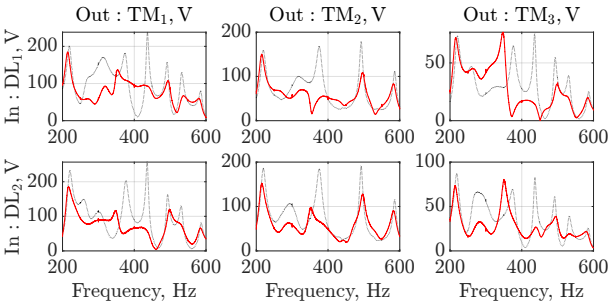

Fig. 2: Measured transfer functions from voltage on disturbance loudspeakers to voltage on target microphones:
without ANC (dotted black line), with ANC (thick red line)

2 System Description and Methodology

The experimental setup shown in Figure 1 consists of a U-shaped air-filled pipe with a circular cross section terminated in loudspeakers, T-junctions (which are also terminated in loudspeakers), and microphones mounted air-tight and flush with the inner wall of the pipe. In order to obtain a parametric state space model for control system design, the Adaptive Antoulas-Anderson algorithm [3] is used on the estimated frequency response function (FRF) data. For μ-synthesis, uncertainty weights for the transfers involving the control loudspeakers are generated from the differences between the model and estimated FRFs. The performance weights are selected as band-pass filters for FM_2 and FM_3 and small constant values for FM_1 and FM_4. Small constant values are used also as noise filters on all microphones.

3 Results

Although the controller primarily attempts to minimize the values at FM_2 and FM_3, TM_1, TM_2 and TM_3 have their magnitudes reduced as well. In the target frequency band locally the magnitude is reduced by more than a factor of 2.5, showing that disturbances originating both upstream and downstream from the target region can be successfully mitigated in this way.

References

[1] Martin Larsson. *Active noise control in ventilation systems: practical implementation aspects*. Department of Signal Processing, School of Engineering, Blekinge Institute of Technology, Karlskrona, 2008.

[2] A. J. Hull, C. J. Radcliffe, and S. C. Southward. Global Active Noise Control of a One-Dimensional Acoustic Duct Using a Feedback Controller. *Journal of Dynamic Systems, Measurement, and Control*, 115(3):488–494, September 1993.

[3] Yuji Nakatsukasa, Olivier Sète, and Lloyd N. Trefethen. The AAA Algorithm for Rational Approximation. *SIAM Journal on Scientific Computing*, 40(3):A1494–A1522, January 2018.

Conic programming duality with applications to Lyapunov inequalities and linear quadratic control

Robin Bulsink, Henk J. van Waarde, M. Kanat Camlibel
Bernoulli Institute, University of Groningen, 9747AG Groningen, the Netherlands
Email: r.m.bulsink@rug.nl, h.j.van.waarde@rug.nl, m.k.camlibel@rug.nl

1 Introduction

Conic programming is the optimization of a linear function subject to feasibility constraints, in terms of inclusion in a cone. Inspired by the rich existing duality theory in semidefinite programming (SDP) [1], we formulate a theory of duality in the conic programming context. We find that this both generalizes SDP duality and considerably simplifies some notable proofs. Finally, it is shown how this duality theory can be applied to study problems in systems and control theory; specifically solutions Lyapunov inequalities and the (infinite-horizon) linear quadratic regulator (LQR) problem.

2 Primal-dual formulation

We consider the *primal* conic program

$$\begin{aligned} \min_{x} \quad & \langle \xi, x\rangle_X \\ \text{s.t.} \quad & x \in C, \\ & \eta + L(x) \in K, \end{aligned} \tag{CP}$$

where X, Y denote finite-dimensional inner product spaces, $C \subset X$ and $K \subset Y$ are closed convex cones and $L : X \to Y$ is a linear map. $\xi \in X$ and $\eta \in Y$ are the defining parameters of the problem. We refer to the primal program as $\mathrm{CP}(\xi, \eta, C, K, L)$. Correspondingly, the *dual* program is

$$\begin{aligned} \min_{y} \quad & \langle \eta, y\rangle_Y \\ \text{s.t.} \quad & y \in K^\star, \\ & \xi - L^\star(y) \in C^\star, \end{aligned} \tag{CP*}$$

where $C^\star$ denotes the dual cone of C, and $L^\star : Y \to X$ denotes the adjoint map of L. We denote the dual program $\mathrm{CP}^\star(\xi, \eta, C, K, L)$. Our aim is to study feasibility and optimality of these problems, exploiting useful relations between the primal and dual. Noting that the set of positive semidefinite matrices is a closed convex cone, it is straightforward to see that SDP is recovered from the above setup by specifying C and K. It should also be noted that the structure of the primal and dual programs is exactly equivalent:

$$\mathrm{CP}^\star(\xi, \eta, C, K, L) = \mathrm{CP}(\eta, \xi, K^\star, C^\star, -L^\star).$$

3 Conic programming duality

We proceed by studying relations between the optimal values of both problems. We denote optimal value of the primal program as follows:

$$\mu(\mathrm{CP}) := \inf_{\substack{x \in C, \\ \eta + L(x) \in K}} \langle \xi, x\rangle_X.$$

The optimal value of the dual, denoted $\mu(\mathrm{CP}^\star)$, is defined exactly analogously. Of particular interest is the equality known as *strong duality*: $\mu(\mathrm{CP}) = -\mu(\mathrm{CP}^\star)$

Theorem 1. *If the primal problem (CP) is strictly feasible, i.e. there exist $x \in int(C)$ such that $\eta + L(x) \in int(K)$, and its optimal value is bounded below, strong duality holds. Moreover, the dual optimum is attained, i.e. there exists $y^+ \in K^\star$ such that $\xi - L^\star(y^+) \in C^\star$ and $\langle \eta, y^+\rangle = \mu(CP^\star)$.*

The main tools in our analysis are *theorems of alternatives*, such as the following:

Theorem 2. *Let $C \subseteq X$, $K \subseteq Y$ be closed convex cones with nonempty interior and let $\eta \in Y$. Then, exactly one of the following is true:*

1. *There exists some $x \in int(C)$ such that $\eta + L(x) \in int(K)$.*
2. *There exists some $y \in (K - L(C))^\star \setminus \{0\}$ such that $\langle y, \eta\rangle \leq 0$.*

4 Control-theoretic applications

By making the proper definitions, we show how we can define a primal-dual pair such that under strong duality, the conic programs recover exactly the solution to a version of the linear quadratic regulator problem. In particular, we show that the assumptions of Theorem 1 are satisfied under particular system-theoretic conditions. Under these conditions, the primal program recovers exactly the maximal solution to the corresponding Algebraic Riccati equation [2]. We also show how Theorem 2 recovers a Lyapunov stability result for linear autonomous sytems.

References

[1] V. Balakrishnan, L. Vandenberghe, "Semidefinite programming duality and linear time invariant systems", *IEEE Transactions on Automatic Control 48.1 (2003): 30-41*

[2] H. Trentelman, A. Stoorvogel, M. Hautus, "Control Theory for Linear Systems", University of Groningen, Bernoulli Institute, 2002.

Distributed Optimization of Charging Energy Hub Networks

Diego Fernández-Zapico, Theo Hofman, Mauro Salazar
Control Systems Technology section
Eindhoven University of Technology
Email: {d.fernandez.zapico, t.hofman, m.r.u.salazar}@tue.nl

Abstract

The adoption of electrical vehicles (EV) requires improving the charging infrastructure to reduce the strain on the electricity grid and to improve user experience. In this context, a Charging Energy Hub (CEH) can help support EV charging stations by introducing renewable energy sources and battery energy storage systems (BESS). The optimal operation of a network of CEHs introduces challenges regarding user privacy, scalability and robustness. Distributed optimization can overcome these challenges by solving local optimization problems of each CEH and reaching a consensus over their shared connections.

We propose solving this consensus distributed optimization problem with ADMM. In Fig. 1, we present the DC network of three CEHs, which are composed of different EV charging demands (P_{ev}), photovoltaic generations (P_{pv}), limits on power exchanges with the grid (P_{g}) and limits on BESS power injections (P_{b}) and BESS energy levels (E_{b}). We present the preliminary results of the optimal operation of the CEH network over a 36 h horizon in Fig. 2.

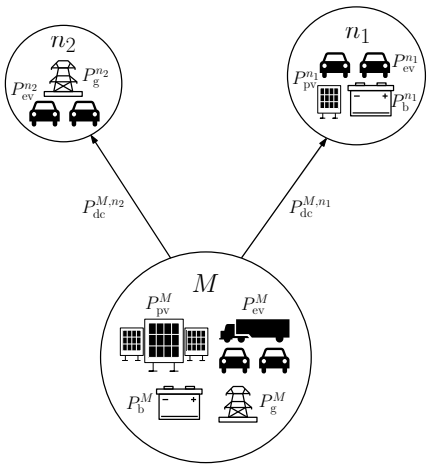

Figure 1: CEH Network composed three hubs (M, n_1, n_2) connected by DC power lines $(P^{M,n_1}_{\text{dc}}, P^{M,n_2}_{\text{dc}})$.

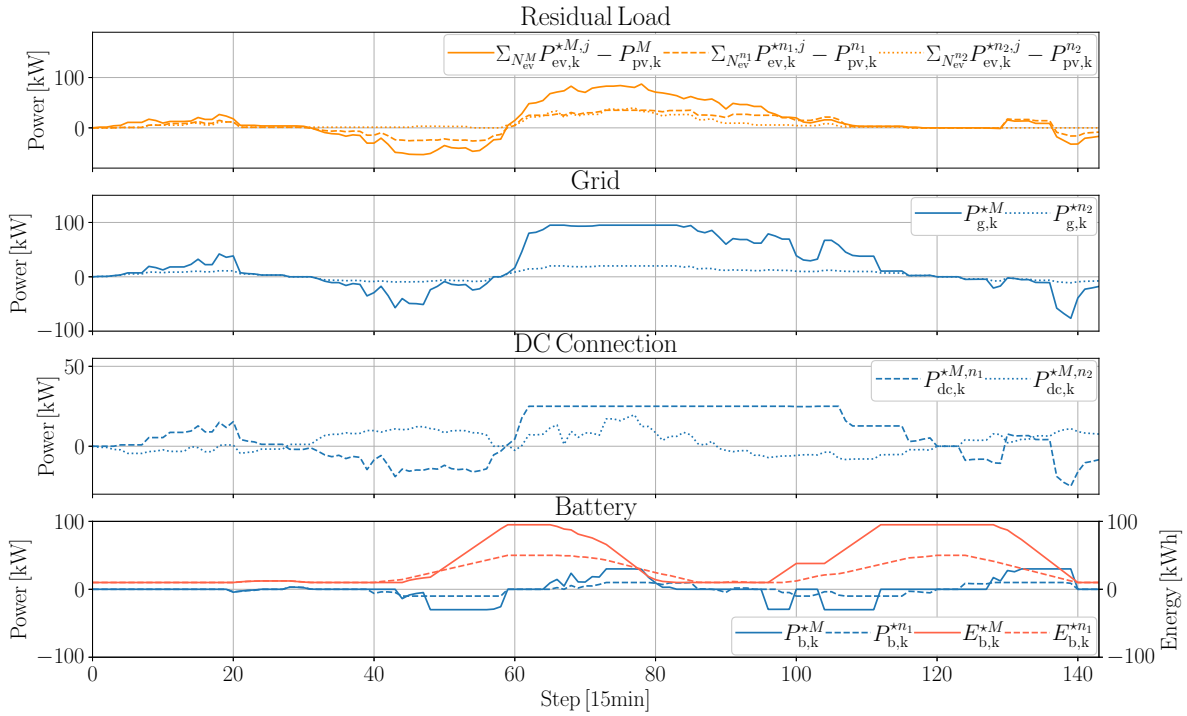

Figure 2: Optimal operation over a 36 h horizon of the CEH network.

Hidden Convexity in Linear Neural Networks

Adeline Colson
UCLouvain
adeline.colson@uclouvain.be

Benoît Legat
UCLouvain
benoit.legat@uclouvain.be

Chenyang Yuan
Toyota Research Institute
ycy@mit.edu

1 Introduction

While nonconvexity is traditionally viewed as a challenge in optimization, many machine learning models exhibit benign nonconvexity, where nonconvex formulations are surprisingly tractable and often more scalable than their convex counterparts. Deep linear neural networks are a canonical example. Despite their simplicity, these models capture key difficulties of deep learning while remaining analytically tractable. Although the objective function is nonconvex in the network parameters, it has been shown that the optimization landscape admits no spurious local minima: every local minimum is global [1, 2, 3]. Existing analyses typically establish this property by explicitly characterizing all stationary solutions [2].

In this work, we adopt a different perspective. We show that the absence of spurious local minima in deep linear networks can be explained by an underlying hidden convexity: local optimality conditions implicitly define a convex optimization problem, which allows global properties such as the optimality of local minima to be inferred. This approach enables the derivation of an explicit global optimality certificate derived directly from local conditions through a semi-automated proof procedure.

Our framework applies to deep linear networks with or without bias terms, naturally extending previous analyses. Biases are important in practice, enabling affine transformations, but unlike weights, they cannot be rearranged or factorized, breaking the symmetries used in previous proofs.

2 Linear Neural Network

We consider deep linear neural networks trained with the mean squared error loss. For a network with an input matrix X, an output matrix Y, l hidden layers and weight matrices $W_1, \dots, W_{l+1}$, the model is defined through the loss

$$\mathscr{L}(W) = ||W_{l+1}W_l....W_1X - Y||_F^2$$

Although such models are linear for $l = 1$, increasing depth introduces nonconvexity, making them a tractable starting point for building theory and later generalizing to more complex models.

Previous works have shown that, under suitable assumptions, all local minima of this objective are global [1, 2, 3]. We revisit this result using a different approach and extend it to networks with bias terms $b_1, \dots, b_{l+1}$, for which the loss takes the form

$$\mathscr{L}(W,b) = \|W_{l+1}W_l....W_1X + \sum_{i=1}^{l+1}(W_{l+1}....W_{i+1}b_i\mathbf{1}) - Y\|_F^2$$

Our analysis distinguishes between full-rank and rank-deficient regimes, where the rank refers to that of the weight matrices. In the rank-deficient case, we show that the local optimality conditions guarantee zero training error. The full-rank case is addressed in the following section.

3 Hidden Convexity

We propose a general methodology for certifying global optimality in benign nonconvex problems based on hidden convexity.

The first-order optimality condition $\nabla\mathscr{L}(W) = 0$ characterizes the structure of the solutions. In particular, it enables the singular value decomposition of the output matrix Y, which in turn allows a reformulation of the problem as the convex program

$$\min_Z \langle YY^\top, Z\rangle \quad s.t. \quad \mathrm{tr}(Z) = r, \quad 0 \preceq Z \preceq I$$

Here, Z represents a projection matrix, and r denotes the rank of the weight matrices.

We then show that the second-order optimality condition $\nabla^2\mathscr{L}(W) \succeq 0$ guarantees the feasibility of the dual problem. The resulting dual certificate establishes that any point satisfying the local conditions is a global minimum of the nonconvex objective.

Unlike prior analyses, our approach provides a unified and flexible framework for deep linear networks, both with and without biases.

References

[1] K. Kawaguchi. "Deep learning without poor local minima". In: Advances in neural information processing systems. Vol. 29. 2016.

[2] Y. Zhou and Y. Liang. "Critical points of linear neural networks: Analytical forms and landscape properties". In: International conference on learning representations. 2018.

[3] L. Zhang. "Depth creates no more spurious local minima". In: International conference on learning representations. Vol. abs/1901.09827. CoRR. 2020.

An Efficient Metaheuristic for Optimal Sensor Placement

Raphaël Goetz[1], Y. Dwaraga[1], Nathan van de Wouw[1], Tom Oomen[1], Hans Zwart[1,2]

[1] Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands

[2] Department of Applied Mathematics, University of Twente, The Netherlands

Email: r.p.p.f.goetz@tue.nl

Background

The sensor placement problem that maximizes the estimation performance of the Kalman filter is NP-hard [1]. Several strategies have been proposed to solve this problem, including the use of surrogate objectives that exhibit properties guaranteeing convergence to near-optimal solutions [1], or the use of convexification techniques that enable the use of convex optimization solvers [2]. However, both strategies yield only suboptimal solutions.

Here, we adopt a strategy based on metaheuristics as these algorithms may find the global optimum of NP-hard problems. Specifically, we focus on the Artificial Bee Colony (ABC) algorithm and show its effectiveness and efficiency in solving the NP-hard sensor placement problem for a 3D thermoelastic system.

Sensor placement problem formulation

The sensor placement problem is formulated for the discrete-time stochastic system,

$$\begin{bmatrix} x_{k+1} \\ y_k \\ z_k \end{bmatrix} = \begin{bmatrix} A & B_u & B_w & 0 \\ C & 0 & 0 & I \\ L & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_k \\ u_k \\ w_k \\ v_k \end{bmatrix},$$

with the discrete time instant k, state x, measurements y, known input u and performance output z to be estimated. The disturbance w and measurement noise v are zero-mean white noises and the initial state x_0 is a Gaussian random variable. They are assumed to be mutually independent.

The sensor layout is characterized by C and is assumed binary where each row corresponds to a sensor. A 1 entry indicates the presence of a sensor to measure the corresponding state, a 0 indicates its absence. This binary nature is captured by the constraint $C \in \Omega_C$.

The desired estimation performance is the average output estimation error $e_k = z_k - \hat{z}_{k|k}$ on the time interval $[1, k_f]$. The sensor placement optimization problem is given by

$$\min_{C \in \Omega_C} \frac{1}{k_f} \sum_{k=1}^{k_f} \mathbb{E}\left[\|e_k\|^2\right] \Leftrightarrow \min_{C \in \Omega_C} \frac{1}{k_f} \sum_{k=1}^{k_f} \mathrm{tr}[P_{z_{k|k}}], \qquad (1)$$

where $\hat{z}_{k|k}$ is the estimate of z_k obtained by the Kalman filter and $P_{z_{k|k}} = \mathbb{E}\left[(z_k - \hat{z}_{k|k})(z_k - \hat{z}_{k|k})^\top\right]$ is the associated error covariance matrix depending on the sensor layout C.

A binary variant of the ABC

The Novel Binary Artificial Bee Colony (NBABC) algorithm is used as it is shown to outperform nine other algorithms on high-dimensional and NP-hard problems [3].

As any other metaheuristics, the NBABC has a crossover mechanism to generate new solutions. This mechanism is adapted to preserve the fixed number of sensors. First, the binary matrix C is expressed as a binary vector c where the jth entry equals 1 if a sensor is placed at the jth candidate location and 0 otherwise. To generate a new sensor configuration from the agent's one c_i, first, a partner solution c_j is randomly chosen from the set of agents' solutions. The indices at which $c_i \neq c_j$ are identified. From these indices, one entry where $c_i = 1$ is randomly selected and flipped to 0, and another one where $c_i = 0$ is randomly selected and flipped to 1. The newly generated solution replaces c_i if it has a better cost.

Results

The NBABC is used to solve the sensor placement problem (1) for a 3D thermoelastic system. Two sensors must be positioned optimally among 124 candidate locations. Because the NBABC is probabilistic, its convergence is assessed using the median cost over 50 runs. While there is more than 7500 possible configurations – making an exhaustive search inefficient – Figure 1 shows that the NBABC converges to a near-optimal solution only after 500 function evaluations.

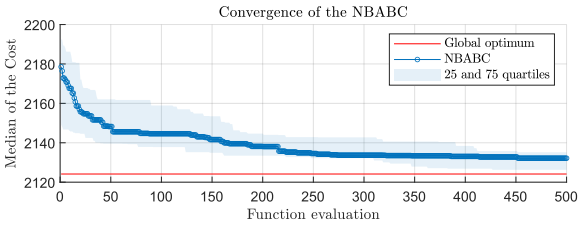

Figure 1: Median of the cost per function evaluation over 50 runs.

References

[1] H. Zhang, R. Ayoub and S. Sundaram, "Sensor selection for Kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms", Automatica, 2017.

[2] N. Darivandi, K. Morris and A. Khajepour, "An algorithm for LQ optimal actuator location", Smart Materials and Structures, 2013.

[3] C.J. Santana, M. Macedo, H. Siqueira, A. Gokhale and C.J.A. Bastos-Filho, "A novel binary artificial bee colony algorithm", Future Generation Computer Systems, 2019.

Benign nonconvexity in sparse low-rank PSD matrix completion

Sophie Lequeu
UCLouvain
sophie.lequeu@uclouvain.be

Chenyang Yuan
MIT
ycy@mit.edu

Benoît Legat
UCLouvain
benoit.legat@uclouvain.be

We study the nonconvex Burer-Monteiro reformulation of the low-rank positive semidefinite (PSD) matrix completion problem. Although this reformulation is nonconvex, it is often solved with first-order methods and exhibits strong empirical performance. Existing theoretical results explain this behavior by showing that, under suitable regularity assumptions or with the addition of explicit regularization, the associated objective function admits no spurious local minima. The goal of this research is to understand this benign nonconvexity without relying on such assumptions, by instead exploiting structural properties of the underlying sparsity pattern.

1 Introduction

Low-rank PSD matrix completion can be formulated as a semidefinite program (SDP) with affine constraints (to match partially observed entries, denoted with b). Given a graph G encoding the sparsity pattern, with edges corresponding to observed entries, let $\mathscr{A}^G : \mathbb{S}^n \to \mathbb{S}^n$ denote the associated masking operator. The problem consists in

$$\begin{aligned} &\text{finding } X \in \mathbb{S}^n \\ &\quad \text{s.t. } X \succeq 0, \quad \mathscr{A}^G(X) = b, \end{aligned} \tag{1}$$

with the additional objective of recovering a solution X of low rank.

To address scalability, a common approach is to apply the Burer-Monteiro factorization $X = UU^\top$ with $U \in \mathbb{R}^{n\times r}$, which transforms the SDP into the smooth but nonconvex optimization problem

$$\min_{U\in\mathbb{R}^{n\times r}} \|\mathscr{A}^G(UU^\top) - b\|^2. \tag{2}$$

In practice, this formulation is solved using first-order methods and, despite its nonconvexity, often converges to globally optimal solutions. Existing theoretical results [2] explain this behavior by imposing regularity assumptions or explicit regularization. Our focus is on the unregularized formulation, where we seek to explain benign nonconvexity through structural properties of the sparsity pattern instead.

2 Chordal sparsity patterns

Our analysis builds on classical results in PSD matrix completion over chordal graphs. A partial symmetric matrix admits a PSD completion if and only if all fully specified principal submatrices associated with the maximal cliques are PSD [3], and such a completion can be chosen with rank bounded by the maximum clique size [1]. This naturally raises the question of whether the benign nonconvexity observed in the Burer-Monteiro formulation can be understood through local clique conditions, just as feasibility and rank results.

3 The complete graph case

We first consider the case of a complete sparsity graph G, which serves as a baseline for analyzing the general case.

Lemma. *For a complete sparsity graph, (2) becomes*

$$\min_{U\in\mathbb{R}^{n\times r}} \|UU^\top - b\|^2 \tag{3}$$

and every second-order critical point (SOCP) is globally optimal.

Although this setting is no longer a matrix completion problem, as all entries are observed and the factorization reduces to a Cholesky decomposition, the proof helps the extension to chordal sparsity patterns.

4 The chordal graphs case

For the more general case of chordal sparsity patterns, the problem decomposes over maximal cliques.

Lemma. *Let $U^\star$ be a SOCP of formulation (2) associated with a chordal graph. If, for each maximal clique, the submatrix of $U^\star$ (restricted to the corresponding nodes) has full row rank, then $U^\star$ is globally optimal.*

This result provides a sufficient condition for global optimality of a point U^* based on local rank properties, but does not characterize all SOCP in general. The remaining cases still require further analysis and may require additional conditions on the measurements b.

As a first step, our research focuses on chordal graphs with a small number of maximal cliques, with the goal of extending these arguments and conditions to general chordal patterns. We formulate conditions under which local minima are also guaranteed to be global ones, to characterize the possibly benign nonconvexity of the problem.

References

[1] Jerome Dancis. "Positive semidefinite completions of partial Hermitian matrices" (1992).

[2] Rong Ge et al. "Matrix Completion has No Spurious Local Minimum" (2018).

[3] Robert Grone et al. "Positive definite completions of partial Hermitian matrices" (1984).

Quadrotor Stabilization in Offshore Wind Conditions: Assessing the Benefit of Future Wind Information

E. Wisse, B.G.L. Meere, S. Doodeman, P. Chanfreut Palacio
Control Systems Technology group, Department of Mechanical Engineering, Eindhoven University of Technology
Email: e.k.wisse@tue.nl

1 Introduction

Autonomous drone-based inspection is a promising alternative to manual offshore wind turbine inspection, which is costly, time-consuming, and hazardous. Internal blade inspection using two drones carrying a radiographic source and detector requires precise alignment and stable hovering, both of which are strongly affected by the wind disturbances present in offshore environments [1].

2 MPC with Wind-Preview

This research investigates Model Predictive Control (MPC) for quadrotor stabilization [2], leveraging its predictive capabilities to incorporate disturbance forecasts and improve robustness to wind disturbances. To assess the benefit of future wind information, two wind-preview strategies are incorporated into the MPC framework. The first approach assumes exact wind-preview information over a limited time of the prediction horizon. The second approach employs a probabilistic tree-based structure, where multiple possible future wind realizations are considered within the prediction horizon, weighted by their likelihood. Both methods aim to improve robustness against wind uncertainty by enabling anticipatory control actions.

3 Results

Simulation results show that incorporating wind-preview information enables the controller to anticipate disturbances before they affect the system, resulting in reduced position errors. The most significant performance improvements are achieved with a short wind preview, as even limited future wind information leads to substantial error reduction, whereas additional preview length yields diminishing gains. The robustness of the wind-preview strategy is further evaluated under noisy and delayed wind measurements; although overall tracking error increases in these scenarios, wind preview remains beneficial. The tree-based strategy shows potential, especially in environments with systematic wind patterns; however, this comes at the cost of increased computational complexity and a strong dependence on the accuracy of the assumed wind probability distribution.

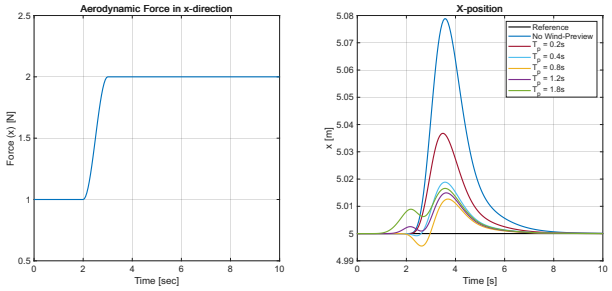

Figure 1: Quadrotor position response with exact wind-preview for varying wind-preview lengths.

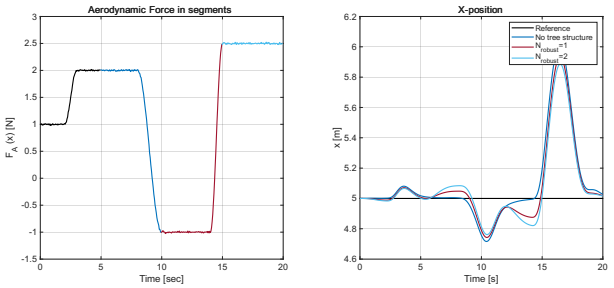

Figure 2: Quadrotor position response obtained with the tree-based structure using a probability distribution defined over four segments.

4 Future Work

Future research steps for MPC with Wind-Preview include (i) verification using real-world wind data, (ii) the development of realistic sensing strategies for acquiring and transmitting future wind information to the quadrotor, (iii) the construction of reliable methods for generating probabilistic wind forecasts, and (iv) experimental validation in real-world environments to assess the practical feasibility of the proposed control strategy.

References

[1] B. Meere, E.Torta, and D. Antunes, *Autonomous Aerial Radiography Inspection of Wind Turbine Blades*. 2024.

[2] H.Merabti, I. Bouchachi, and K. Belarbi, *Nonlinear Model Predictive Control of Quadcopter*. IEEE, 2015.

Explainable Model Predictive Control: HVAC Case Study

Boulaid Boulkroune*, Branimir Mrak, Bruno Depraetere
Flanders Make, Belgium. E-mail: boulaid.boulkroune@flandersmake.be

1 Motivation

Model Predictive Control (MPC) is a powerful optimisation-based control approach for MIMO systems, where it can balance competing objectives—such as tracking accuracy and operational cost—while explicitly enforcing system and input constraints. Historically, the main deterrents to the widespread adoption of MPC have been the availability of sufficiently accurate system models and limitations in computational hardware. In recent years, technological advances have significantly reduced these barriers. However, a lack of explainability remains a significant obstacle to adoption, similar to challenges observed in the deployment of artificial intelligence systems. MPC can generate complex control actions that require explanation to different stakeholders, including control experts, engineers, and end users. This challenge is further amplified in MPC approximations that incorporate artificial intelligence or data-driven models, which often reduce transparency and increase the difficulty of interpreting control decisions.

2 Approach

Several explainability methods have been proposed in the literature, which can be broadly categorized into model-based approaches—leveraging MPC structure and system models (e.g., regularization, primal–dual feasibility analysis, etc.) and black-box approaches (e.g., SHAP[1]).

In this work, a data-based analysis is proposed to explain the closed-loop behavior of the MPC by linking changes in system features to changes in the control input over time. The method is based on the Pruned Exact Linear Time (PELT) algorithm, which efficiently detects multiple change points by minimizing a penalized segmentation cost. The procedure starts by applying PELT separately to each measured features and to the control input, to identify time instants at which their statistical properties, such as the mean level, change significantly. These individual change points are then combined into a single unified set, which divides the full trajectories into a common sequence of segments shared by all variables. For each segment, feature-wise summaries are computed and stored, including: segment identifier, start time, end time, mean value in the preceding segment, current segment mean, absolute magnitude of change, and percentage change with respect to the previous segment. The direction of change is then categorized into qualitative labels such as stable, high increase, increase, small increase, high decrease, decrease, and small decrease. This categorization enables a compact and interpretable description of how each feature and the control action evolve from one segment to the next.

The proposed methodology is demonstrated on an HVAC system with integrated renewable energy sources, specifically regulating the temperature of a single room equipped with a photovoltaic panel and a heat pump under time-varying reference profiles, fluctuating electricity prices, and outdoor temperatures. MPC is applied to track the temperature setpoint while minimizing energy costs and maximizing solar energy utilization.

3 Results on HVAC case and conclusions

The PELT-based explanation divides the trajectories into segments and describes both each segment and the transitions between them. The results indicate that, within a segment, the control action is driven by simultaneous changes in several variables, making it hard to see which one dominates. To quantify each variable's influence on the change in control between segments, a simple linear regression model might be used. A consistency-check block can then verify that these inferred relationships remain coherent across segments and scenarios. Furthermore, incorporating information from the available models, for example, the dynamic model equations (CasADi) used in the MPC, could help to further improve the explainability.

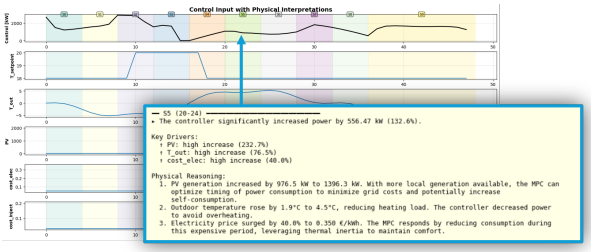

Figure 1: PELT-based explainability approach

References

[1] A. Pathaka, B. Srinivasanb, and R. Srinivasana, "Enhancing Human Operators Trust in MPC Controllers Through an Explainable AI Methodology", Comput. Aided Chemical Eng., vol. 53, pp. 1969–1974, 2024.

Acknowledgements: This research was supported by Flanders Make, the Flemish strategic research centre for the manufacturing industry, within the framework of the CTO2025-03 Human-centric AI project.

A Model Predictive Track Keeping Pilot

Amirreza Haqshenas, Jan Swevers
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium
amirreza.haqshenasmojaveri@kuleuven.be

1 Introduction

Inland waterways expand over 31,000 km, connecting numerous cities and industrial centers across Europe. Navigating through these waterways is a demanding task and can lead to efficiency and reliability challenges. One way to reduce navigational complexity is to equip vessels with additional safety and autonomous features [1, 2].

2 Motivation

Modern inland vessels already benefit from a range of automation technologies that support increasingly autonomous operation. Among these systems is the track keeping pilot, which works alongside the autopilot to keep the vessel on a predefined track. The track keeping pilot computes the required turning rate command, which is then translated by the autopilot into a rudder angle to steer the vessel. However, this control structure does not explicitly account for the vessel's dynamic characteristics [3]. As a result, the track keeping pilot may generate commands that the autopilot is unable to execute, potentially leading to safety critical situations such as grounding or collision.

3 Approach

To address these shortcomings, this research focuses on integrating the vessel's dynamic constraints directly into the Track Keeping Pilot through a nonlinear model predictive control framework. The proposed approach consists of four main components: path smoothing, Kalman filtering, model identification, and model predictive path following. Fig. 1 plots the schematic of the proposed solution.

Path smoothing ensures an even distribution of given waypoints utilizing a predefined distance and cubic interpolations. Kalman filter generates synchronized state estimates from the obtained measurements for both control and identification algorithms. Furthermore, These estimates are employed for dead reckoning. Assuming a deep water operation of the vessel and away from strong currents and wind, the model identification uses the Nomoto model as backbone to map the relationship between the input command and turning rate. It identifies the Nomoto parameters based on a Recursive Least Square algorithm.

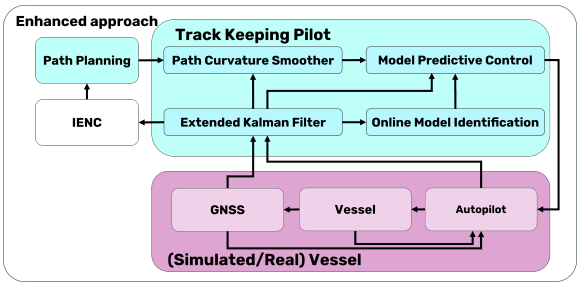

Figure 1: The structure of the proposed solution

This model, combined with the rigid body's planar motion formulas, creates the maneuvering model for an inland vessel. The track keeping task comprises a constrained nonlinear model predictive control based path following control relying on the worked out maneuvering model in the model identification subsystem and calculates the required turning rate as input for the vessel's autopilot.

This research provides a comprehensive approach to enhance the safety and efficiency of inland water navigation through integrating surrounding information with track keeping pilot. The solution has been validated through simulations, comparison with experimental data and implementation on a full scale class VI/a inland vessel in various rivers and canals under different loading condition and water levels. This ongoing work promises significant enhancements in autonomy and safety of inland waterways transportation.

References

[1] A. Dhyani and et al, "Autobargesim: Matlab(r) toolbox for the design and analysis of the guidance and control system for autonomous inland vessels," 2025.

[2] R. Saha and et al, "Autonomous inland waterway transport for a safer and sustainable tomorrow: The autobarge project," *Societal Impacts*, vol. 7, p. 100161, 2026.

[3] A. Haqshenas M. and et al, "Model identification and path following for an inland vessel using ienc data," in *2024 18th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, pp. 846–851, 2024.

Acknowledgment: This research is supported by the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie grant agreement No 955.768 (MSCA-ETN AUTOBarge). This publication reflects only the authors' view, exempting the European Union from any liability. https://etn-autobarge.eu/

On Local Deformations of Explicit MPC Regions Under Parametric Uncertainty

Carlos J. G. Rojas, Valentina Breschi, Leyla Özkan
Control Systems Group, Electrical Engineering Department
Eindhoven University of Technology
Email: c.j.gonzalez.rojas.@tue.nl

1 Introduction

Explicit model predictive control (eMPC) enables real-time implementation of constrained optimal control by computing offline a piecewise-affine law defined over a partition of the state space into critical regions. These partitions and the associated control laws are typically computed for a nominal model. However, in practice, uncertainty can alter the number and boundaries of the critical regions, affecting feasibility and control performance. In this work, we propose a method for computing the first-order effects of model parametric uncertainty and obtaining approximate control laws.

2 System Description

Consider a discrete-time linear system described by the state-update equation:

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k, \quad x_k \in X_c, \quad u_k \in U_c \tag{1}$$

where x_k, u_k are the state and control variables, θ represents a time invariant scheduling parameter and $X_c \subset \mathbb{R}^{n_x}, U_c \subset \mathbb{R}^{n_u}$ are the constraints sets. In this work, we examine static parametric dependencies that typically arise from material properties or design parameters of systems. For simplicity, we investigate a scalar uncertainty description on the dynamic matrix:

$$A(\theta) = A_0 + \theta\Delta A, \quad \theta \in (-1,1) \tag{2}$$

with A_0 obtained by re-centering and ΔA by re-scaling the limits of the varying physical parameter.

3 Problem Formulation

We formulate an eMPC for a general regulation (or tracking) problem:

$$J_0^*(x(t)) = \min_{U_0} \quad J_0(x(t), U_0) = p(x_N) + \sum_{k=0}^{N-1} q(x_k, u_k) \tag{3}$$
$$\text{s.t.} \quad \textit{Eq. (1)},$$

where $U_0 = \{u_0, ..., u_{N-1}\}$, $x_0 = x(t)$ is the initial state, and p and q are functions defined according to the control requirements, as well as the chosen MPC cost (here assumed quadratic). In eMPC the optimization problem given by Eq. (3) is solved offline as a multiparametric quadratic program (mpQP), where the system state, references, and previous inputs can be treated as parameters. Lifting the system dynamics over the prediction horizon results in condensed matrices that depend nonlinearly on θ through repeated matrix products. Consequently, variations in A cannot be treated as affine parameters of a standard mpQP. Our goal is to analyze the first order effects of θ in the critical regions and provide an approximate explicit control law in terms of it.

4 Effects of Model Parametric Uncertainty

The lifting of the system dynamics is written as:

$$X = S_x x_0 + S_u U_0, \tag{4}$$

with S_x and S_u containing the powers A^p with $p = \{1, ..., N\}$. We propose using the Fréchet derivative L_f [1] to have a linear approximation of the powers depending on θ:

$$A^p \approx {A_0}^p + \theta L_f(\Delta A, p) \tag{5}$$

Next, we decompose the lifting matrices $S_x = \tilde{S}_x + \theta\Delta S_x$ and $S_u = \tilde{S}_u + \theta\Delta S_u$, by grouping the terms that depend on A_0 and ΔA and we plug these expressions into the QP (cost and constraints). The cost function obtained based on this procedure depends quadratically on the parameter θ. Therefore, the second order terms should be truncated if an explicit affine expression is desired. We solve this mpQP using the KKT conditions and the region-based approach for the construction of the critical regions [2]. As a simulation example, we apply the described approach to the system given in Example 12.3 in [3]. Figure 1 shows the difference in the critical regions when the element (2,2) ($\rho \in [0.7, 1.3]$) in the dynamic matrix has a parametric uncertainty with $\theta = -0.33$ versus the exact model.

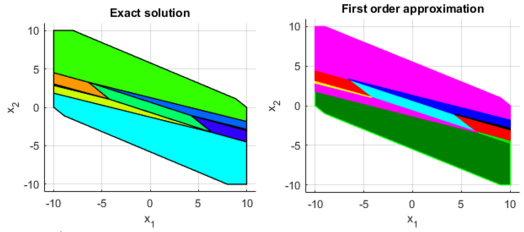

Figure 1: Effects of Parametric Uncertainty on Critical Regions

References

[1] A. H. Al-Mohy and N. J. Higham, "Computing the Fréchet derivative of the matrix exponential, with an application to condition number estimation," *SIAM Journal on Matrix Analysis and Applications*, 2009.

[2] M. Kvasnica, B. Takacs, J. Holaza, and S. Di Cairano, "On region-free explicit model predictive control," in *Proc. IEEE Conf. on Decision and Control (CDC)*, 2015.

[3] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*, Cambridge University Press, 2017.

Real-time fleet rebalancing: a comparative study of MPC and economic MPC

Arthur Mélot
ICTEAM, UCLouvain
Louvain-la-Neuve, Belgium
Email:arthur.melot@uclouvain.be

Gianluca Bianchin
ICTEAM, UCLouvain
Louvain-la-Neuve, Belgium
Email:gianluca.bianchin@uclouvain.be

1 Introduction

Rapid urbanization drives the adoption of Mobility on Demand (MoD) systems, which provide shared transportation adapted to real-time user requests. However, the differences in demand patterns imply the need for active fleet rebalancing to ensure vehicle availability across the network [1]. While Carron et al. [2] addressed this using scalable Model Predictive Control (MPC) on an Autonomous MoD fluid model, this work implements Economic MPC (EMPC) and compares its performance against their baseline.

2 EMPC algorithm

EMPC suggests taking an economic objective as the objective function of the optimization problem instead of the deviation from a known setpoint as in MPC. The main advantage of this method is that it does not require any previous knowledge about an existing setpoint to track, but just an economic criterion to optimize [3].

3 Application and first results

We use the same discrete-time disturbed LTI model as in [2] (Equation 2), where the state vector $\mathbf{x}$ and the input vector $\mathbf{v}$ are decomposed as $\mathbf{x}_k = \begin{bmatrix}\mathbf{c}_k & \mathbf{w}_k & \mathbf{s^u}_k & \mathbf{s^r}_k\end{bmatrix}^T$ and $\mathbf{v}_k = \begin{bmatrix}\mathbf{u}_k & \mathbf{r}_k\end{bmatrix}^T$ where $\mathbf{c}_k$ is a vector representing all the customers waiting on each edge, $\mathbf{w}_k$ the idle vehicles at each node, $\mathbf{s^u}$ and $\mathbf{s^r}$ the vehicles (resp. with and without customers) currently traveling each edge, $\mathbf{u}_k$ the vehicles with customer that are starting their trip on each edge, $\mathbf{r}_k$ the rebalancing vehicles sent through each edge at time k. We define the economic objective as minimizing the number of customers waiting in the system and the total travel time of empty vehicles. Finally, we will also penalize deviation from steady-state dynamics at the end of the prediction horizon. Mathematically, this leads to the following optimization problem:

$$\min_{\mathbf{v}} \sum_{i=0}^{N-1} \left[\|\mathbf{c}_{k+i}\|^2 + \mathbf{r}_{k+i}^T \tilde{T} \mathbf{r}_{k+i}\right] + \sigma_N^2 \qquad (1)$$

$$\text{s.t. } \forall i = 0,\dots,N-1$$

$$\mathbf{x}_{k+i+1} = \mathscr{A}\mathbf{x}_{\mathbf{k+i}} + \mathscr{B}\mathbf{v}_{\mathbf{k+i}} + \mathscr{W}\lambda$$

$$\mathbf{x}_{k+i}, \mathbf{v}_{k+i} \in \mathscr{X}, \mathscr{U}$$

with $\sigma_N = \|(I - \mathscr{A})\mathbf{x}_N - \mathscr{B}\mathbf{v}_N - \mathscr{W}\lambda\|$ quantifying the deviation from steady state dynamics at the end of the prediction horizon. Simulating the two algorithms on a small system with 5 nodes and 8 edges gives the following preliminary results

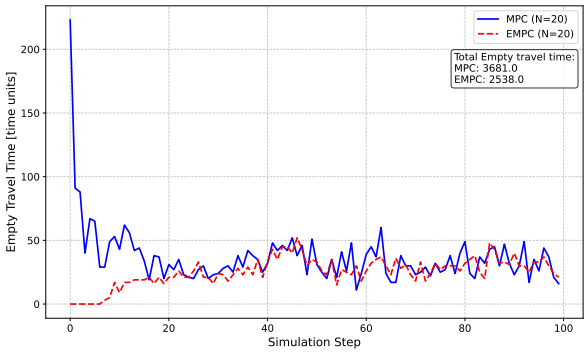

Figure 1: Comparison of total distance traveled by rebalancing vehicles at each time step of the simulation

where we observe that, despite similar performances at steady-state, EMPC provides a lower total empty travel time, implying less kilometers traveled and thus less fuel wasted during the transient than MPC (Figure 1). The initial state is an equilibrium point computed in [2] disturbed with 15 additional customers on the first edge for both algorithms. Further analysis may focus on different objectives and additional metrics related or not to the application (e.g. the average waiting time of customers on each edge, execution time comparison, or even total fuel/electricity consumption).

References

[1] F. Viadero-Monasterio et al., "Low-Cost vehicle rebalancing control for an autonomous mobility on demand system," 2025. doi: 10.1016/j.jfranklin.2025.108333.

[2] A. Carron et al., "Scalable Model Predictive Control for Autonomous Mobility-on-Demand Systems," 2021, doi: 10.1109/TCST.2019.2954520

[3] J. B. Rawlings et al., "Fundamentals of economic model predictive control," 2012, doi: 10.1109/CDC.2012.6425822.

Dual MPC for Active Learning of Nonparametric Uncertainties

Tren Baltussen, Maurice Heemels, Alexander Katriniok
Control Systems Technology, Department of Mechanical Engineering
Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands
Email: {t.m.j.t.baltussen, m.heemels, a.katriniok}@tue.nl

1 Introduction

One of the reasons for the high interest in model predictive control (MPC), in particular for safety-critical systems, is its ability to systematically handle constraints on the states and inputs. Yet, the performance and safety of MPC heavily depend on the accuracy of the model of the system. In various applications, the system dynamics may be uncertain or partially unknown, and need to be identified online. This requires sufficiently informative data, and hence, the excitation of the system. However, these system excitations may conflict with the control objective and safety requirements. Balancing these objectives, identification and control, is the focus of *dual control* [1]. Therefore, studying dual control in the context of MPC is of high relevance for simultaneous identification and control of safety-critical systems.

2 Problem Formulation

We consider a discrete-time control system $\mathscr{S}$ of the form

$$x_{k+1} = f(x_k, u_k) + \underbrace{Bg(x_k, u_k) + v_k}_{=:w_k}, \tag{1}$$

where $x_k \in \mathbb{R}^{n_x}$ and $u_k \in \mathbb{R}^{n_u}$ denote the state and input vector, respectively, at time step $k \in \mathbb{N}$. We assume to have access to exact measurements of the state x_k and input u_k. The function $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ is a nominal model of the system dynamics. The function $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_g}$ represents the unknown system dynamics that have to be identified. We assume f and g are continuously differentiable and can be generally nonlinear. The matrix $B \in \mathbb{R}^{n_x \times n_g}$ is a known matrix with full column rank. Lastly, $v_k \in \mathbb{V} \subset \mathbb{R}^{n_x}$ is a bounded disturbance at time $k \in \mathbb{N}$. The system is subject to state constraints $x_k \in \mathbb{X} \subset \mathbb{R}^{n_x}$ and input constraints $u_k \in \mathbb{U} \subseteq \mathbb{R}^{n_u}$, $k \in \mathbb{N}$. The uncertainty $w_k := Bg(x_k, u_k) + v_k$ is assumed to satisfy $w_k \in \mathbb{W} \subset \mathbb{R}^{n_x}$, for all $x \in \mathbb{X}$, $u \in \mathbb{U}$ and $k \in \mathbb{N}$, where $\mathbb{W}$ is a bounded set and $0 \in \mathbb{W}$.

We consider the problem of designing a dual controller that balances system identification and control. Specifically, we aim to safely regulate the system $\mathscr{S}$ (1) to a desired setpoint while actively learning the unknown dynamics $g(x, u)$.

3 Contributions

We present a GP-MPC that *actively* explores the state space and learns the unknown system dynamics while ensuring robust constraint satisfaction. The main contributions of this work are as follows. 1) We propose an explicit dual MPC that actively learns system dynamics using the active learning framework from [2]. 2) We prove recursive feasibility of the actively learning MPC by incorporating a *contingency horizon* [3]. 3) We demonstrate the effectiveness of the proposed dual MPC through a numerical study on a nonlinear system with nonparametric, unknown dynamics. In addition, we compare our dual MPC with a robust MPC (RMPC), a passively learning dual MPC [3], and the active learning framework from [2] with a single horizon.

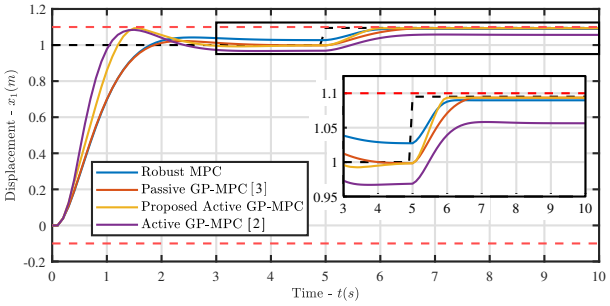

Figure 1: Simulation of nonlinear mechanical system.

4 Results

The results of the simulation study are shown in Fig. 1. The passively learning GP-MPC [3] with a contingency horizon successfully regulates the system to the desired setpoint while adapting to nonparametric unmodeled dynamics, where learning occurs only as a side effect of control. In contrast, the proposed actively learning GP-MPC explicitly excites the system in a safe manner to learn the unknown dynamics and then exploits the acquired model to achieve similar performance. We expect the actively learning dual MPC will outperform the passive MPC in more complex problems that require exploration to learn the unknown dynamics.

References

[1] A. Feldbaum, "Dual control theory problems," *IFAC Proceedings Volumes*, vol. 1, no. 2, pp. 541–550, 1963.

[2] R. Soloperto, J. Kohler, and F. Allgöwer, "Augmenting MPC Schemes With Active Learning: Intuitive Tuning and Guaranteed Performance," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 713–718, Jul. 2020.

[3] M. Geurts, T. Baltussen, A. Katriniok, and M. Heemels, "A contingency model predictive control framework for safe learning," *IEEE Control Systems Letters*, vol. 9, pp. 1075–1080, 2025.

Towards cooperative model predictive control in tokamaks

Hari Prasad Varadarajan[1,2,†], Domenico Frattolillo[3], Marco De Baar[1,2], Dinesh Krishnamoorthy[4]

[1] DIFFER - Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands

[2] Eindhoven University of Technology, Control Systems Technology, Eindhoven, the Netherlands

[3] University of Naples Federico, Naples, Italy

[4] Norwegian University of Science and Technology, Trondheim, Norway

[†] h.varadarajan@differ.nl

1 Introduction

Tokamaks are among the more mature pathways towards the realization of magnetically confined fusion energy. Sustaining the extreme temperatures and densities of the plasma required for fusion reactions presents substantial control challenges across multiple domains. Failure to adequately regulate these processes can lead to degraded performance or, in extreme cases, damage to the device. Moreover, future reactors will be operated close to their operational and safety limits, leaving limited margins for control errors.

Significant progress has been made in the development of domain-specific controllers for tokamak plasmas, with magnetic shape control [1] and density [2] control typically treated as independent problems. While this modular approach has practical advantages, it neglects the fact that changes in magnetic geometry directly affect plasma transport properties and confinement. Such cross-domain interactions are often either ignored or treated implicitly as disturbances, rather than being modeled and exploited in the control design. As a consequence, controllers may at times act against each other, or fail to leverage the full actuation potential available. In some cases, it may be advantageous to employ actuators from one domain to assist in regulating quantities traditionally assigned to another. As a first step towards achieving an integrated control framework for tokamaks, this work focuses on an underexplored cross-domain interaction: the effect of plasma volume on the plasma electron density profile. We develop a coupled model that explicitly captures the influence of magnetic shaping actuators on core density evolution. Building on this model, we investigate the use of magnetic shaping actuators as virtual actuators [3] for density regulation and assess the benefits of a coupled model predictive control strategy relative to conventional domain-separated approaches.

2 Methods

A centralized model predictive control (MPC) scheme formulated in velocity form is designed for the coupled system. The controller optimizes both fueling and magnetic coil voltage trajectories while explicitly incorporating the coupling term between magnetic and density dynamics. To assess the impact of explicitly accounting for inter-domain coupling, the coupled MPC is compared against a baseline in which the magnetic and density subsystems are regulated independently using domain-specific controllers that neglect the coupling term. This representation abstracts the underlying plasma physics and is intended to capture the dominant coupling relevant for control design. It must be noted that although the controller plant model is linear, the closed loop simulation is performed with a nonlinear plant, thereby demonstrating the controller's robustness to plant-model mismatch.

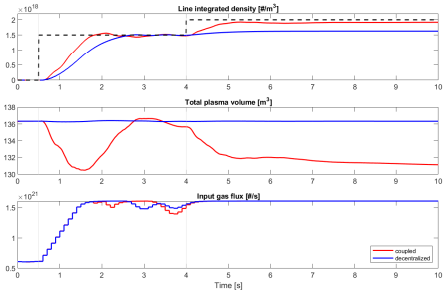

Figure 1: Coupled MPC vs Domain based decentralized MPC: the coupled MPC is able to track higher density references, with a saturated fueling valve, as it can use shape changes to affect the density. The steady state errors observed reflect the higher relative weighting for the density error compared to the shape(volume) error.

References

[1] G. De Tommasi, "Plasma magnetic control in tokamak devices," *Journal of Fusion Energy*, vol. 38, no. 3, pp. 406–436, 2019.

[2] T. O. S. J. Bosman, "Density profile and heat-exhaust control in nuclear fusion devices: Development of control methods in view of the operation of fusion power plants," 2025.

[3] L. E. di Grazia, E. Fable, M. Mattei, M. Siccinio, F. Maviglia, and H. Zohm, "Burn control strategies using plasma elongation in demo," *Fusion Engineering and Design*, vol. 215, p. 114976, 2025.

Comparison of Systems Across Time Domains

Armin Pirastehzad Bart Besselink

Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence
University of Groningen
Email: a.pirastehzad@rug.nl,b.besselink@rug.nl

1 Introduction

Modern engineering systems often operate in accordance with intricate specifications, such as those expressed by temporal logic formulae. For a continuous-time system, these specifications are usually enforced by controllers that are synthesized for a discrete-time model of the original system. Such controllers, however, guarantee specifications only at the sampled time instances rather than all times. In fact, it is not clear to what extent the specifications are satisfied/violated in between these instances. Motivated by this, we develop a *formal* framework for comparison of the controlled input-state behavior of a continuous-time system with that of its discrete-time model. This framework can be utilized for control synthesis to 1) guarantee specifications at sampled time instances and 2) measure the extent of (possible) specification violation in between these instances.

2 System Interpolation

Consider the continuous-time system

$$\mathbf{\Sigma}_c : \dot{x}_c(t) = A_c x_c(t) + B_c u_c(t),$$

with state $x_c \in \mathbb{R}^n$ and control input $u_c \in \mathbb{R}^m$. In correspondence to (2), we consider the discrete-time system

$$\mathbf{\Sigma}_d : x_d(i+1) = A_d x_d(i) + B_d u_d(i),$$

with state $x_d \in \mathbb{R}^n$ and $u_d \in \mathbb{R}^m$.

We conceive a method that compares the input-state behavior of $\mathbf{\Sigma}_c$ with that of $\mathbf{\Sigma}_d$. To accomplish this, we first perform *piece-wise polynomial* interpolation to compare continuous-time signals with discrete-time ones.

Definition 1. For a sampling time $\tau \in \mathbb{R}_{>0}$, a polynomial degree $N \geq 1$, and integers $\ell, p \in \mathbb{Z}_{>0}$, a continuous-time signal $s_c : [0, \ell\tau] \to \mathbb{R}^p$ is said to be an N-th order interpolation of a discrete-time signal $s_d : [\![0, \ell]\!] \to \mathbb{R}^p$ if there exist polynomial functions $P_0, P_1, \cdots, P_{\ell-1} \in \mathbb{P}_N^p$ such that for all $i \in [\![0, \ell-1]\!]$, we have that

$$s_c(i\tau) = s_d(i),$$
$$s_c\big((i+1)\tau\big) = s_d(i+1)$$

and

$$\forall t \in [0, \tau] : \quad s_c(i\tau + t) = P_i(t).$$

We denote by $\mathbb{I}_N^\tau(s_d)$ the set of all N-th order interpolations of a given discrete-time signal $s_d : [\![0, \ell]\!] \to \mathbb{R}^p$.

We then compare the behavior of $\mathbf{\Sigma}_c$ and $\mathbf{\Sigma}_d$ by inspecting whether $\mathbf{\Sigma}_c$ is able to generate piece-wise polynomial interpolations of the input-state behavior of $\mathbf{\Sigma}_d$.

Definition 2. For a sampling time $\tau \in \mathbb{R}_{>0}$ and an integer $N \geq 1$, a continuous-time system $\mathbf{\Sigma}_c$ is said to be an N-th order interpolator of a discrete-time system $\mathbf{\Sigma}_d$ if for any integer $\ell \in \mathbb{Z}_{>0}$, any initial condition $x_0 \in \mathbb{R}^n$, and any discrete-time input $u_d : [\![0, \ell]\!] \to \mathbb{R}^m$, there exists a continuous-time input $u_c \in \mathbb{I}_N^\tau(u_d)$ such that

$$x_c(\cdot; x_0, u_c)\big|_{[0,\ell\tau]} \in \mathbb{I}_N^\tau\Big(x_d(\cdot; x_0, u_d)\big|_{[\![0,\ell]\!]}\Big). \tag{2}$$

The notion of system interpolation can be characterized as a subspace inclusion, in terms of the parameters of $\mathbf{\Sigma}_c$ and $\mathbf{\Sigma}_d$. This characterization then enables discretization of $\mathbf{\Sigma}_c$ into $\mathbf{\Sigma}_d$ such that $\mathbf{\Sigma}_c$ is an N-th order interpolator of $\mathbf{\Sigma}_d$.

Given a specification, system interpolation enables a *hierarchical* control scheme that splits synthesis into the following steps:

1. Construction of $\mathbf{\Sigma}_d$ such that $\mathbf{\Sigma}_c$ is an N-th order interpolator of $\mathbf{\Sigma}_d$;
2. Synthesis of a discrete-time control u_d such that $\mathbf{\Sigma}_d$ satisfies the specification;
3. Design of a continuous-time control $u_c \in \mathbb{I}_N^\tau(u_d)$ that establishes (2).

This hierarchical control scheme guarantees that $\mathbf{\Sigma}_c$ satisfies the specification at sampled time instance, while it gives a measure of the extent to which $\mathbf{\Sigma}_c$ may violate the specification in between these instances.

3 Conclusion

We introduced and characterized the notion of system interpolation, which determines whether a continuous-time dynamics is able to generate piece-wise polynomial interpolations of input-state trajectories of a discrete-time one. After conducting discretization according to system interpolation, we proposed a hierarchical control scheme to enforce specifications.

On the Effect of Actuator Dynamics on the Stability of Sampled Data Systems

Pournami Padmalatha[1,2]
Wim Michiels [1]
[1]Department of Computer Science, NUMA, KU Leuven, B-3001 Leuven, Belgium
pournami.padmalatha@kuleuven.be
wim.michiels@kuleuven.be

Laurentiu Hetel[2]
Daniel Denardi Huff[2]
[2]UMR 9189 CRIStAL, CNRS, Université de Lille, Centrale Lille, F-59000 Lille, France
laurentiu.hetel@univ-lille.fr
daniel.denardi-huff@univ-lille.fr

1 Abstract

The effect of actuator dynamics on the stability of Sampled Data Systems (SDS) is studied. The system has a Linear Time Invariant (LTI) continuous time plant, actuator and a digital controller with aperiodic sampling. Such a system merits an investigation into how the stability is affected due to the finite response time of the actuator and the sampling. We probe into how slow the actuator dynamics can be and how high the sampling period can be before the system loses its stability. Building upon prior works that treat sampling as a perturbation, we extend the analysis to include non-ideal actuator dynamics as an additional perturbation. The Small Gain Theorem (SGT) from the robust control theory framework is used to analyze the problem.

2 Problem Formulation

Consider the LTI system represented by

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{A}_s\mathbf{x}(t)+\mathbf{B}_s\mathbf{v}(t) \\ \varepsilon\dot{\mathbf{z}}(t) &= \mathbf{A}\mathbf{z}(t)+\mathbf{B}\mathbf{u}(t) \\ \mathbf{v}(t) &= \mathbf{C}\mathbf{z}(t) \end{aligned} \tag{1}$$

where $\mathbf{x}\in\mathbf{R^n},\mathbf{z}\in\mathbf{R^m}$, matrices $\mathbf{A_s},\mathbf{B_s},\mathbf{A},\mathbf{B},\mathbf{C}$ are of appropriate dimensions and $\mathbf{u}$ is the control signal. $\mathbf{v}$ is the actuator output. Matrices $\mathbf{A_s},\mathbf{B_s}$ define the nominal continuous time plant dynamics and $\mathbf{A},\mathbf{B}$ the actuator dynamics. The latter evolves on a fast scale, characterized by the singular perturbation parameter, ε. As $\varepsilon\to 0$, the actuator dynamics becomes infinitely fast and approaches the ideal instantaneous actuator. For a set of ordered sampling instants $\{t_k\}_{k\geq 0}, t_0=0$, let the actuator input $\mathbf{u}(t)$ be composed of linear sampled state feedback as

$$\mathbf{u}(t)=\mathbf{K}\mathbf{x}(t_k),\quad t\in[t_k,t_{k+1}) \tag{2}$$

where $\mathbf{K}$ is the gain matrix. The sampling period $h_k=t_{k+1}-t_k$ satisfies $0\leq\underline{h}\leq h_k\leq\overline{h}\ \forall k>\mathbf{N}$. The matrices $\mathbf{A}$ and $\mathbf{A}_{cl}=\mathbf{A_s}+\mathbf{B_s}\mathbf{K}$ are assumed to be Hurwitz and the actuator is assumed to have unity steady state gain ($-\mathbf{C}\mathbf{A}^{-\mathbf{1}}\mathbf{B}=\mathbf{I}$). The objective is to obtain an upper bound on ε for a given maximum sampling period $\overline{h}$ (or vice versa) for which robust stability is preserved. The following transfer function definitions are required to state the result

$$G_n(s)=-\mathbf{K}\mathbf{A}_{cl}(s\mathbf{I}-\mathbf{A}_{cl})^{-1}\mathbf{B_s}-\mathbf{K}\mathbf{B_s} \tag{3}$$

$$\Delta_1(s)=\mathbf{C}(s\mathbf{I}-\mathbf{A})^{-1}(-\mathbf{A}^{-\mathbf{1}}\mathbf{B}) \tag{4}$$

$$\Delta_2(s)=\mathbf{C}(s\mathbf{I}-\mathbf{A})^{-1}\mathbf{B} \tag{5}$$

3 Result

Theorem 1 *Let γ_1,γ_2 represent the $\mathscr{H}_\infty$ norms of the transfer functions Δ_1, Δ_2 respectively and let $\|G_n\|_{\mathscr{H}_\infty}$ be the $\mathscr{H}_\infty$ norm of G_n. Then the system represented by (1), (2) is stable if*

$$\varepsilon\gamma_1+\frac{2}{\pi}\overline{h}\gamma_2<\frac{1}{\|G_n\|_{\mathscr{H}_\infty}}. \tag{6}$$

Proof. The proof is based on remodeling the system with the finite actuator bandwidth and sampling considered as perturbations to the ideal closed loop system. G_n accounts for the ideal closed loop system and Δ_1,Δ_2 relate to the finite actuator response. The effect of sampling is captured by an operator ([1][2]) with an induced L_2 norm of $\frac{2}{\pi}\overline{h}$. SGT is used to obtain the stability condition (6) of this feedback interconnection system.

4 Conclusion and Future Work

By expressing the stability condition (6) explicitly in ε and $\overline{h}$, we set the stage for future developments in controller synthesis and the study of SDS with more than two time scales.

References

[1] L. Mirkin, "Some remarks on the use of time-varying delay to model sample-and-hold circuits," *IEEE Transactions on Automatic Control*, vol. 52, no. 6, pp. 1109–1112, Jun. 2007.

[2] H. Omran, L. Hetel, J.-P. Richard, and F. Lamnabhi-Lagarrigue, "On the stability of input-affine nonlinear systems with sampled-data control," in *Proc. 2013 European Control Conference (ECC)*, pp. 2585–2590, 2013.

Bounds on Scaled Graphs of Reset Systems

T. de Groot*, W.P.M.H. Heemels*, T. Oomen*,**, S.J.A.M. van den Eijnden*

*Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands

**Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands

{t.d.groot2, W.P.M.H.Heemels, t.a.e.oomen, s.j.a.m.v.d.eijnden}@tue.nl

1 Introduction

Reset control systems, which are systems whose internal state is instantaneously reset according to a reset map based on the state and/or input, are able to overcome fundamental limitations of linear time-invariant control [1]. In Fig. 1a an example of the output of a reset system, subject to a sinusoidal input, is shown. Due to the nonlinear nature of the reset systems, classical graphical design methods based on frequency response functions, such as Bode and Nyquist diagrams, cannot be used. The lack of such intuitive (graphical) design tools for reset systems, limits their adoption in industry. A promising tool to develop intuitive graphical design methods for reset systems is the scaled graph (SG) [2]. The SG of a reset system $\mathfrak{R}$ is a subset of $\mathbb{C}$, defined as

$$\mathrm{SG}(\mathfrak{R}) = \left\{ \frac{\|y\|}{\|u\|} \exp\left(\pm j \arccos \frac{\langle u,y \rangle}{\|u\|\|y\|} \right) \mid y \in \mathfrak{R}u, u \in \mathscr{L}_2 \right\},$$

with $\langle u,y \rangle := \int_0^\infty u(t)y(t)dt$ and $\|u\| = \sqrt{\langle u,u \rangle}$. SGs provide a graphical representation of the input-output behaviour of reset systems in terms of gain and phase, and in that sense generalize the classical Nyquist diagram for linear systems.

2 Problem Statement

A main problem in SG analysis for reset control is calculating the SG. Using the definition directly results in the need to calculate an uncountable number of input-output pairs, which is not possible. This problem is mitigated by constructing over-bounds of the SGs. To this end, we present an efficient method to calculate over-bounds of SGs of reset systems. To investigate the conservatism of the over-bound, we present a lower-bound on the over-bound, based on samples of the SG.

3 Approach

To calculate over-bounds of SGs we exploit the link between dissipativity and SGs explored in [3,4]. We extend the class of storage functions from common quadratic storage functions, as considered in [3, 4], to piecewise quadratic storage functions, with the aim of reducing conservatism. The verification of dissipativity with respect to different supply rates can be done efficiently using linear matrix inequalities (LMIs). Each supply rate maps to a region in $\mathbb{C}$ with a circular boundary containing the SG. Taking the intersections of these regions results in a over-bound $\mathscr{G}$ of the SG.

To construct the under-bounds based on samples of the SG we exploit the fact that $\mathscr{G}$ is constructed as the intersection of circular regions, which implies that $\mathscr{G}$ is hyperbolic convex, see [2]. This is exploited to show that the hyperbolic convex hull of samples of the SG provides a lower-bound on $\mathscr{G}$.

4 Example

Consider an example reset system. In Fig. 1b an over-bound $\mathscr{G}$ is shown together with a set of samples and the resulting under-bound. It can be seen that the outer boundary of $\mathscr{G}$ is close to the under-bound. Thus the over-approximation cannot be significantly improved for this example system by considering the intersection of circular regions.

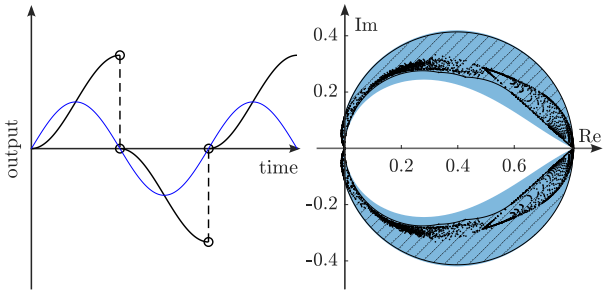

(a) (b)

Figure 1: (a) Example of the output of resetting integrator (black) to an input (blue). (b) Over-approximation of SG (blue region), samples of SG (black points) and hyperbolic convex hull of the samples (black hatched region)

5 Outlook

Future work focuses on extending the methods towards general nonlinear and multivariable systems.

References

[1] O. Beker, C. Hollot, and Y. Chait, "Plant with integrator: an example of reset control overcoming limitations of linear feedback," *IEEE Trans. Autom. Control*, vol. 46.

[2] T. Chaffey, F. Forni, and R. Sepulchre, "Graphical nonlinear system analysis," *IEEE Trans. Autom. Control*, vol. 68.

[3] S. van den Eijnden, T. Chaffey, T. Oomen, and M. Heemels, "Scaled graphs for reset control system analysis," *Eur. J. Control*.

[4] T. de Groot, M. Heemels, and S. van den Eijnden, "A Dissipativity Framework for Constructing Scaled Graphs." arXiv:2507.08411 [math].

Minimal Security Intervals for Resilience of Nonlinear Systems under Impulsive Attacks

Kangping Gao[1,2], Xiaodong Cheng[2], Yishu Wang[1], Jianquan Lu[1]
[1]Southeast University, China
[2]Wageningen University & Research, Netherlands
Emails: kangping.gao@wur.nl, Xiaodong.Cheng@wur.nl, wongyisu@gmail.com, jqluma@seu.edu.cn

1 Introduction

In information-rich or networked environments, impulsive attacks resulting from external disturbances generally compromise system stability [1, 2]. Existing attack models have inherent limitations, since impulsive attacks are difficult to characterize, but can be detected through abrupt changes in system states. The ability of a system to recover from impulsive attacks and maintain stability is referred to as impulsive resilience. In our work, we determine the minimal security interval required between two consecutive attacks to ensure *impulsive resilience*, based on the system states immediately before and after each impulsive instant.

2 Impulsive Resilience and Minimal Security Intervals

Consider the nonlinear system subject to impulsive attacks:

$$\begin{cases} \dot{x}(t) = f(x(t)), & t \geq t_0, t \neq t_k, \\ x(t_k) = g_k(x(t_k^-)), & k \in \mathbb{Z}_+, \end{cases} \tag{1}$$

where $x(t) \in \mathbb{R}^n$ represents the state, and $\dot{x}(t)$ stands for its right-hand derivative. The function $f \in C(\mathbb{R}^n, \mathbb{R}^n)$ is a locally Lipschitz continuous flow. The function $g_k(\cdot)$ captures the instantaneous change caused by an impulsive attack at time t_k, mapping state $x(t_k^-)$ before the impulsive attack to state $x(t_k)$ after the attack. The sequence $\{t_k, k \in \mathbb{Z}_+\}$ specifies unknown impulsive attack moments. We assume that the intensity of impulsive attacks is bounded and that $x(t)$ is right continuous at time t_k, in the sense that $x(t_k) = x(t_k^+)$. The origin is an equilibrium point of (1), ensured by $f(0) = 0$ and $g_k(0) = 0$. The solution of system (1) starting from (t_0, x_0) is denoted by $x(t, t_0, x_0)$ or simply $x(t)$ when no confusion arises.

The system (1) is said to be *impulsively resilient* if, for any $\varepsilon > 0$ and $t_0 \geq 0$, there exists some $\delta = \delta(\varepsilon, t_0)$ such that whenever the initial state satisfies $|x_0| \leq \delta$, the solution satisfies $|x(t)| \leq \varepsilon$ for all $t \geq t_0$, despite the occurrence of impulsive attacks, provided that the inter-attack intervals are sufficiently large to allow the system to recover from prior impulses.

In our work, we have demonstarted that if there exist functions $\gamma_1, \gamma_2 \in \mathscr{K}_\infty$, constants $b > 0$, $\beta_k > 0$ and a locally Lipschitz continuous function $V : \mathbb{R}^n \to \mathbb{R}^+$ for any $x \in \mathbb{R}^n$ such that $\gamma_1(\|x\|) \leq V(x) \leq \gamma_2(\|x\|)$, $V(g_k(x)) \leq e^{\beta_k} V(x)$, and $D^+V[x(t)]f(x(t)) \leq -bV(x(t)), t \neq t_k$ for $k \in \mathbb{Z}_+$. Then, there exist constants $c_k > 0$ for $k \in \mathbb{Z}_+$ satisfying

$$t_k^- = \inf\{t \geq t_{k-1} : V(x(t)) \leq e^{-c_k} V(x(t_{k-1}))\}, \tag{2}$$

and system (1) can achieve impulsive resilience provided that the condition $c_k \geq \beta_k$ holds. Moreover, the minimal security intervals between two consecutive impulsive attacks satisfy $t_k - t_{k-1} \geq \beta_k / b, k \in \mathbb{Z}_+$.

To better understand impulsive resilience, we further illustrate it by considering three system states: the post-attack state following the previous impulsive attack $x(t_{k-1})$, the pre-attack state immediately before the next impulse $x(t_k^-)$, and the post-attack state immediately after the impulse $x(t_k)$, as shown in Fig. 1. When the interval between consecutive impulsive attacks equals the minimal security interval, the decay factor c_k of the pre-attack system state equals the impulsive intensity parameter β_k.

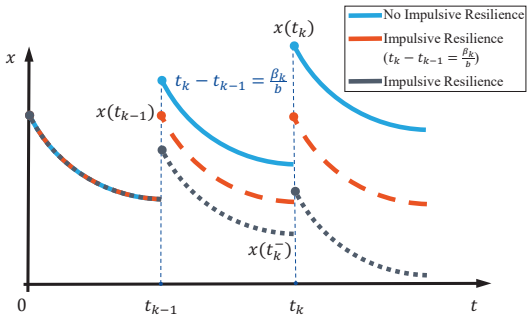

Figure 1: Illustrations of impulsive resilience.

References

[1] X.-K. Liu, C. Wen, Q. Xu, and Y.-W. Wang, "Resilient control and analysis for DC microgrid system under DoS and impulsive FDI attacks," *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 3742–3754, 2021.

[2] H. Su, Y. Wu, and W. X. Zheng, "Distributed observer for LTI systems under stochastic impulsive sequential attacks," *Automatica*, vol. 159, p. 111370, 2024.

Incremental input-to-state stability for a class of discontinuous piecewise linear systems

M. J. Erlandsen, W. P. M. H. Heemels, and S. J. A. M. van den Eijnden
Control Systems Technology, Eindhoven University of Technology, The Netherlands
E-mail: {m.j.erlandsen, m.heemels, s.j.a.m.v.d.eijnden}@tue.nl

1 Introduction

The notion of *incremental input-to-state stability* (δISS) provides a powerful tool for the analysis and design of nonlinear systems. In this paper, we present conditions under which a class of discontinuous piecewise linear (PWL) systems is δISS. The conditions arise from a common quadratic incremental Lyapunov function approach, and are formulated in terms of numerically tractable matrix inequalities. We derive an insightful frequency-domain counterpart of the matrix inequalities, which reveals that, essentially, the class of PWL systems that is considered in this paper is δISS, if the underlying linear subsystems possess a strong passivity property. PWL systems have become important within a wide range of applications, including mechanical systems with unilateral constraints and electrical circuits.

2 System description and main result

Consider a PWL system of the form

$$\dot{x} = A_i x + Bw \quad \text{if } x \in \mathscr{S}_i, i \in \{1,2\}, \tag{1}$$

where $x \in \mathbb{R}^n$ is the state, $w \in \mathbb{R}^m$ is the input, and where we assume that $A_1 = A_2 - bc^\top$, for some vectors $b, c \in \mathbb{R}^n$. The sets $\mathscr{S}_i$, $i \in \{1,2\}$, in (1) that determine the switching are defined through two linearly independent vectors $v, c \in \mathbb{R}^n$, as indicated by Figure 1. Examples of such systems include generic Lur'e type of systems, and control systems with variable-gain elements and split-path integrators.

In this paper, we aim to find conditions for uniform incremental exponential stability (δUES), δISS [1], and the stronger property *quadratic* δISS (q-δISS) [2] of (1).

Theorem 1 *Consider the PWL system* (1)*. Suppose that there exist a matrix* $P = P^\top \succ 0$ *and a number* $\alpha > 0$ *that satisfy the conditions*

$$\begin{bmatrix} A_1^\top P + PA_1 + \alpha P & Pb + c \\ \star & -2 \end{bmatrix} \preceq 0, \tag{2a}$$

$$Pb = v. \tag{2b}$$

Then, the piecewise linear system in (1) *is q-δISS.*

The matrix conditions in (2) are reminiscent of classical conditions for passivity of linear systems. However, the inequality in (2a) coupled with the equality in (2b) requires a somewhat stronger system property than classical passivity. The next theorem reveals this property in terms of a frequency-domain inequality.

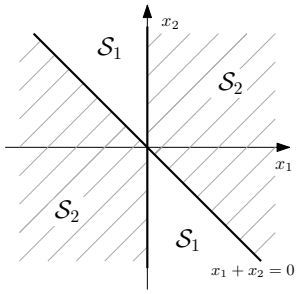

Figure 1: Illustration of the sets in (1) with $x = [x_1\, x_2]^\top \in \mathbb{R}^2$, $c = [-1\, 0]^\top$, and $v = [1\, 1]^\top$.

Theorem 2 *Suppose the pair* (A_1, b) *is controllable. There exist* $P = P^\top \succ 0$ *and* $\alpha > 0$ *that satisfy the matrix conditions in* (2) *if and only if*

$$4\mathrm{Re}\left\{v^\top (j\omega I - A_1)b\right\} \geq \|(c^\top + v^\top)(j\omega I - A_1)b\|^2 \tag{3}$$

is satisfied for all $\omega \in \mathbb{R}$*.*

The frequency-domain condition in (3) requires something more than passivity of the system $G_v(s) = v(sI - A_1)b$ alone, i.e., it is not enough for $G_v(s)$ to be positive real.

3 Future work

Interesting directions for future work include the extension of the results to more general classes of PWL systems and exploring the possibility to relax the conditions through the use of piecewise quadratic functions.

References

[1] D. Angeli, "A Lyapunov approach to incremental stability properties," *IEEE Trans. on Aut. Cont.*, 47(3), 2002.

[2] A. Pavlov, A. Pogromsky, N. van de Wouw, H. Nijmeijer. "On convergence properties of piecewise affine systems," *International Journal of Control*, 80. pp. 7-12, 2005.

The impact of sensor placement on graph-neural-network-based leakage detection

J.J.H. van Gemert
Control Systems group
University of Technology
The Netherlands
j.j.h.v.gemert@tue.nl

V. Breschi, M. Lazar
Control Systems group
University of Technology
The Netherlands
{v.breschi, m.lazar}@tue.nl

D.R. Yntema
Wetsus, Centre of excellence for sustainable water technology
The Netherlands
doekle.yntema@wetsus.nl

K.J. Keesman
Systems and Control group
Wageningen University
The Netherlands
karel.keesman@wetsus.nl

1 Introduction

Leakages in water distribution networks (WDNs) cause significant water losses, operational costs, and potential public-health risks. As water scarcity increases, efficient leakage detection becomes increasingly more urgent. Recent advances in graph neural networks (GNNs) enable pressure reconstruction and leakage detection from sparse measurements by exploiting the WDN topology. However, their performance strongly depends on sensor placement, which is often fixed or chosen arbitrarily and rarely analyzed explicitly. To address this gap, we introduce a scalable, topology-based, and model-free sensor placement strategy based on PageRank centrality, inspired by [2]. Using the EPANET Net1 benchmark, the resulting sensor configuration is evaluated within a GNN-based reconstructor–predictor framework for pressure reconstruction and one-step-ahead prediction, which highlights the impact of PageRank-based sensor placement.

2 PageRank-Centrality-based Sensor Placement

We propose a sensor placement approach based on the WDN topology by viewing it as a graph and exploiting PageRank centrality. For an adjacency matrix $A \in \mathbb{R}^{n\times n}$ and degree matrix D, the PageRank vector p is obtained from

$$p^{(k+1)} = \alpha M p^{(k)} + (1-\alpha)v, \quad M = AD^{-1}, \quad v = \tfrac{1}{n}\mathbf{1}^n, \quad (1)$$

which evaluates how accessible each junction is under a random walk. Sensors are placed at junctions with minimum PageRank centrality, as their pressures are hardest to reconstruct from neighboring measurements.

3 GNN for Pressure Reconstruction and Prediction

The sparse pressure measurements obtained from the PageRank-based placement are used as input to a GNN for pressure reconstruction and one-step-ahead prediction. A Chebyshev GNN (ChebNet) introduced in [1] is employed, so that information is propagated over the network topology using truncated Chebyshev polynomial filters, i.e.,

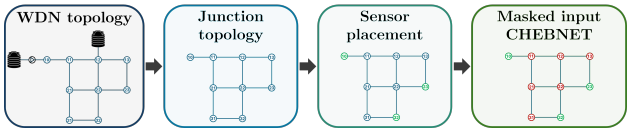

Figure 1: From WDN topology to Chebnet input.

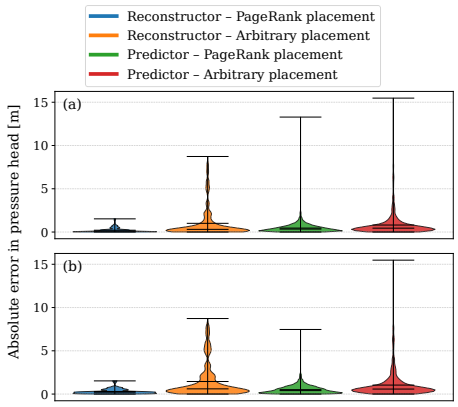

Figure 2: Residual distributions for the four GNNs on Net1, shown as violin plots over all junctions in (a) and over non-sensed junctions in (b).

$$g_\Theta(\hat{L})x = \sum_{k=0}^{K} \Theta_k T_k(\hat{L})x. \quad (2)$$

To assess the impact of sensor placement on the reconstruction and prediction of pressure, the PageRank-Centrality-based configuration is compared with an arbitrary placement on the EPANET Net1 network. For each sensor configuration, a ChebNet pressure reconstructor and predictor are trained, resulting in four models in total. Figure 2 shows the residual distributions for pressure reconstruction and prediction over all junctions and over non-sensed junctions. The PageRank-based placement yields narrower and more concentrated residual distributions, indicating improved pressure reconstruction and prediction accuracy. These improvements lead to more reliable residual-based leakage detection, with reduced sensitivity to demand-driven fluctuations.

References

[1] Gergely Hajgató, Bálint Gyires-Tóth, and György Paál. Reconstructing nodal pressures in water distribution systems with graph neural networks. *arXiv preprint arXiv:2104.13619*, 2021.

[2] J.J.H. van Gemert, V. Breschi, D. R. Yntema, K. J. Keesman, and M. Lazar. Exploiting structural observability and graph colorability for optimal sensor placement in water distribution networks. In *2025 European Control Conference (ECC)*, pages 546–551, 2025.

Persistent versus fading memory in serial recall: an analysis through recurrent neural networks

Lucas Stordeur[1], Guillaume Drion[1], Steve Majerus[2], Pierre Sacré[1]
[1]Montefiore Institute, University of Liège, Belgium [2] PsyNCog, University of Liège, Belgium
{Lucas.Stordeur, gdrion, smajerus, p.sacre}@uliege.be

Memory is a fundamental cognitive process involved in many everyday activities, such as learning new information, retaining past experiences, and solving problems. However, memory is not a unitary function: it encompasses a variety of processes and mechanisms that may be differentially engaged depending on the task and temporal demands. Among these mechanisms, two broad forms can be distinguished: persistent memory, which supports the long-term maintenance of information, and fading memory, in which information is maintained only transiently.

In psychology, serial recall is an experimental task requiring the recall of items in their original order, which is extensively studied to investigate the mechanisms of working memory. For this reason, a wide range of models has been developed to account for experimentally observed results. While most models rely on separate representations of item and order, Botvinick and Plaut (2006), motivated by new electrophysiological data from monkeys, proposed a dynamical model based on recurrent neural networks that employs a unified representation of item and order with a very simple cell dynamic [1].

Over recent years, recurrent neural networks have become increasingly prominent and effective through the introduction of new mechanisms and cell dynamics that enable more complex network behavior. In particular, fading memory can now be implemented using gated recurrent unit (GRU) cells [2], while persistent memory can be achieved with bistable recurrent cells (BRCs) [3], which rely on bistability.

In this work, we investigate which type of memory is better suited to represent this task of serial recall. While both fading and persistent memory can reproduce the main experimental findings, certain experiments show that bistable recurrent cells better replicate the observed experimental results and vise-versa. By combining both types of cells within the same network and modulating their relative proportions, we aim to better understand the respective roles of these memory mechanisms and their impact on experimental replication (Figure 1). In addition, we perform internal network analyses to examine how bistability is exploited during task execution. Finally, we apply lesions to the network, revealing that lesions affecting persistent memory lead to different outcomes than lesions affecting fading memory.

References

[1] M. M. Botvinick and D. C. Plaut. *Short-Term Memory for Serial Order: A Recurrent Neural Network Model*. Psychological Review, 113(2):358–389, 2006.

[2] K. Cho, et al. *Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation*. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1724–1734, 2014.

[3] N. Vecoven, D. Ernst, and G. Drion. *A bio-inspired bistable recurrent cell allows for long-lasting memory*. PLoS One, 16(6):e0252676, 2021.

Acknowledgment This work is supported by the Belgian Government through the Federal Public Service Policy and Support. LS is a Research Fellow of the Fonds de la Recherche Scientifique – FNRS.

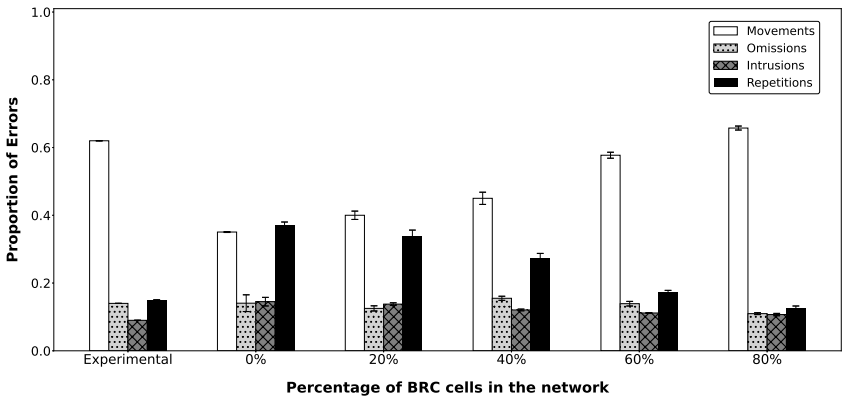

Figure 1: Proportion of errors for networks with varying BRC cell ratios. Errors include movements (item in the wrong position), omissions (item not recalled), intrusions (item not in the sequence), and repetitions (same item recalled multiple times).

Emulation-based Spiking Control for LTI Systems

Elena Petri[1], Koen J. A. Scheres[2], Erik Steur[1], W. P. M. H. (Maurice) Heemels[1]

[1] Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands.

[2] Department of Electrical Engineering (ESAT), STADIUS, KU Leuven, Leuven 3001, Belgium.

Email: e.petri@tue.nl

1 Introduction

Brain-inspired spiking technologies can offer important advantages over classical digital clock-based technologies, including energy efficiency, low-latency, adaptability and robustness in various domains. In systems and control engineering, the interest in neuromorphic control is growing in the recent years, where the aim is to design controllers inspired by biological neurons in the brain, see, e.g., [1–3]. A key fundamental gap in the literature is the lack of a formal framework for the modeling, design and analysis of spiking-based controllers. Focusing on these questions, we present a new mathematically rigorous approach for stabilizing linear time-invariant (LTI) systems using neuromorphic controllers, designed as a network of integrate-and-fire neurons. We consider the scenario where the input to each neuron is (part of) the measured output from the plant, and the neurons generate spiking control signals, as shown in the control configuration in Figure 1 for a single-input single-output setup.

2 Emulation-based approach and contributions

To prove a stability property for the closed-loop system, we propose an emulation-based design approach. We consider output feedback stabilizable LTI systems in closed loop with spiking controllers. Our goal is that the closed-loop trajectory emulates, in a sense that will be clarified below, the continuous-time closed-loop trajectory we would have obtained if the LTI plant would have been stabilized using a static output feedback controller. Our emulation-based design approach consists of two steps. In the first step, we establish conditions on the neuron's parameters (spike amplitude and firing threshold) to guarantee that the spiking signal, modeled as a sequence of Dirac delta pulses, generated by a pair of integrate-and-fire neurons, approximates any continuous-time signal input to the neurons with arbitrary accuracy. In particular, we prove that the error between the continuous-time signal input and the spiky signal output of the pair of neurons can be made arbitrarily small in an integral sense by tuning the neuron's parameters. The second step of the proposed emulation-based design approach consists in the definition of a novel stability notion, called integral spiking-input-to-state-stability (iSISS), building on the specific integral metric from the first step. This new stability property differs from the existing notion of integral input-to-state stability (iISS) [4], as it allows to consider spiking input signals, and, in some cases, iSISS is stronger than iISS. We also prove that an asymptotically stable LTI system is iSISS. This implies that the error system describing the difference between the two closed-loop trajectories is iSISS with respect to the spiky emulation error, i.e., the difference between the spiking signal generated by the neurons and the continuous-time static output feedback control signal to be emulated. By combining these two steps, we establish a practical stability property of output feedback stabilizable LTI systems in closed loop with neuromorphic controllers. Simulations illustrate the results.

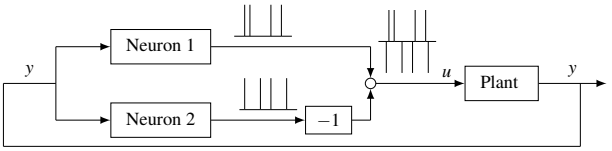

Figure 1: Block diagram representing the system architecture for a single-input single-output setup.

This abstract is based on [5].

References

[1] R. Sepulchre, "Spiking control systems," *Proceedings of the IEEE*, vol. 110, no. 5, pp. 577–589, 2022.

[2] L. Ribar and R. Sepulchre, "Neuromorphic control: Designing multiscale mixed-feedback systems," *IEEE Control Systems Magazine*, vol. 41, no. 6, pp. 34–63, 2021.

[3] E. Petri, K. J. A. Scheres, E. Steur, and W. P. M. H. Heemels, "Analysis of a simple neuromorphic controller for linear systems: A hybrid systems perspective," *IEEE Conference on Decision and Control,* Milan, Italy, pp. 8578–8583, 2024.

[4] E. D. Sontag, "Comments on integral variants of ISS," *Systems & Control Letters*, vol. 34, no. 1-2, pp. 93–100, 1998.

[5] E. Petri, K. J. A. Scheres, E. Steur, and W. P. M. H. Heemels, "Emulation-based neuromorphic control for the stabilization of LTI systems," *arXiv:2511.11875*, 2025.

Simulation of Excitable Neuromorphic Systems Across Event Scales

Amir Shahhosseini, Rodolphe Sepulchre
Center for Dynamical Systems, Signal Processing, and Data Analytics (STADIUS)
Department of Electrical Engineering (ESAT), KU Leuven
amir.shahhosseini@kuleuven.be, *rodolphe.sepulchre@kuleuven.be*

1 Motivation

An efficient and scalable simulation framework is essential for transforming the neuromorphic paradigm into an engineering technology. Acknowledging and exploiting the geometry and structural invariances of the neuromorphic systems is central to attaining such a framework. This geometry stems from the inherent modularity and the presence of well-defined motifs in both the temporal and spatial domains and is present at every scale.

A scalable framework must therefore exploit these structures while modeling only what is relevant at the event scale of interest, while permitting that scale to shift. When a person walks, the observed event is the coordinated swing of the leg. Analyzing the swing requires modeling the central pattern generators driving muscle contractions, not the individual neurons that spike within them. Nevertheless, one can descend to finer scales if deemed necessary (e.g., neuromodulation or design). Crucially, unlike mean-field models, these higher-level events are not detached from the behavior of their building blocks and emerge directly from the dynamics of the underlying elements, which remain the source of the observed behavior and can be directly altered by sliding the scale of the solver.

2 Spatio-Temporal Structures Across Scales

The temporal structures of an excitable neuromorphic system are rooted in the underlying mechanism. The mixed feedback principle [1], consisting of elements of positive feedback at faster timescales and elements of negative feedback at slower timescales, generates excitable events. Precisely speaking, an event scale requires elements of faster positive feedback and slower negative feedback that dictate the timescale of the event. These two timescales define the temporal scale of the event. The spatial scale of the event can be defined with a consistent logic.

3 An Operator-Theoretic Approach

The sampling resolution of the neuromorphic system can act as an averaging mechanism, effectively suppressing the filter banks with timescales that evolve faster than the sampling resolution. Thus, by sliding the resolution, it is possible to transition between the event scales without any need for ad-hoc abstractions of the events or extensive computational expenses.

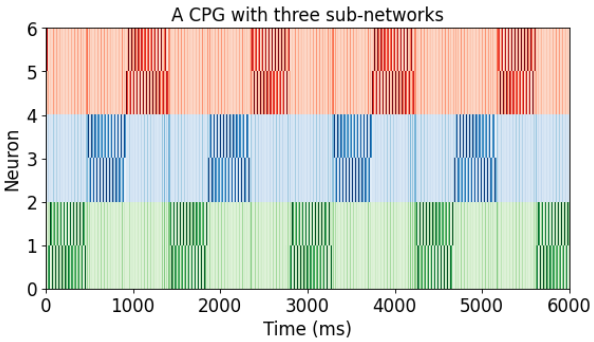

Figure 1: Sub-network rhythms arise from organized spikes [2], but at coarse scales, the spikes themselves are unimportant and only the higher-level event matters. Think of it as recognizing an object through pixels: each pixel contributes to the photo, but it is the assembly of them that provides the meaningful context.

Appreciating the structural invariances of neuromorphic systems becomes significantly easier when the representation of the dynamics shifts from ODEs to operators. This shift has several justifications. First, the size of the ODE models grows superlinearly with the size of the neuromorphic network, creating a scalability bottleneck. Furthermore, these ODEs are stiff as they represent multi-timescale events such as spiking or bursting. Finally, it is impractical to impose prior knowledge of the behavior (e.g., expert knowledge) into the solver. An operator-theoretic representation addresses all these bottlenecks [3].

References

[1] R. Sepulchre, Spiking control systems. *Proceedings Of The IEEE*. **110**, 577-589 (2022)

[2] Y. Huo, F. Forni, & R. Sepulchre, A Neuromorphic Architecture for Scalable Event-Based Control. *ArXiv Preprint ArXiv:2511.11924*. (2025)

[3] A. Shahhosseini, T. Chaffey, & R. Sepulchre, Operator-Splitting Methods for Neuromorphic Circuit Simulation. *ArXiv Preprint ArXiv:2505.22363*. (2025)

Acknowledgment This work was supported by European Research Council under the Advanced ERC Grant Agreement SpikyControl 101054323.

Spiking Control for Stabilization and Tracking in LTI Systems

W.M. Klip, E. Petri, W.P.M.H. Heemels
Control Systems Technology, Department of Mechanical Engineering
Eindhoven university of Technology, 5600 MB Eindhoven, The Netherlands
Email:{w.m.klip, e.petri, w.p.m.h.heemels}@tue.nl

Introduction

Neuromorphic control is an emerging field that aims to mimic the workings of biological neurons to harness their advantages such as low energy consumption, quick event-based responses, adaptability and robustness. The key difference with conventional controllers is the event-based nature where neurons fire as a response to an event and communicate through very short pulses, also called spikes. Control design methods already exist for either event-based or impulsive control but there is a lack of design methods for controllers that are both event-based and impulsive. In this work, this topic is investigated by constraining the control input to a train of spikes which are modelled as Dirac pulses with a fixed amplitude. Note that the only remaining control freedom is the timing of when to fire which spike. We have developed methods for practical stabilization and tracking for linear time-invariant (LTI) systems controlled by such spikes. This abstract is based on submissions [1] and [2].

Problem Formulation

Consider a stabilizable LTI system

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx, \end{aligned} \tag{1}$$

with A invertible, where $x \in \mathbb{R}^{n_x}, u \in \mathbb{R}^{n_u}, y \in \mathbb{R}^{n_y}, n_x, n_u, n_y \in \mathbb{Z}_{>0}$. Given a constant reference y^*, define u^* as the input required needed to remain at y^*, $y^* = -CA^{-1}Bu^*$. The idea is to control this system with bidirectional Dirac spikes such that each input channel is given by, for any $t \geq 0$,

$$u_k(t) = \sum_i \alpha_k \delta(t - t_{k,i}^+) - \sum_j \alpha_k \delta(t - t_{k,j}^-), \tag{2}$$

where α_k is the amplitude of the spike, $t_{k,i}^+$ and $t_{k,j}^-$ are respectively the positive and negative spiking times and $k \in \{1, 2, \ldots, n_u\}$. Note that due to the properties of the Dirac function, this control signal effectively leads to instantaneous jumps in the state which allows us to formulate it as a hybrid system. The system is then in open-loop in between spikes and instantly changes state at spikes. The goal is to formulate a spiking strategy to determine these spiking times such that the state (practically) asymptotically converges to a predefined reference.

Lyapunov-based spiking strategy

The idea for the spiking strategy is based on a recent paper [3], but there, a formal analysis was lacking. This work generalizes and extends this with a formal proof. The idea is to first choose an appropriate quadratic Lyapunov function $V(e) := e^\top Pe$. The spiking condition is then chosen as

$$V(e \pm \alpha_k B_k) + c_k \leq V(e), \tag{3}$$

such that the Lyapunov function always decreases by at least $c_k > 0$ at jumps. B_k refers to column k of B. Between spikes, the Lyapunov function may increase, but it is shown that this increase is relatively small as a result of the spiking condition. It follows that there exists a level set of the Lyapunov function ($V \leq d$) for which $\dot{V} \leq 0$. Since the Lyapunov function also decreases at spikes, this set is an attractor around the specified reference. Therefore, global practical tracking is achieved.

Extended methods

A potential problem with this method is that many spikes can be triggered in the same time instant if $|B_k^\top Pe|$ is large. To this end, a time-regularized method is also developed where the spiking times are constrained to have a minimum fixed inter-spike time. Semi-global practical tracking is proven in this case.

Another method to combat this problem is to extend the method to fit the framework of a spiking neural network (SNN). This is worked out for the stabilization case ($y^* = u^* = 0$). This method is proven to globally practically stabilize the system while always having a positive inter-spike time in the single-input case.

Future work

Future work will focus on extending the SNN implementation to also do tracking, and potentially completely removing the problem of having multiple spikes in one time-instant. Also, the limiting behaviour for small spike amplitudes and the shape of the ultimate trajectory and possible existence of a limit cycle will be investigated. A potential extension to (a class of) non-linear systems is also of interest.

References

[1] W.P.M.H. Heemels, W.M. Klip, A. Aksoy and E. Petri, "Spiking Control for Stabilization and Tracking in Linear Systems: A Greedy Lyapunov-based Approach", under review.

[2] W.M. Klip, E. Petri and W.P.M.H. Heemels, "Stabilizing Linear Time-Invariant Systems with Recurrent Spiking Neural Networks", under review.

[3] P. Agliati et al., "Spiking neurons as predictive controllers of linear systems", arXiv preprint arXiv:2507.16495, 2025.

ISS for PWA systems

Radu Cimpean, Stephan Trenn, Alef Sterk
Bernoulli Institute, University of Groningen, 9747AG Groningen, The Netherlands
r.cimpean@rug.nl s.trenn@rug.nl a.e.sterk@rug.nl

1 Introduction

Starting from the biological-inspired Wanner-Friedrich (WF) ODE model proposed in [1], we are interested in its stability properties. In our previous work [2], we showed stability of a planar version of the model, utilizing methods developed for piecewise affine (PWA) systems. We observed difficulties of extending the method directly to larger dimensions and, therefore, we propose a divide-and-conquer approach, where the larger network is viewed as the interconnection of simple (planar) PWA systems. This approach requires to add inputs to the planar PWA (the coupling terms), and instead of stability of an autonomous system, we are now interested in input-to-state stability (ISS) of PWA systems.

2 Problem statement

We consider PWA systems of the form

$$\dot{x} = A_s x + b_s + B_s u, \quad x \in X_s, \tag{1}$$

where $\bigcup_s X_s = \mathbb{R}^n$ is a suitable state space decomposition, for details see [2]. In line with the properties of the motivating WF system, we make the following assumptions

(A1) Continuity: $A_i x + b_i + B_i u = A_j x + b_j + B_j u$ for all $x \in X_i \cap X_j$, $i, j \in \{0, 1, \ldots, N\}$, and all $u \in \mathbb{R}^m$.

(A2) Average dwell time (ADT): For every input bound $\|u\|_\infty \leq U_{\max}$ the corresponding solutions of (1) induce switching signals with an average dwell time $\tau_a > 0$ and chatter bound N_0.

(A3) Existence of regional quadratic ISS Lyapunov functions: For each region X_s there exists regional ISS Lyapunov functions of the form $V_s(x) = x^\top P_s x_s + 2v_s^\top x + \omega_s$ with $P_s = P_s^\top \in \mathbb{R}^{n\times n}$, $v_s \in \mathbb{R}^n$, $\omega_s \in \mathbb{R}$. Furthermore, there exists $0 < \varepsilon_{1,s} < \varepsilon_{2,s}$ and $\lambda_s > 0$, $\Gamma_s > 0$, $\Theta_s \geq 0$ such that

$$\varepsilon_{1,s}|x|^2 \leq V_s(x) \leq \varepsilon_{2,s}(|x|^2 + 1) \qquad \forall x \in X_s,$$

and, for all $u \in \mathbb{R}^m$,

$$\dot{V}_s(x,u) \leq -\lambda_s V_s(x) + \Gamma_s |u|^2 + \Theta_s \quad \forall x \in X_s.$$

(A4) Crossings: For every solution of the PWA (1) which leaves region X_i and enters region X_j at some time $t_{ij} \geq t_0$ the values of the corresponding Lyapunov functions *do not increase*, i.e.

$$V_j(x(t_{ij})) \leq V_i(x(t_{ij})).$$

(A5) Origin: $0 \in X_0$ and $0 \notin X_s$ for all $s \in \{1, \ldots, N\}$. Furthermore, the regional ISS Lyapunov function V_0 from Assumption (A3) takes the form $V_0(x) = x^\top P_0 x$ and for all $x \in X_0$ and all $u \in \mathbb{R}^m$, satisfies

$$\varepsilon_{1,0}|x|^2 \leq V_0(x) \leq \varepsilon_{2,0}|x|^2,$$
$$\dot{V}_0(x,u) \leq -\lambda_0 V_0(x) + \Gamma_0 |u|^2.$$

2.1 Regional ISS to practical ISS

We begin with $(A3)$, which means that in each region we have an ISS Lyapunov function. We want to arrive at a global ISS type bound, for which we then make use of $(A1), (A2), (A4)$. Assumption (A1) avoids technicalities concerning existence and uniqueness of solutions. Similarly, implying the Assumption (A2) avoids Zeno behavior and chattering. Assumption (A4) is a key ingredient towards achieving practical ISS, given that it allows us to patch together the regional ISS Lyapunov functions to obtain a global one.

2.2 Practical ISS to global ISS

Proving practical ISS is not sufficient to conclude that solutions come to rest once the input vanishes. For that, Assumption (A5) together with a certain quantitative assumption on the overall practical ISS bounds are required. The idea is to make use of practical ISS to ensure that solutions enter X_0 when the input is sufficiently small and then indeed converge to zero if the input vanishes, thus concluding the PWA system exhibits global ISS.

3 Conclusion

To conclude, we have established a computationally feasible method to certify ISS for a certain class of PWA system which are inspired by the WF model of neural interactions. The proposed method first patches together the regional ISS Lyapunov functions to conclude practical ISS, and then, utilizing the ISS property of the region around the equilibrium, establishes the global ISS result under the assumption that the practical ISS properties ensure that the solution enters the region around the equilibrium for sufficiently small inputs.

References

[1] A. Wanner and R.W. Friedrich, "Whitening of odor representations by the wiring diagram of the olfactory bulb," Nat Neurosci. 23(3), 433-442, 2020.

[2] R. Cimpean and S. Trenn, "Utilizing discontinuous piecewise Lyapunov function for a biological PWA system and extending the analysis via input-to-state type stability," 2025 European Control Conference (ECC), Thessaloniki, Greece, 2025, pp. 194-200.

Volume-contracting systems are entropy-reducing

Gustave Bainier
Allée de la Découverte 10, Liège, Belgium
gustave.bainier@uliege.be

Alessio Franci
Allée de la Découverte 10, Liège, Belgium
afranci@uliege.be

At the beginning of cybernetics, *homeostasis* was roughly described as an *entropy-reducing mechanism*: achieved through negative feedback loops, it actively counteracts disorder and uncertainty to maintain a system behavior. Although the concept of entropy played a central role in the early foundations of cybernetics (most notably in the work of Norbert Wiener [1]) it has largely faded from mainstream control theory, except in specialized areas such as stochastic and information-theoretic control (e.g. [2, 3]). In recent decades, control theory has focused on *contracting systems* (along with other incremental concepts) marking a shift away from classical notions of stability [4, 5, 6]. Contraction emphasizes convergence between trajectories rather than convergence to a specific equilibria. In this short note, it is demonstrated that volume-contracting dynamics induce a reduction of entropy in probability distributions. This simple result reconnects modern contraction-based control with the original entropy-based perspective of early cybernetics.

Take a n-dimensional, time-invariant, smooth dynamical system with both forward-complete and backward-complete solutions:

$$\frac{d}{dt}x(t) = f(x(t)). \tag{1}$$

Let φ_t denote the flow of (1). It is a diffeomorphism from R^n to itself. The quantity $|\det(J(\varphi_t)(x))|$, where J stands for the Jacobian, measures the change in volume of a small open neighborhood of x after it has evolved under the flow of (1) for a time t. Volume-contraction for infinitesimal volume elements implies:

$$\lim_{t\to+\infty} |\det(J(\varphi_t)(x))| = 0 \text{ for a.e. } x \in R^n. \tag{2}$$

The Liouville-Ostrogradski formula [5] yields:

$$|\det(J(\varphi_t)(x))| = \exp\Big(\int_0^t \operatorname{div}(f)(\varphi_s(x))ds\Big),$$

so (2) implies:

$$\lim_{t\to+\infty} \int_0^t \operatorname{div}(f)(\varphi_s(x))ds = -\infty \text{ for a.e. } x \in R^n.$$

We make the technical assumption that there exist $t^* > 0$ and $M \in R$ such that for all $t \geq t^*$:

$$\int_0^t \operatorname{div}(f)(\varphi_s(x))ds < M \text{ for a.e. } x \in R^n. \tag{3}$$

Both (2) and this assumption are easily verified for a uniformly negative divergence, $\operatorname{div}(f) \leq -\lambda < 0$.

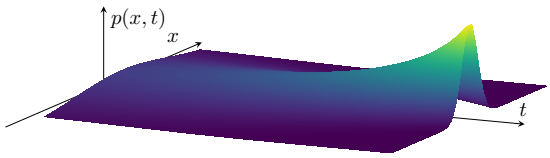

Figure 1: Volume-contracting flows cause probability distributions to compress, thus decreasing their entropy.

Now consider (1) at a random initial condition X_0, with a probability density function p. The entropy of X_0 is:

$$h(X_0) := -\int_{R^n} p(x)\ln(p(x))dx.$$

For all $t \geq 0$, $X_t := \varphi_t(X_0)$. Since φ_t is a diffeomorphism from R^n to itself, a well-known result in information theory [7] expresses the difference of entropy between X_t and X_0 in terms of the Jacobian of φ_t as:

$$h(X_t) - h(X_0) = \int_{R^n} p(x)\ln|\det(J(\varphi_t)(x))|dx.$$

Applying the Liouville-Ostrogradski formula yields:

$$\boxed{h(X_t) - h(X_0) = \int_{R^n} p(x)\Big(\int_0^t \operatorname{div}(f)(\varphi_s(x))ds\Big)dx.}$$

If (1) is such that (2) and assumption (3) hold, the reverse Fatou's lemma demonstrates that the system is entropy-reducing:

$$\lim_{t\to+\infty} h(X_t) - h(X_0) = -\infty.$$

References

[1] N. Wiener, "Cybernetics: or control and communication in the animal and the machine."; 1948, revised ed. 1961.

[2] T. Bourdais, N. Oudjane, F. Russo. "An entropy penalized approach for stochastic control problems."; 2025.

[3] H. Touchette, S.-Lloyd. "Information-theoretic approach to the study of control systems."; Physica A, 2004.

[4] W. Lohmiller, J-J E. Slotine, "On contraction analysis for non-linear systems"; Automatica, 1998.

[5] C. Wu, I. Kanevskiy, M. Margaliot, "k-contraction: theory and applications."; Automatica, 2022.

[6] F. Bullo, "Contraction theory for dynamical systems."; 2024.

[7] S. Ihara. "Information theory for continuous systems."; World Scientific Publishing, 1993.

Semi-physical water balance model for agricultural polder systems

Maulana, F.R.[1,2,*], Wagterveld, R.M.[2,3], Cheng, X.[1], Mayer, M.J.J.[3], van Heijster, P.[1], and Keesman, K.J.[1,2]
[1]Mathematical and Statistical Methods (Biometris), Wageningen University and Research,
PO Box 16, 6700 AA, Wageningen, The Netherlands
[2]Wetsus, European Centre of Excellence for Sustainable Water Technology,
Oostergoweg 9, 8911 MA, Leeuwarden, The Netherlands
[3]Aqa.earth, Tweede Sluisweg 35, 8413 NN Oudehorne, The Netherlands
*Email: fakhira.maulana@wur.nl

1 Introduction

The escalating degradation of surface water quality due to overfertilization of Nitrogen and Phosphorus in agricultural soils has necessitated intensive, real-time monitoring of soil and surface water quality [1]. However, existing mechanistic hydrological models used in the monitoring operations often underrepresent small water bodies, such as agricultural ditches, and exhibit large uncertainties due to complex model structures and difficult-to-measure physical parameters. We propose a simplified dynamic model of soil and surface-water balance for agricultural polder systems. The model adopts a semi-physical approach, combining physics-based soil-water relations with empirical parameters to represent local hydrological processes in flat polder systems. Quantifying soil water fluxes is critical for coupling with water-quality models to predict the dynamics of nutrient pollutants, thereby supporting water-quality monitoring and sustainable fertilization.

2 Data Collection

We used input and validation data from a multisensor water-quality network at a dairy farm in Lemmer, the Netherlands [2]. The sensors recorded hourly soil moisture in the topsoil and subsoil, ditch water level, and weather variables. Given the fluctuating groundwater table in the polder systems, we used groundwater level data from a nearby monitoring well.

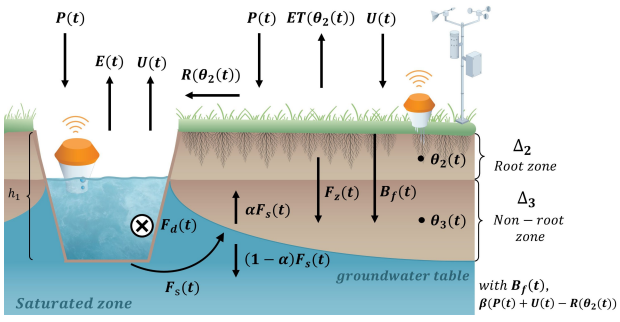

Figure 1: Schematization of the soil-water balance model

3 Model Formulation

The model is divided into three compartments: a ditch, an unsaturated zone with a root zone, and non-root-zone layers. Soil moisture states in the root zone $(\theta_2(t))$ and non-root zone $(\theta_3(t))$ are predicted using soil water balance **(Figure 1)**. Ditch evaporation $(E(t))$ and grass evapotranspiration $(ET(\theta_2(t)))$ are estimated based on the Penman-Monteith method. Percolation between the two zones $(F_z(t))$ is modeled using Darcy-Buckingham law with a non-constant unsaturated hydraulic conductivity $(K(\theta_2(t)))$. The hydraulic head (ψ) is derived from the Mualem-van Genuchten equation $(\theta(\psi))$. Surface irrigation is applied as $U(t)$, and the lateral flow from the adjacent ditch is denoted by $F_d(t)$.

We introduce new empirical models to model the local processes. Seepage $(F_s(t))$ is partitioned into an upward flow and a downward flow using an α parameter. The bypass flow $(B_f(t))$ is calculated as a fraction β of effective rainfall and irrigation. Surface runoff $(R(\theta_2(t)))$ is generated using linear penetration resistance and saturation excess scenario.

4 Methodology

The structural and practical identifiability of the parameters (α, β) are investigated. The analytical expressions for both parameters are derived using steady-state analysis. A recursive joint state–parameter estimation is performed. The model is validated using time-series soil moisture data.

References

[1] Sudriani, Y., Sebestyén, V., & Abonyi, J. (2023). *Surface water monitoring systems—the importance of integrating information sources for sustainable watershed management.* IEEE Access, 11, 36421-36451

[2] Mayer M.J.J., Akkerman F., Oudakker M., and Wagterveld R.M. (2023). *Network and method for the monitoring of soil and / or surface water quality* (Patent No. 2023211277A)

Acknowledgement This work was performed in the cooperation framework of Wetsus, European Centre of Excellence for Sustainable Water Technology (www.wetsus.nl). Wetsus is co-funded by the European Union (Horizon Europe, LIFE, Interreg and EDRF), the Province of Fryslân and the Dutch Government: Ministry of Economic Affairs (TTT, SBO & PPS-I/TKI Water Technology), Ministry of Education, Culture and Science (TTT & SBO) and Ministry of Infrastructure and Water Management (National Growth Fund - UPPWATER). The authors like to thank the participants of the research theme "Monitoring & Quality" for the fruitful discussions and their financial support.

Equivalent circuit model parameter identification for lithium-ion batteries and adaptation to aging

Ali Rahdarian
Depart. SAAS
Université libre de Bruxelles
CP 165/55, 50 Av.F.D. Roosevelt
B-1050 Brussels
Belgium
Email: ali.rahdarian@ulb.be

Michel Kinnaert
Depart. SAAS
Université libre de Bruxelles
CP 165/55, 50 Av.F.D. Roosevelt
B-1050 Brussels
Belgium
Email: michel.kinnaert@ulb.be

1 Introduction

Compared with other commonly used batteries, lithium-ion batteries continue to draw more attention as a promising energy storage technology. In most applications, a battery system consists of the battery and the battery management system (BMS), which is made up of hardware and software. A BMS controls the charging and discharging of the battery while guaranteeing reliable and safe operation. A BMS is also responsible for other tasks, such as cell balancing and thermal management of the battery pack. Designing BMS software for Li-ion batteries relies on a model that can describe the battery dynamics. For this purpose, equivalent circuit models (ECMs) or physics-based models can be used. ECMs are more popular due to their simplicity and lower computational cost compared to physics-based models.

Using the model, the State of charge (SOC) of the battery can be monitored, which is one of the most important tasks of a BMS. SOC represents the remaining energy of the battery and is essential for preventing overcharge and overdischarge. In recent years, research has been focused on model-based SOC estimation approaches based on ECMs or physics-based models. The main challenge for maintaining accurate SOC estimation over the battery's lifetime is aging [1]. As a result of aging, models that have been identified based on the fresh cell data lose their accuracy. ECM parameters, including open circuit voltage (OCV), capacity, and overpotential parameters, change. This decreases the accuracy of model-based SOC estimation methods such as the extended Kalman filter (EKF). The impact of aging on the accuracy of SOC estimation has been studied extensively in the literature. In addition, it has been shown that to maintain an accurate SOC estimate over time, model parameters must be updated upon aging. For instance, in [2], the authors analysed the impact of aging on the accuracy of the SOC estimation using EKF. It was seen that when the battery ages, if parameters are not updated, the estimator's performance deteriorates significantly. However, there are significant challenges in updating ECM parameters over time. In practice, diagnostic tests such as hybrid pulse power characterization (HPPC) and OCV characterization tests can not be performed during normal battery operation, as they are time-consuming. Therefore, parameter updating must be based on data obtained from BMS sensors, which exhibit lower precision compared to testers. In response to these challenges, multiple studies have been focused on updating ECM parameters, including capacity. In [3], a method based on weighted total least squares (WTLS) was developed for estimating capacity from operational data. By taking into account the impact of noise on current and uncertainity in SOC estimate, the method provides an unbiased estimate of the capacity. Owing to its low computational complexity, the method enables real-time capacity estimation during normal battery operation. However, updating capacity alone is insufficient to fully compensate for aging effects, as the OCV characteristics and overpotential parameters also evolve with aging. In most existing studies, the OCV characteristics are assumed to be independent of aging; however, this assumption does not hold in practice. The aim of our work is to develop a straightforward method to identify ECM parameters from fresh cell data and update all the ECM parameters, including the OCV characteristics, using the operational data. For this purpose, the system identification for the fresh cell is formulated as a linear regression problem and solved using the data obtained from the tester. As the battery ages, parameter update is carried out using total least squares to account for the measurement inaccuracies and noise in the BMS sensor data.

References

[1] J.P Rivera-Barrera, N. Muñoz-Galeano, H. O Sarmiento-Maldonado, "SOC Estimation for Lithium-ion Batteries: Review and Future Challenges," Electronics, vol. 6, no. 4, 2017.

[2] X. Hu, S. Li, H. Peng, and F. Sun, "Robustness analysis of State-of-Charge estimation methods for two types of Li-ion batteries," Journal of power sources, vol. 217, pp. 209–219, 2012.

[3] G. L. Plett, "Recursive approximate weighted total least squares estimation of battery cell total capacity," vol. 196, pp. 2319-2331, 2011.

Exploring frequency-tunable actuators for NTM control in future nuclear fusion reactors

David Becker[†], and Matthijs van Berkel[†‡]
[†]DIFFER- Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands[1]
[‡]Eindhoven University of Technology, Eindhoven, the Netherlands
Email: d.l.c.becker@differ.nl and m.vanberkel@differ.nl

1 Nuclear fusion

The field of nuclear fusion aims to produce net energy using nuclear reactions between light nuclei, most commonly the hydrogen isotopes deuterium and tritium. In magnetic confinement fusion, this is achieved by confining a plasma, consisting of the reactants, in a reactor such as a tokamak. Control plays a major role in magnetic confinement fusion. For example, control is crucial to prevent plasma instabilities and disruptions.

2 NTM control

The Neoclassical Tearing Mode (NTM) is a plasma instability that can reduce reactor performance and can result in the loss of the entire plasma [1]. Each NTM corresponds to a 'magnetic island' at the plasma flux surface where the safety factor q is equal to a characteristic value. For future reactors, the NTMs with mode numbers $q = 2/1$ and $q = 3/2$ are expected to impact reactor performance the most [2]. Furthermore, the NTMs are characterized by the full island width w. NTMs are metastable for small $w < w_{crit}$. However, other plasma instabilities can introduce NTMs where $w > w_{crit}$, which will increase in width. The control objective is to detect and control the island width w before it reaches a threshold value for which the magnetic island cannot be removed anymore. For the ITER reactor, this results in a maximum allowed latency of 1050 ms for the $q = 2/1$ mode.
An NTM can be detected using the Electron Cyclotron Emission (ECE) diagnostic [2] within 250 ms. The actuation in NTM control is done using Electron Cyclotron Heating (ECH) [1]. In this scheme, illustrated in Figure 1, externally generated electron cyclotron (EC) waves are injected into the plasma, where the deposition location should be aligned with the center of the magnetic island. This reduces the NTM width. The dynamics of the interaction depend on the EC power P_{ECRH} and the displacement x_{dis} between the island center and EC deposition location. The conventional microwave source used for ECH is the gyrotron, which operates at a fixed frequency [2]. As a consequence, the deposition location can only be tuned by manipulating the orientation of the final mirror in the EC transmission line. The settling time of the mirror can be as long as 2500 ms, which makes control of the $q = 2/1$ mode challenging [3].

3 Frequency-tunable actuator

The NTM control system can be described as two SISO systems. The power control loop has P_{EC} as inputs and controls the island width w. The positioning loop has the mirror angle θ as input and x_{dis} as output. Given the slow dynamics of the mirror actuator and therefore the positioning loop, **this work explores the use of an alternative microwave source, which is adjustable in frequency**. It is possible to control x_{dis} by varying the EC frequency instead of the mirror angle. Using this approach, the deposition location can be tuned non-mechanically, which is much faster (ms timescale). Using MATLAB Simulink, the (approximated) dynamics of the frequency-tunable NTM control system will be compared to those of the old system. The final aim of this work will be to design the actuator and characterize the actual dynamics of the new system.

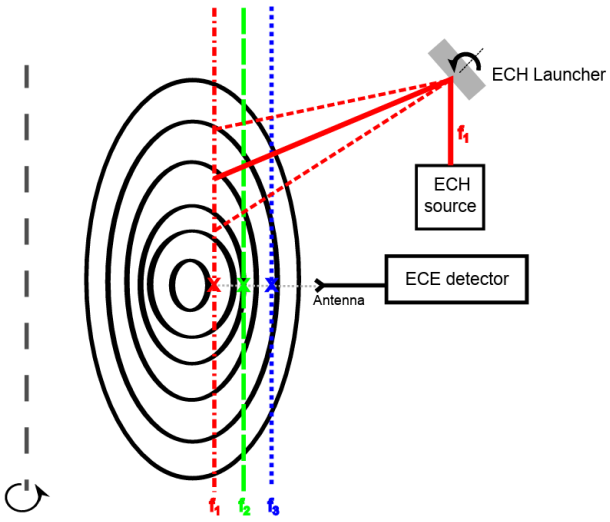

Figure 1: Schematic of the NTM control system [2].

References

[1] La Haye, R. J., Physics of Plasmas **13**, (2006)
[2] van den Brand, H., PhD Dissertation, (2016)
[3] Collazos, A., et al., IEEE Trans. Plasma Sci. **38**, (2010)

[1] DIFFER is part of the institutes organisation of NWO.

Optimal Co-Design of Charging Energy Hubs for Heavy-Duty Electric Transport

Maedeh Izadi, Mauro Salazar, Theo Hofman
Control Systems Technology section
Eindhoven University of Technology
Email: {m.izadi.najafabadi1, m.r.u.salazar, t.hofman }@tue.nl

Abstract

Driven by European and Dutch climate regulations, the logistics sector in the Netherlands is transitioning toward electrification. This shift will substantially increase electricity demand, particularly for heavy-duty vehicles that require large batteries and fast charging. The rapid growth of both renewable energy and electric vehicle charging demand has already led to grid congestion, which could slow or even block the transition to zero-emission heavy transport. Charging Energy Hubs (CEHs) mitigate these impacts by integrating charging infrastructure with renewable energy sources and battery storage. Determining the optimal size of CEH components is therefore a critical investment decision, yet remains challenging due to the strong coupling between design choices and operational dynamics.

This study formulates the CEH design optimization problem as a mixed-integer linear programming (MILP) model to determine the optimal number of charging units, photovoltaic systems, wind turbines, and battery energy storage systems. A co-design framework is adopted in which plant design decisions and the energy management strategy, which distributes power among different energy sources, are optimized simultaneously. This approach is essential because sizing decisions strongly affect operational flexibility and costs, while operational constraints such as grid limits, renewable generation variability, and electric vehicle (EV) logistics directly influence which component sizes are feasible and cost-optimal. The objective function minimizes the total cost of ownership of the CEH while ensuring that the energy demand of the EV fleet is fully satisfied. A modular system configuration is adopted to ensure practical implementability and scalability, in which each component is modeled as a discrete unit with known capacity and technical characteristics. We model all individual EV charging sessions using a smart-charging framework inspired by [1], in which each charging session occurs within the vehicle's arrival and departure window. The proposed model incorporates grid power constraints that impose upper bounds on electricity import and export according to the CEH grid-connection contract. To reduce model dimensionality and computational burden, we adopt a scenario-based approach [2]. Representative scenarios capture seasonal weather variations and temporal differences in EV charging demand. Each scenario represents a typical day and is weighted by the number of times it occurs over the year.

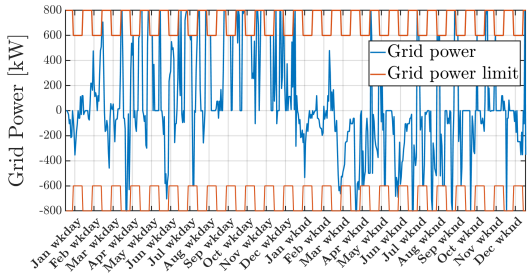

Figure 1: Grid power across all representative scenarios (positive: withdrawals, negative: injections).

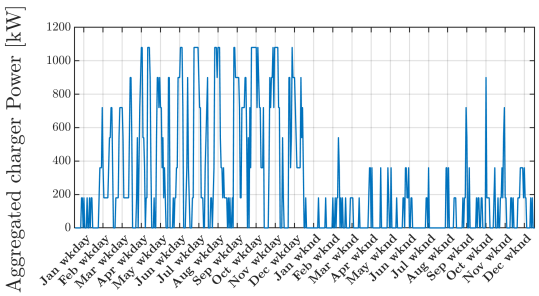

Figure 2: Aggregated EV charging power across all representative scenarios, showing charging-demand variability.

We apply the optimization framework to a heavy-duty fleet over a one-year horizon with 24 representative scenarios (weekday and weekend profiles for each month). The resulting CEH configuration includes 1110 solar units with a nominal power of 550 W each, 1 wind turbine with a rated power of 500 kW, 6 battery units with a capacity of 580 kWh each, 4 chargers rated at 180 kW, and 1 charger rated at 360 kW. Fig. 1 and Fig. 2 illustrate the optimal operational profiles of the CEH configuration.

References

[1] A. Khaksari, G. Tsaousoglou, P. Makris, K. Steriotis, N. Efthymiopoulos, and E. Varvarigos, "Sizing of electric vehicle charging stations with smart charging capabilities and quality of service requirements," *Sustainable Cities and Society*, vol. 70, p. 102872, 2021.

[2] M. Izadi, D. Fernandez Zapico, M. Salazar, and T. Hofman. (2026) Optimal sizing of charging energy hubs for heavy-duty electric transport through co-optimization. Available at https://arxiv.org/abs/2602.01502.

Model Predictive Control of Radiator Heating Systems with Thermostatic Valves

Erfan Shakhesi, W. P. M. H. (Maurice) Heemels, Paula Chanfreut
Control Systems Technology Section, Department of Mechanical Engineering,
Eindhoven University of Technology, The Netherlands.
{e.shakhesi, m.heemels, p.chanfreut.palacio}@tue.nl

1 Introduction

Buildings are increasingly required to provide thermal comfort and energy flexibility while simultaneously reducing their environmental footprint. In many buildings, space heating is provided by radiators equipped with thermostatic radiator valves (TRVs), which are supplied with hot water from one or more central boilers (see Figure 1). Within this architecture, each TRV acts as a proportional controller that regulates the local heat flow based on the room temperature relative to its setpoint [1]. Although this provides basic comfort control, the lack of coordination among the heating systems and the absence of predictive capability (e.g., for outdoor temperature) lead to reduced energy efficiency. To address these limitations, model predictive control (MPC) for buildings has emerged as a promising approach [2], [3].

2 Thermal Modeling

The MPC design relies on a thermal model of buildings. This model can be described using room temperatures, wall temperatures, and lumped radiator temperatures as state variables. The dynamics of the room temperature $T_{\text{room},i}$ and wall temperature $T_{\text{wall},j}$ can be written as

$$\begin{aligned} C_{\text{room},i}\dot{T}_{\text{room},i} &= \text{room-wall heat exchange} \\ &+ \text{room-outdoor heat exchange} + \text{radiator heat output} \\ &+ \text{internal heat gains} + \text{solar radiation}, \end{aligned}$$

and

$$\begin{aligned} C_{\text{wall},j}\dot{T}_{\text{wall},j} &= \text{wall-room heat exchange} \\ &+ \text{wall-outdoor heat exchange} + \text{solar radiation}, \end{aligned}$$

where $C_{\text{room},i}$ and $C_{\text{wall},j}$ denote the thermal capacitance of room i and wall j, respectively [4]. Here, the radiator heat output is the only adjustable term, and it is governed by the corresponding radiator temperature $T_{\text{rad},k}$ evolving as

$$C_{\text{rad},k}\dot{T}_{\text{rad},k} = \text{heat supplied by boiler} - \text{radiator heat output},$$

where $C_{\text{rad},k}$ is the thermal capacitance of radiator k. Thus, the boiler supply temperature provides control over the radiator heat output and is therefore selected as the control input in the proposed MPC framework. Note that the heat delivered to each room also depends on the water flow rate through the corresponding radiators, which is controlled by the TRV and must be considered in the model.

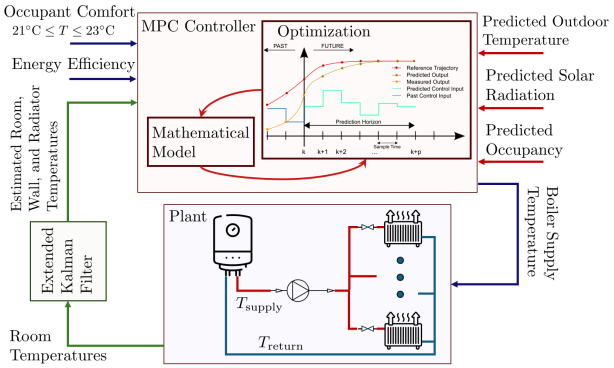

Figure 1: Overall topology of the control architecture.

3 Proposed Model Predictive Control

We propose an MPC-based controller that regulates the boiler supply temperature to keep each target room within a prescribed comfort range while minimizing energy costs. The controller explicitly incorporates forecasts of outdoor temperature, solar radiation, and internal heat gains. Moreover, in typical building installations, only the room temperatures are measured. Therefore, an observer, such as an extended Kalman filter, is required to estimate the unmeasured states, particularly the wall and radiator temperatures [3]. The overall topology is depicted in Figure 1.

Future work includes exploring different MPC topologies, accounting for pipe delays and losses, and validating the controller on a real building.

References

[1] B. Xu *et al.*, "Dynamic simulation of space heating systems with radiators controlled by TRVs in buildings," Energy and Buildings, 2008.

[2] J. Drgoňa *et al.*, "All you need to know about model predictive control for buildings," Annual Reviews in Control, 2020.

[3] A. J. Khabbazi *et al.*, "Lessons learned from field demonstrations of model predictive control and reinforcement learning for residential and commercial HVAC: A review," Applied Energy, 2025.

[4] D. Sturzenegger *et al.*, "BRCM Matlab Toolbox: Model generation for model predictive building control," American Control Conference, 2014.

Improved load estimation in hybrid testing of floating wind turbines

B. Gerlagh, W. Hakvoort
Control and Mechatronics, University of Twente, P.O Box 217, 7500 AE Enshede, The Netherlands
Email: {b.gerlagh, w.b.j.hakvoort}@utwente.nl

1 Introduction

Floating Offshore Wind Turbines (FOWTs) require scaled testing to cut costs and validate simulations. While wind tunnels are ideal for aerodynamic load studies, scaling both aerodynamic and hydrodynamic loads simultaneously is fundamentally infeasible. Hybrid testing overcomes this by using a physical scale model turbine and a hexapod to emulate hydrodynamic loads [2]. However, accurate aerodynamic load measurement, which critical for proper hexapod control, is hindered by load distribution and sensor limitations (e.g., post-sensor inertia). This study aims to improve aerodynamic load estimation in hybrid FOWT testing by leveraging disturbance and noise characteristics for greater accuracy and reliability.

2 Methods and Results

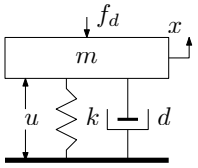

Figure 1: Mass (m) spring (k) damper (d) system where disturbance, f_d and mass are unknown. u characterises the control force and x the output position.

In a simplified problem formulation, we assume the dynamics of a second order mass spring damper system subjected to measurement and load disturbances (Figure 1). The goal is to estimate both the mass and disturbance signals using only position information. In the context of hybrid testing of floating wind turbines, the wind loads on the turbine follow a Kaimal spectrum, S_{ww} ,which is characterised by its dominant low-frequency content [1]. The thrust that acts on the nacelle scales linearly with the Kaimal spectrum and is considered to be the disturbance f_d in Figure 1. In contrast, wave loads that are emulated by control input u, adhere to either a JONSWAP or ISSC spectrum, S_{uu}, both of which exhibit higher-frequency behaviour and contain most of their energy beyond the first rigid-body buoyancy mode of the floating structure. We model a linear system $V_i(s)$ on the given spectra, S_{ww} and S_{uu} by matching the energy of the spectrum as

$$V_i(s) = \arg\min_{V_i} \int_{-\infty}^{\infty} \left| |H_x(j\omega)|^2 - S_{xx}(\omega) \right|^2 d\omega \,, \quad (1)$$

where $i \in \{w, u\}$. For S_{ww}, we fit a lowpass filter and for S_{uu} we fit a second order bandpass filter (Figure 2). The modelled disturbance spectra can be used to extend the original system such that the input disturbance is zero mean Gaussian. The states of the extended state system are estimated using an Extended Kalman filter. Results show that the mass

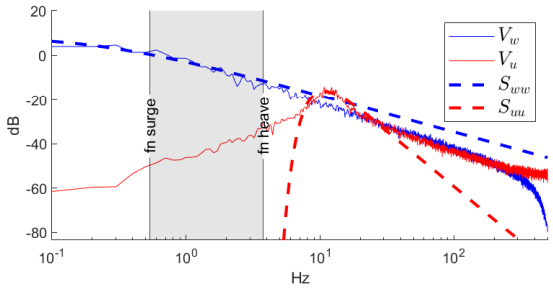

Figure 2: Modelled spectra of S_{uu} and S_{ww}. The range of first natural frequencies in floating structures is indicated in grey.

can be estimated correctly (Figure 2). However, the estimate of the aerodynamic forces still contains substantial error in the order of 20% of the actual signal. The dominant spectral content of the input signal u is above the first mode, while the dominant spectral contents of the disturbance w is below the first mode. Future work is in improving the observer and including a force sensor. Force sensors that use strain gauges often exhibit drifting behaviour and a noise floor characterised by S_{vv}.

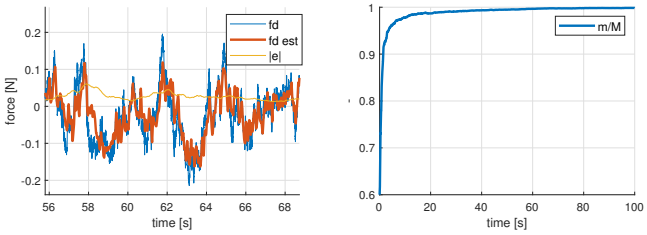

Figure 3: Extended Kalman filter results, disturbance forces (left) and estimated mass over the actual mass (right).

References

[1] L. Arany. et al., 2015 J. Wind Energ., 18: 2171–2197
[2] I. Bayati et al., 2018 J. Phys.: Conf. Ser. 1037 052025

Introducing Crospi: an open-source Constraint-based Reactive and Orchestrated Sensor-driven PIpeline for robot control

Santiago Iregui, Federico Ulloa Rios, Herman Bruyninckx, Erwin Aertbeliën
Robotics Research Team
Department of Mechanical Engineering, KU Leuven
Flanders Make@KU Leuven, Leuven, Belgium
Email: santiago.iregui@kuleuven.be

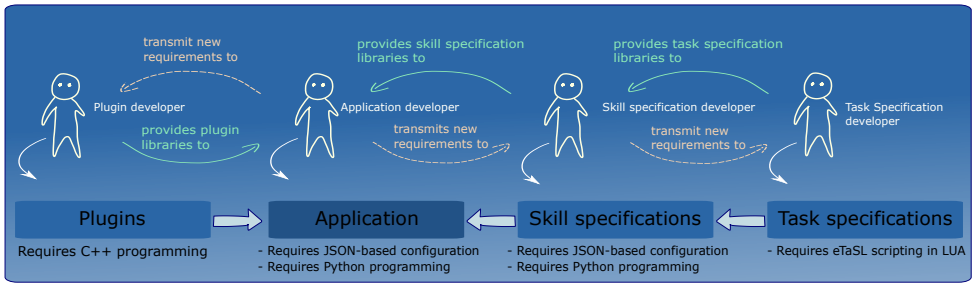

Figure 1: Architecture that separates the concerns of robot application development into four blocks, each with its corresponding stakeholder.

1 Introduction

Crospi[1] is an open-source software that enables reactive robotic behaviors, while reducing the expertise required for application development (see Fig. 1). Multiple sensors (e.g. force/torque, vision, IMU, tactile, etc.) can be easily integrated (via ROS2 or other communication channels) into robot control loops to generate real-time sensor-driven behaviors.

2 Methodology

Crospi relies on a constraint-based approach to specify the desired robot behavior at a high level of abstraction based on task specifications written in eTaSL [1]. Crospi reads these task specifications in configuration states of the orchestrator [2] to automatically generate real-time controllers that execute the desired behavior in multi-robot systems. Crospi is designed to be modular and extensible, allowing to easily integrate new sensors, robots, solvers, etc. It is implemented in C++ with real-time performance and ROS2 integration.

Fig. 1 illustrates Crospi hierarchical architecture, which reduces the expertise required by the different stakeholders. The reusability of the different blocks is facilitated, e.g. by automatically-generated JSON Schemas that define configurable parameters of each block.

[1]Code and tutorials available at: https://rob.pages.gitlab.kuleuven.be/crospi/.

[2]We provide integration for Behavior trees and FSMs using BetFSM https://github.com/Robotics-Research-Group-KUL/betfsm, but other alternatives such as YASMIN can be also used.

3 Application example

Crospi has been used in several use cases, including a mobile manipulation use case for cable harnessing, combining Neura's MAV mobile base with the MAiRA robotic arm (see Fig. 2). The task relies on force-based skills that regulate cable tension during manipulation, requiring low-level closed-loop control at 1 kHz. By coordinating MAV's mobility with MAiRA's fast dynamics, the system robustly executes sensor-driven cable routing, integrating vision for coarse localization with force feedback to handle geometric and mechanical uncertainties. This use case obtained the second place in the Neura Robotics Challenge 2025. A video is available in the following link: https://youtu.be/vRGSpFZ4Pgo?si=g5OYncd4LnIHBNLy

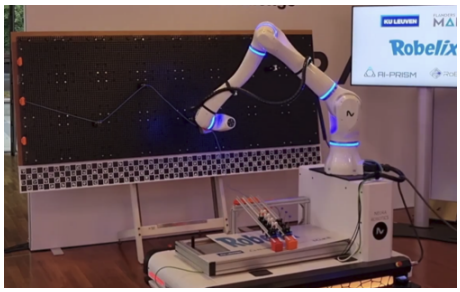

Figure 2: Application example setup.

References

[1] E. Aertbeliën, J. De Schutter, "eTaSL/eTC: A constraint-based Task Specification Language and Robot Controller using Expression Graphs" 2014 IEEE/RSJ IROS, 2014.

Learning UAV Navigation with Collision-Aware Predicted-Motion Embedding

Mohammad Javad Zallaghi, Bryan Convens, Kelly Merckaert, Adrian Munteanu, and Bram Vanderborght
BruBotics, Vrije Universiteit Brussel (VUB), and IMEC, Belgium
Email: Mohammad.Javad.Zallaghi@vub.be; Bryan.Convens@vub.be

1 Introduction to the problem

Navigation of autonomous Unmanned Aerial Vehicles (UAVs) in unknown, obstacle-cluttered environments requires motion planning under limited sensing and strict safety constraints. Although learning-based navigation approaches have demonstrated promising performance, many rely on purely reactive obstacle avoidance or indirect reward shaping (e.g., end-to-end learning or deep-learning-based sensor encoding) [1, 2], without explicitly incorporating predictive collision information into the learning process. This work introduces a conceptual framework for learning collision-aware UAV navigation over a predicted motion horizon.

2 Collision-aware navigation learning framework

Instead of feeding encoded depth observations into a deep learning model, we propose a method that performs collision checking on predicted UAV motion and uses these predictions as a learning signal for navigation. LiDAR and depth sensing are used to build an online map. A closed-form model of the nominally controlled UAV predicts forward motion to check trajectory-level collisions. This yields a computationally cheaper and more interpretable, explicitly collision-aware policy. Figure 1 visualizes the framework.

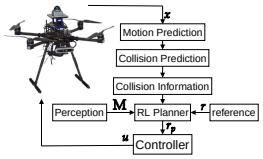

Figure 1: Framework of learning-based UAV navigation with collision-aware predicted-motion embedding.

3 Reinforcement learning to train planner

The planner is a Neural Network (NN) π_θ, that is trained via Reinforcement Learning (RL) to minimize the objective in (1) and to reach a feasible reference $\mathbf{r}$ in randomly cluttered environments. Training is conducted in AerialGym [3], where at each time step t, the policy outputs a planner reference $\mathbf{r_p}$ that minimizes squared Euclidean norm of reference tracking error $\|\mathbf{r}-\mathbf{p}\|_2^2$, planning effort $\|\mathbf{r_p}-\mathbf{p}\|_2^2$ with UAV position $\mathbf{p}$, and control effort $\|\mathbf{u}-\mathbf{u}_{\text{eq}}(\mathbf{r_p})\|_2^2$ with control input $\mathbf{u}$ and steady-state control effort $\mathbf{u}_{\text{eq}}$, while penalizing predicted safety-distance d_{safe} violations via minimum distance to environment $d_{\min}(\mathbf{x}, \mathbf{r_p}, \mathbf{M})$ as squared positive-part penalty of safety distance $\max(0,\, d_{\text{safe}} - d_{\min})^2$ computed from collision checking on the online occupancy map $\mathbf{M}$

$$\begin{aligned}\min_{\theta}\ \mathbb{E}\Big[\sum_{\tau=t}^{t+T}\gamma^{\tau}\Big(&w_r\|\mathbf{r}(t)-\mathbf{p}(\tau|t)\|_2^2+w_p\|\mathbf{r_p}(t)-\mathbf{p}(\tau|t)\|_2^2\\&+w_u\|\mathbf{u}(\tau|t)-\mathbf{u}_{\text{eq}}(\mathbf{r_p}(t))\|_2^2\\&+w_d\max\big(0,\, d_{\text{safe}}-d_{\min}\big(\mathbf{x}(\tau),\mathbf{r_p}(t),\mathbf{M}(t)\big)\big)^2\Big)\Big]\end{aligned} \tag{1}$$

where $\tau \in \{t,\ldots,t+T\}$ is prediction horizon time step, T is prediction horizon, $\gamma \in (0,1]$ is the discount factor, θ denotes the NN policy parameters, $\mathbb{E}[\cdot]$ is the expectation over trajectories induced by the policy, and $w_r, w_p, w_u, w_d \geq 0$ are weights.

4 Conclusion and contributions

This navigation framework proposed on a learning-based planner that is explicitly collision-aware, operating on predicted motion without the need to encode sensor data with deep learning models. Ongoing work investigates implementing this framework.

5 Acknowledgment

This work was supported by the EU-funded SPEAR project under Grant 101119774, by the Research Program Artificiële Intelligentie Vlaanderen from the Flemish Government, and by Research Foundation Flanders (FWO) postdoc fellowships 1287725N and 1283624N.

References

[1] Mihir Kulkarni and Kostas Alexis. Reinforcement learning for collision-free flight exploiting deep collision encoding. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15781–15788, 2024.

[2] Hang Yu, ChristopheDe Wagter, and Guido C. H. E de Croon. Mavrl: Learn to fly in cluttered environments with varying speed. *IEEE Robotics and Automation Letters*, 10(2):1441–1448, 2025.

[3] Mihir Kulkarni, Welf Rehberg, and Kostas Alexis. Aerial gym simulator: A framework for highly parallelized simulation of aerial robots. *IEEE Robotics and Automation Letters*, pages 1–8, 2025.

Matrix-Free Delassus Operations: Scalable and Memory-Efficient Algorithms

Ajay Suresha Sathya[1], Louis Montaut[1], Yann de Mont-Marin[1], and Justin Carpentier[1 12]

This abstract investigates efficient matrix-free algorithms for computing the Delassus operator and its damped inverse for articulated rigid-body systems with arbitrary kinematic loops and contact constraints. Matrix-free implementations of linear operators are highly beneficial when the explicit construction of a matrix or its inverse is computationally prohibitive. Explicitly computing the Delassus matrix can require up to $O(n^3+m^2n+mn^2)$ operations, where n is the number of degrees of freedom (DOFs) and m is the number of constraints. Typically, the Delassus matrix must also be factorized to solve for constraint forces or impulses, incurring an additional $O(m^3)$ cost. Current efficient recursive algorithms can compute the Delassus matrix in $O(n+m^2)$ operations but cannot handle loop-closure constraints or are computationally expensive for high-DOF robots in contact-rich scenarios.

The Delassus operator Λ^{-1} and its damped inverse Λ_R (with $R \in \mathbb{S}^m_{++}$) are defined as:

$$\Lambda^{-1} = J_c M^{-1} J_c^\top, \tag{1}$$

$$\Lambda_R = (R^{-1} + J_c M^{-1} J_c^\top)^{-1}, \tag{2}$$

where $M \in \mathbb{S}^n_{++}$ is the joint-space inertia matrix, $J_c \in \mathbb{R}^{m\times n}$ is the constraint Jacobian and $R \in \mathbb{S}^m_{++}$ is a diagonal damping matrix.

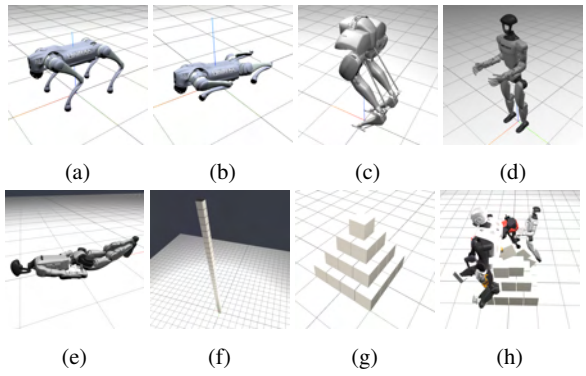

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 1: Illustrating the contact scenarios benchmarked in Table 1.

By leveraging recursive rigid-body dynamics, specifically the Loop-Constrained Articulated-Body Algorithm (LCABA) [1], our algorithms achieve linear memory and time complexity in practice. Extensive benchmarking confirms significant speedups compared to state-of-the-art explicit matrix methods [2], ranging from 2× up to 400×. While explicit methods remain competitive for systems with very few constraints, our approach is vastly superior in contact-rich scenarios and more efficient when downstream applications require frequent updates to damping parameters. The generality and efficiency of these algorithms facilitate their integration into physics engines for real-time control and simulation of complex robotic systems, which remains the focus of our future work. For more details, the reader is referred to [3].

Table 1: Computation time (in μs) and memory footprint (in kB) are reported in dedicated columns for the Delassus operator and its inverse on various robotic platforms and scenarios (Go2: $n_v = 18$, Cassie: $n_v = 32$, G1: $n_v = 35$, Boxes: $n_v = 180$, Pyramid+Robots: $n_v = 285$). The rows corresponding to the joint friction-enabled scenarios are shaded in gray. The proposed Matrix-Free method is compared against the explicit LTL-OSIM baseline.

Scenario	n_c	n_{con}	Delassus Operator (Λ^{-1}) Matrix-Free time	Matrix-Free mem	Explicit time	Explicit mem	Inverse Delassus Operator (Λ_R) Matrix-Free time	Matrix-Free mem	Explicit time	Explicit mem
Go2 standing	4	12	**1.55**	**26**	3.11	33	**1.64**	**29**	3.65	32
Go2 standing	4	24	**1.61**	**28**	4.64	47	**1.63**	**32**	5.95	43
Go2 fallen	9	27	**1.74**	**38**	5.89	62	**2.42**	**42**	7.63	56
Go2 fallen	10	42	**1.77**	**45**	8.10	97	**2.44**	**49**	12.00	81
Cassie	8	24	**2.94**	**53**	9.27	84	**4.26**	**69**	10.40	88
Cassie	8	50	**3.07**	**57**	16.00	121	**4.51**	**72**	21.72	143
G1 standing	8	24	**3.66**	**91**	9.37	195	**3.58**	**110**	10.60	235
G1 standing	8	53	**3.96**	**94**	18.53	254	**3.83**	**117**	25.00	334
G1 fallen	23	69	**4.27**	**47**	25.98	75	**6.22**	**62**	37.88	80
G1 fallen	25	101	**4.41**	**50**	42.09	107	**7.33**	**65**	70.40	127
Stack of boxes	120	360	**9.47**	**241**	1214	1400	**31.38**	**269**	2482	1850
Box pyramid	288	864	**14.88**	**755**	5925	9694	**65.20**	**826**	23092	15530
Pyramid + Robots	108	324	**20.34**	**473**	1511	4533	**37.12**	**730**	2430	5630

References

[1] A. S. Sathya and J. Carpentier, "Constrained articulated body algorithms for closed-loop mechanisms," *IEEE Trans. Robot.*, pp. 1–20, 2026.

[2] R. Featherstone, *Rigid body dynamics algorithms*. Springer, 2014.

[3] A. Suresha Sathya, L. Montaut, Y. de Mont-Marin, and J. Carpentier, "Matrix-Free Delassus Operations: Scalable and Memory-Efficient Algorithms," Jan. 2026, working paper or preprint. [Online]. Available: https://hal.science/hal-05457476

[1]Inria and Département d'Informatique de l'École Normale Supérieure, PSL Research University in Paris, 75013 Paris, France.

[2]emails: ajay.sathya@inria.fr louis.montaut@inria.fr yann.de-mont-marin@inria.fr justin.carpentier@inria.fr,

Tracklet association using a neural network model: long short-term approach

Ruth David[1], Jeroen Ploeg[2], Mohsen Alirezaei[2], and Tom van der Sande[1]
[1]Dynamics and Control group, TU Eindhoven
[2]Siemens Industry Software Netherlands B.V., Helmond

1 Introduction

Using recorded real-world data is essential for the evaluation of these driving functionalities. However, during data acquisition, objects may be occluded for a short period of time, resulting in incomplete trajectories. As a result of these occlusions, fragmented trajectories are formed, known as tracklets.These gaps make it difficult to reconstruct scenarios and thereby, system validation is complicated. The aim of the contribution is to associate the tracklets that belong to the same object using a long short-term memory (LSTM)-based model.

2 Proposed approach

It's assumed each trajectory contains a single occlusion, and the goal is to classify tracklet pairs that belong to the same object. The tracklet pairs are different combinations of two tracklets. An LSTM-based model is used to learn tracklet dynamics from sequential time, position, and velocity data leveraging gated memory to capture temporal dependencies. The network outputs match and non-match probabilities for each pair, which are refined into one-to-one associations using the Hungarian algorithm [1].

3 Long short term memory network

The network consists of an input layer with each neuron representing time t_r, positions $p_{x,r}$, $p_{y,r}$, and velocities $v_{x,r}$, $v_{y,r}$ of a tracklet pair r. For each r, the tracklets i and j are combined as potential matches. Tracklet i has t_i, $p_{x,i}$, $p_{y,i}$, $v_{x,i}$, $v_{y,i}$ and tracklet j has t_j, $p_{x,j}$, $p_{y,j}$, $v_{x,j}$, $v_{y,j}$. To illustrate, $p_{x,i}$ denotes all x-position samples of tracklet i. The target labels are either 1 (class match) or 0 (class non-match) for each pair. After the input layer, the LSTM layer is defined, with 64 neurons. The input data are transformed using weights, biases, and activation functions to be passed on to the next layer. This process continues for consecutive layers. Each neuron consists of hidden states, cell states, forget gates, input gates, and output gates enabling it to handle long-term memory by reducing the vanishing gradient problem. The LSTM layer is followed by three fully connected layers with 64, 32, and 2 neurons, respectively. ReLU activations are used in the first two layers which introduces non-linearity resulting in unormalize scores.

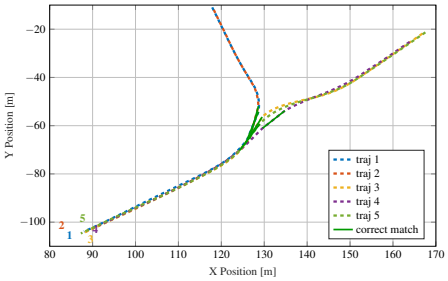

Figure 1: Tracklet association based on two-tracklet network (rounD data).

A softmax activation is used in the final layer generating normalized scores for both classes.

3.1 Two-tracklet network

The network introduced here is known as the two-tracklet network. Each input neuron represents a feature vector for each tracklet pair defined as,

$$F = \begin{bmatrix} t_i^\mathsf{T} & p_{x,i}^\mathsf{T} & p_{y,i}^\mathsf{T} & v_{x,i}^\mathsf{T} & v_{y,i}^\mathsf{T} \\ t_j^\mathsf{T} & p_{x,j}^\mathsf{T} & p_{y,j}^\mathsf{T} & v_{x,j}^\mathsf{T} & v_{y,j}^\mathsf{T} \end{bmatrix}. \quad (1)$$

The rows correspond to the two tracklets as a potential match, and columns represent feature inputs. This reduces the input dimensionality to five. Each neuron represents the time, positions, and velocities of a tracklet pair, i, j, resulting in five input neurons. A cross-entropy loss function is used.

4 Results

The average metric scores (recall, precision, and F1-score) based on the rounD dataset is used for evaluations. The network perform well, achieving metric values above 0.9. The non-match class achieves marginally higher metric score than the match class. The results based on a scenario with 10 tracklets (5 vehicle trajectories) is given in Fig. 1, which shows that all tracklets are associated correctly.

References

[1] R. David, J. Ploeg, M. Alirezaei, and T. van der Sande, "Multi-Vehicle Tracking Through Occlusions Using a Long Short-Term Memory Approach" IEEE Transactions on Intelligent Transportation Systems,2025, submitted.

Velocity-Through-Water Estimation for Overactuated Surface Vessels Using Electric-Drive Telemetry

Jef Billet[1], Herman Bruyninckx[2,3], Peter Slaets[1]
[1]IMP Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
[2]ROB Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
[3]Flanders Make@KU Leuven, Leuven, Belgium
Email: jef.billet@kuleuven.be

1 Introduction

Recent technological and environmental drivers have pushed surface vessel propulsion toward electrified architectures. These architectures are now found in applications ranging from small marine drones to large (hybrid-)electric ships. Well-known examples include the ferries operated by *De Waterbus* in Antwerp and *De Driehoeksveer* in Rotterdam. For many of these platforms, electrification and overactuation coincide. They become equipped with multiple azimuth thrusters, resulting in substantially improved maneuverability compared to their under-actuated counterparts. This can reduce their (un)docking times and improve operational energy efficiency by providing room for maneuvering optimization. In this setting, the vessel's *through-water-velocity* is a particularly valuable state, as it directly affects propeller inflow and, consequently, thrust and power consumption, making it essential for both monitoring and energy-efficient control tasks. And yet, it remains expensive and difficult to measure in practice. In this work, we exploit available electric-drive telemetry to **(i)** calibrate a vessel's drivetrain and **(ii)** estimate its through-water-velocity. Both steps are demonstrated and validated on the twin-azimuth electric test vessel Maverick [1].

2 Methodology

We formulate the estimation problem from a hydrodynamic loading perspective. For electric thruster i, the measured DC power is modeled as a composition of (1) the propeller torque law and (2) prop-shaft powertrain efficiencies,

$$Q_{\text{prop},i} = \rho D^5 n_i^2 K_Q(J_{ax,i}), \tag{1}$$

$$P_{dc,i} = f_{\text{pt}}(n_i, Q_{\text{prop},i}), \tag{2}$$

where n_i is the (controlled) propeller speed, ρ is the water density, D is the propeller diameter, and $f_{\text{pt}}(\cdot)$ captures calibrated drivetrain, motor, and inverter losses. The axial advance ratio is

$$J_{ax,i} = \frac{V_{a,i}}{n_i D}, \tag{3}$$

with $V_{a,i}$ the axial inflow at the thruster. The main idea is to group these expressions for all thrusters, yielding a nonlinear map $\mathbf{P}_{dc} = h(\mathbf{v}_{tw}, \mathbf{n}, \boldsymbol{\alpha})$ from the body-frame *through-water* velocity $\mathbf{v}_{tw} = [u_{tw}\ v_{tw}]^T$, to the vector of power measurements. The informativeness of h clearly depends on the instantaneous azimuth configuration $\boldsymbol{\alpha}$. Each thruster's power is driven by the component of $\mathbf{v}_{tw}$ projected onto its thrust axis, so the sensitivity of $\mathbf{P}_{dc}$ to $\mathbf{v}_{tw}$ is inherently directional. When the active thrusters are (near-)collinear in azimuth, these directional sensitivities become redundant. Consequently, the local sensitivity of $\mathbf{P}_{dc}$ to $\mathbf{v}_{tw}$ degrades, making the inversion ill-conditioned, and in the limit locally non-invertible. We, therefore, estimate $\mathbf{v}_{tw}$ with a sliding-window Moving Horizon Estimator (MHE), exploiting temporal variation in α_i and n_i.

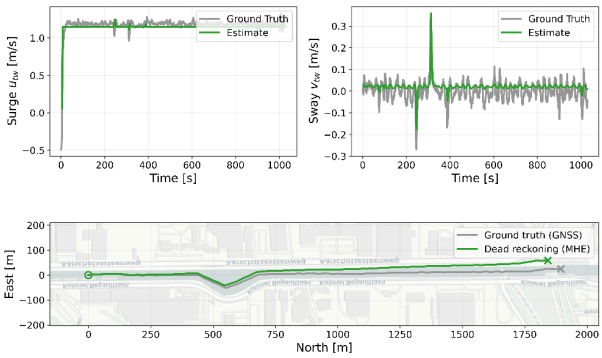

Figure 1: Experimental run with the Maverick test vessel through the Leuven-Dijle canal. The surge velocity (top-left) and sway velocity (top-right) estimates are used for dead reckoning, the resulting trajectory (bottom) is visualized and compared to GNSS-RTK ground-truth.

3 Results

Experiments took place in a waterway where the ambient current was small relative to vessel speed, such that the measured over-ground GNSS-RTK trajectory and velocities provide a practical ground truth for the evaluation of the power-based through-water-velocity estimates. Fig. 1 presents the power-based $\mathbf{v}_{tw}$ estimation, demonstrating how an electrified propulsion system can function as a proprioceptive $\mathbf{v}_{tw}$ estimator. The estimated velocities are further employed for dead reckoning by integrating with the measured heading to obtain a world-frame trajectory for comparison. Future work includes incorporating data-driven models for hard-to-model effects (e.g., thruster-hull interactions) and fusing additional sensors to estimate the ambient current online.

References

[1] Y.-Y. Zhang, J. Shuai, J. Billet, and P. Slaets, "Design and Build of an Autonomous Catamaran Urban Cargo Vessel," J. Phys.: Conf. Ser., vol. 2618, no. 1, p. 012002, Oct. 2023, doi: 10.1088/1742-6596/2618/1/012002.

Anomaly Detection in Electric Vehicles under Dynamic Wireless Charging: Challenges and Research Gaps

Junhwi Mun[1]

[1] Control Systems Technology Group, Dept. of Mechanical Engineering,
Eindhoven University of Technology (TU/e), The Netherlands
j.mun@tue.nl

1 Introduction

Dynamic Wireless Power Transfer (DWPT) enables Electric Vehicles (EVs) to be charged wirelessly while driving via electromagnetic coupling between road-embedded transmitting coils and onboard receiver pads, as shown in Fig. 1. Compared to conventional EVs, DWPT-equipped vehicles incorporate onboard components such as receiver pads and high-frequency power electronics, introducing additional electromagnetic and thermal interactions. This paradigm reduces charging downtime and alleviates battery capacity requirements, yet vehicle-side anomalies related to onboard power reception and energy storage systems under DWPT operation remain largely unexplored.

EV anomalies such as battery degradation and thermal runaway are well known in conventional driving environments. However, DWPT introduces tightly coupled electromagnetic, thermal, and mechanical interactions during simultaneous driving and charging, making DWPT-specific vehicle-side anomalies difficult to distinguish from normal EV behavior.

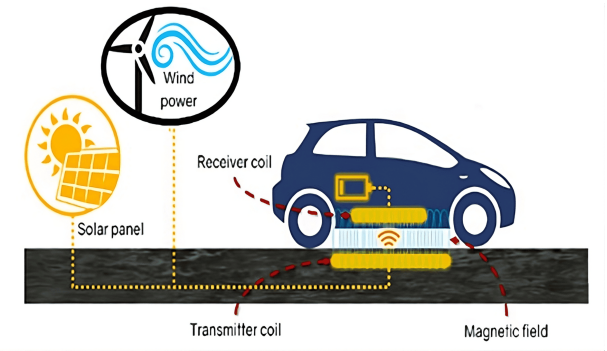

Figure 1: Overview of EV charging via DWPT [1].

2 Vehicle-Side Challenges and Anomaly Characteristics

Despite increasing interest in DWPT, vehicle-side anomaly detection remains underdeveloped, as most existing approaches focus on infrastructure components rather than onboard receiving systems, and assume driving and charging as independent operating modes.

In this study, vehicle-side anomalies of interest include foreign object detection (FOD), battery-related thermal runaway risks, and abnormal thermal behavior caused by receiver pad losses such as ferrite or litz wire degradation. These anomalies differ in required response time: FOD events demand rapid online detection and immediate power interruption, whereas battery- or receiver-related thermal anomalies require early warning to prevent long-term degradation or hazardous conditions.

A key challenge lies in the context-dependent interpretation of observable symptoms. Signals that would be classified as anomalies in conventional EV operation may arise as inherent characteristics under DWPT conditions due to high-frequency power transfer and electromagnetic coupling. Without distinguishing the underlying root causes, such symptoms risk being misclassified as faults, leading to unnecessary interventions or overlooked risks.

Therefore, reliable anomaly detection in DWPT-equipped EVs requires not only identifying abnormal patterns but also classifying their root causes by separating DWPT-induced operational characteristics from genuinely hazardous conditions. This challenge is further amplified by coupled electromagnetic–thermal feedback mechanisms, where degraded receiver efficiency can accelerate thermal stress and battery aging.

3 Methodological Trends and Open Research Gaps

Vehicle-side anomaly detection under DWPT faces a central challenge in distinguishing DWPT-induced electromagnetic–thermal effects from normal driving behavior under severe data scarcity. Recent methodological trends indicate that this problem is increasingly addressed using data-driven and AI-based approaches, as the coupled dynamics during simultaneous driving and charging are difficult to model accurately using conventional methods alone.

To cope with limited availability of labeled vehicle-side fault data, existing works commonly adopt normal-operation modeling and hybrid AI frameworks that integrate physical system knowledge, employing physics-informed learning and physically meaningful feature design to improve interpretability.

References

[1] R. C. Majhi et al., "Optimal allocation of dynamic wireless charging facility for electric vehicles," *Transportation Research Part D: Transport and Environment*, 2022.

Design of an optimal positively stabilizing state feedback for a spatially discretized heat equation

Violaine PIENGEON and Joseph J. WINKIN
Namur Institute for Complex Systems (naXys), Department of Mathematics, UNamur, Belgium
Email: violaine.piengeon@unamur.be, joseph.winkin@unamur.be

1 Introduction

Positive systems are dynamical systems that leave the nonnegative cone of the state space invariant. The study of such systems is of great importance because the positivity property is relevant in many fields such as (bio)chemical engineering, biomedical sciences, or epidemiology. Indeed, physical quantities such as temperature (K), concentration or population density are all nonnegative quantities that can be represented by state variables. When one wants to stabilize a positive system, it is necessary to design state feedbacks such that the closed-loop system maintains this natural positivity property, otherwise it would not have any meaning with respect to the model.

Here, we would like to present recent results on the design of an optimal (according to a well-chosen optimality criterion) positively exponentially stabilizing state feedback for a particular class of LTI systems.

2 Discretized heat equation

Let us consider the following n-dimensional LTI system

$$\dot{x}^{(n)}(t) = A^{(n)}x^{(n)}(t) + b^{(n)}u(t), \quad x^{(n)}(0) = x_0^{(n)}, \qquad (1)$$

where

$$A^{(n)} = \begin{pmatrix} -p_2 & p_2 & 0 & \cdots & 0 \\ p_2 & -2p_2 & p_2 & & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & & p_2 & -2p_2 & p_2 \\ 0 & \cdots & 0 & p_2 & -p_2 \end{pmatrix} \in \mathbb{R}^{n\times n},$$

is a tridiagonal Metzler matrix, $b^{(n)} = p_1\mathbf{e}_1^{(n)} \in \mathbb{R}^n$, $p_1 = (n-1)/L$, $p_2 = D_a p_1^2$, $\mathbf{e}_1^{(n)}$ is the first vector of the canonical basis of $\mathbb{R}^n$, and $L, D_a > 0$. The variables $x^{(n)}$ and u in (1) denote the state trajectory and input, respectively. This system is a spatially discretized (via a finite-difference scheme) version of the well-known heat equation on the spatial domain $[0,L]$, with diffusion coefficient D_a and with a boundary control u in Neumann boundary conditions.

3 Positively exponentially stabilizing state feedbacks

As mentioned in [1, Section 3.2], (1) is positive but not exponentially stable. However, for n sufficiently large, a state feedback law of the form

$$u(t) = p_1\kappa^{(n)}\mathbf{e}_1^{(n)\top}x^{(n)}(t) =: K^{(n)}x^{(n)}(t), \quad t \geq 0, \qquad (2)$$

is positively exponentially stabilizing for (1), or equivalently the matrix $A^{(n)}_{\kappa^{(n)}} := A^{(n)} + b^{(n)}K^{(n)}$ is Metzler and exponentially stable if, and only if, $\kappa^{(n)} < 0$ (see [1, Theorem 4]).

4 Optimal control problem

Since there exists an infinite number of state feedbacks of the form (2) that are positively exponentially stabilizing for (1), our goal is to find the best one according to a well-chosen optimality criterion. More precisely, we consider, for any $\alpha > 0$, the following optimal control problem, denoted $\mathscr{S}_\alpha^{(n)}$:

$$\begin{aligned} \min_{\kappa^{(n)}\in\mathbb{R}} \quad & \kappa^{(n)2}, \\ \text{s.c.} \quad & A^{(n)}_{\kappa^{(n)}} \text{ is Metzler,} \\ & A^{(n)}_{\kappa^{(n)}} \text{ is } \beta\text{-exponentially stable}, \forall\beta \in (-\alpha, 0]. \end{aligned}$$

The $\mathscr{S}_\alpha^{(n)}$ problem may be viewed as a trade-off between a sufficiently high stability margin and the advantages offered by a sufficiently low gain state feedback, while maintaining positivity of the state trajectories.

We shall report simple feasibility conditions on α, as well as the solution of the $\mathscr{S}_\alpha^{(n)}$ problem. Indeed, the latter problem has a unique solution $\kappa_*^{(n)}$, which is given by

$$\kappa_*^{(n)} = \frac{\sqrt{D_a}}{p_1} \cdot \frac{\displaystyle\sum_{l=0}^{n-1} \binom{n+l}{2l+1} \frac{(-1)^{l+1}}{p_2^{l+\frac{1}{2}}} \alpha^{l+1}}{\displaystyle\sum_{l=0}^{n-1} \binom{n+l-1}{2l} \frac{(-1)^l}{p_2^l} \alpha^l}. \qquad (3)$$

This form of the solution has the advantage to be analytic and explicit, and so easily computable for any n. The results will be illustrated by numerical simulations. A special focus will be put on the asymptotic behavior of the problem and its solution as n tends towards ∞.

References

[1] J.N. Dehaye and J.J. Winkin, "Parameterization of positively stabilizing feedbacks for single-input positive systems," Systems & Control Letters, 98:57–64, 2016.

An Optimal Power Management Policy for a Hydrogen-based Hybrid Aircraft Engine[1]

Faezeh Pak, Uto Perra, Frank Willems, Theo Hofman, Mauro Salazar
Control Systems Technology
Eindhoven University of Technology
Emails: {f.pak, u.perra, f.p.t.willems, t.hofman, m.r.u.salazar}@tue.nl

1 Abstract

The aviation industry has committed to become climate neutral by 2050, partly by adopting novel powertrain concepts. The FlyECO team supports these efforts by developing the hybrid electric architecture, as illustrated in Fig. 1, consisting of a hydrogen gas turbine and solid oxide fuel cell (SOFC) that are thermodynamically coupled, to power a propeller via a turboprop and an electric motor, respectively [1]. The gas turbine includes a low- and high-pressure compressor (HPC), which compress ambient air to a higher pressure. A bypass valve strategically directs part of the compressed air to the SOFC to undergo an electrochemical reaction with hydrogen and to produce electrical power and steam, while the remaining air enters the combustion chamber alongside the unreacted hydrogen, and air from SOFC, and steam, which mitigates the NO_x emission. The resulting high-temperature gas expands through two turbines to rotate the shafts, the first turbine rotates compressors shaft while the low-pressure turbine rotates the propeller shaft.

This re-circulation causes thermodynamic and mechanical coupling between the subsystems. Consequently, such a complex aero engine requires a high-level control policy to coordinate the components' operation and to minimize hydrogen consumption. Previously, a fuzzy-PID multi-loop controller [2] was proposed for an SOFC and gas turbine integration, without considering the optimal operation, while optimal control policies, such as the Pontryagin's minimum principle (PMP) [3], were studied for energy management in hybrid vehicles. The analytical non-smooth PMP is solved by defining a Hamiltonian consisting of the objective function and the indicator functions to enforce the problem constraints. To the best of the authors' knowledge, optimal supervisory power management policies, such as PMP, with the aim of minimum fuel consumption in SOFC and gas turbine integrated propulsion systems has not yet been studied for aerospace applications.

In this study, we primarily formulate a fuel minimization problem for a hydrogen hybrid electric aero engine, explicitly incorporating key performance constraints that capture the complex thermodynamic and mechanical couplings between components, such as the influence of bypass valve on overall system dynamics, to ensure a realistic and practical optimization framework. Second, we employ a non-smooth PMP to derive an optimal high-level supervisory controller to coordinate system components. Furthermore, we evaluate the performance of the proposed control strategy through numerical simulations across the flight mission to demonstrate its effectiveness as well as its ability to optimize energy consumption.

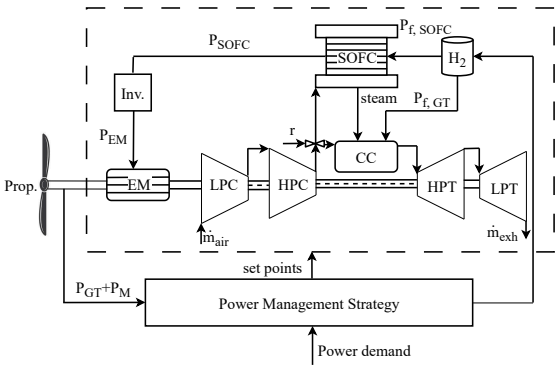

Figure 1: Schematic view of the aero engine, including a gas turbine, the SOFC, and an electrical motor (EM).

[1]Research funded by the European Union (ID 101138488) and by the UK Research and Innovation (ID 10106893). Views expressed are those of the author(s) only and do not necessarily reflect those of the funding entities. Neither the European Union nor the granting authority can be held responsible for them.

References

[1] FlyECO. (2024) Future enabling technologies for hydrogen-powered electrified aero engine for clean aviation. FlyECO. [Online]. Available: https://flyeco-european-project.eu/

[2] X. Wang, X. Lv, X. Mi, C. Spataru, and Y. Weng, "Coordinated control approach for load following operation of sofc-gt hybrid system," *Energy*, 2022.

[3] J. Ritzmann, A. Christon, M. Salazar, and C. Onder, "Fuel-optimal power split and gear selection strategies for a hybrid electric vehicle," *SAE Technical Paper*, 2019.

Model-Based Fault-Tolerant Control of Multi-Engine Rocket Propulsion Systems

Felix Ebert
Chair of Space Mobility and Propulsion
TU Munich
80333 Munich, Germany
Email: felix.ebert@tum.de

Chiara Manfletti
Chair of Space Mobility and Propulsion
TU Munich
80333 Munich, Germany
Email: chiara.manfletti@tum.de

1 Introduction

Liquid rocket engines are thermo-fluid systems in which fuel and oxidizer react in a combustion chamber to generate thrust. The thrust is what gives spacecrafts control authority. From a control perspective, the primary objective is accurate thrust tracking as commanded by the flight controller.

In addition, the oxidizer-to-fuel ratio (ROF) must be regulated precisely. The ROF affects both propulsion efficiency and component loading. Violating admissible limits can lead to severe hardware damage and subsequently mission loss. Hence, rocket engine control must ensure thrust tracking while strictly satisfying ROF constraints.

The propulsion system is responsible for most launch vehicle failures. According to [1], 54% of the launch vehicle failures between 2006 and 2021 were caused by the propulsion system (based on 57 reported launches with anomalies). This motivates the development of fault-tolerant propulsion control strategies to enhance mission reliability.

2 Related Work

State-of-the-art rocket engine control relies on two PI control loops with time-scale separation: the oxidizer valve regulates ROF, while the fuel valve tracks combustion chamber pressure as a proxy for thrust [2]. Although effective under nominal conditions, this architecture does not explicitly accommodate faults.

Fault-tolerant approaches have been investigated within the Intelligent Control System (ICS) framework [3]. Model-based fault accommodation, e.g., for a frozen oxidizer valve, has been demonstrated using pre-designed nominal and faulty controllers combined via controller blending. A command generator adapts references according to remaining actuator capabilities. However, this approach assumes engine-level actuator redundancy, which is usually not available in mass-optimized rocket engines.

3 Fault Tolerant Propulsion Control

We propose the fault-tolerant propulsion control framework shown in Fig. 1, explicitly targeting multi-engine configurations.

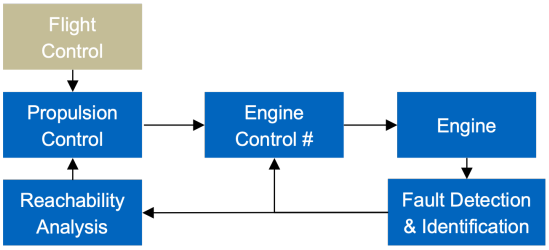

Figure 1: Fault Tolerant Propulsion and Engine Control Framework.

At engine level, a convex model predictive controller (MPC) performs set-point tracking while enforcing combustion chamber pressure and ROF constraints. Actuator faults are detected and isolated by a dedicated Fault Detection and Identification (FDI) module.

Based on the identified actuator limitations, an online reachability analysis determines the admissible operational envelope of each engine. This remains computationally tractable due to monotone behavior in the relevant nonlinear engine dynamics. The propulsion-level controller then allocates the total thrust demand across the engines according to their individual feasible envelopes, enabling continued operation despite actuator faults.

References

[1] L. A. Fernández, C. Wiedemann, and V. Braun, "Analysis of Space Launch Vehicle Failures and Post-Mission Disposal Statistics," *Aerotec. Missili Spaz.*, vol. 101, pp. 243–256, 2022.

[2] E. W. Otto and R. A. Flage, "Control of combustion-chamber pressure and oxidant-fuel ratio for a regeneratively cooled hydrogen-fluorine rocket engine," NASA Technical Note D-82, National Aeronautics and Space Administration, Washington, DC, 1959.

[3] J. L. Musgrave, T.-H. Guo, E. Wong, and A. Duyar, "Real-time accommodation of actuator faults on a reusable rocket engine," *IEEE Transactions on Control Systems Technology*, vol. 5, no. 1, pp. 100–109, 1997.

Safe Abstraction-Based Controller Synthesis From Natural Language

Amir Bayat
ICTEAM Institute
UCLouvain
Louvain-la-Neuve, Belgium
Email: Amir.Bayat@UCLouvain.be

Raphaël M. Jungers
ICTEAM Institute
UCLouvain
Louvain-la-Neuve, Belgium
Email: Raphael.Jungers@UCLouvain.be

1 Introduction

Modern cyber-physical systems are expected not only to operate safely but also to provide intuitive and accessible interfaces for non-expert users. Ideally, control objectives could be specified in Natural Language (NL), enabling a robot, an autonomous vehicle, or, more generally, a safety-critical system to achieve them in a correct and safe manner. Addressing this vision requires bridging high-level NL specifications with certified control synthesis methods.

Large Language Models (LLMs) offer a promising interface for interpreting NL specifications; however, they do not provide formal guarantees of correctness or reliability and are therefore not suitable for direct use in planning and controller design, which involve complex reasoning. To mitigate this limitation, we restrict the role of LLMs to smaller, well-defined reasoning tasks and tune them to act as experts in handling these simpler subtasks. This task decomposition significantly improves the overall success rate. Early attempts at employing LLM-based agents for linear controller synthesis have been reported in [1].

To ensure correctness and safety despite increasing system complexity, we rely on Abstraction-Based Controller Design (ABCD) [2]. In ABCD, the continuous (infinite) state and input spaces are discretized according to a suitable relation, yielding a bisimilar finite-state approximation of the original system, referred to as an abstract system. Control synthesis is performed in this abstract domain, after which the resulting abstract controller is concretized to control the original system. This approach scales well with complexity while providing formal safety guarantees.

Building on this foundation, we propose an LLM-enhanced pipeline enabling interactive ABCD synthesis. The process begins with user requirements expressed in NL, which are translated into Linear Temporal Logic (LTL) specifications using LLMs, including both fine-tuned models [3] and commercially available LLMs. From the resulting LTL formula, a nondeterministic Büchi automaton is constructed. By performing a synchronous product between the abstract transition system and the Büchi automaton, the control synthesis problem is solved, yielding a controller that satisfies the specified requirements.

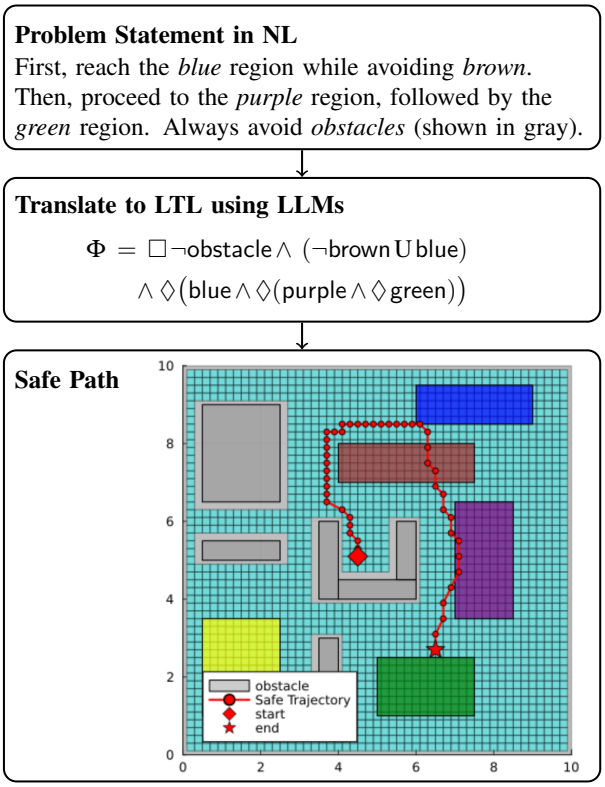

Figure 1: From NL specification to LTL and a safe path.

References

[1] Guo, X., et al. "Controlagent: Automating control system design via novel integration of LLM Agents and domain expertise," in arXiv preprint arXiv:2410.19811, 2024.

[2] A. Girard, G. Pola, P. Tabuada. "Approximately bisimilar symbolic models for incrementally stable switched systems," in IEEE Transactions on Automatic Control, 2009.

[3] Chen, Y., et al, "NL2TL: Transforming Natural Languages to Temporal Logics using Large Language Models," in Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2023.

Supervisor Synthesis for Multi-Purpose Systems

M. Minkenberg[1,*], M. A. Reniers[1], M. A. Goorden[1],
J. M. van de Mortel-Fronczak[1] and W. J. Fokkink[2]
[1]Control Systems Technology, Eindhoven University of Technology, The Netherlands
[2]Theoretical Computer Science, Vrije Universiteit Amsterdam, The Netherlands
[*]Email: m.minkenberg@tue.nl

1 Introduction

With rising sea levels, countries at risk of flooding have built various structures to hold back water. Because failures of such critical water infrastructures have a huge socioeconomic impact, reliable control of their systems is crucial [1]. Figure 1 shows a storm surge barrier in Denmark.

The supervisory control of storm surge barriers is challenging due to inherently conflicting control goals. On one hand, a barrier must close to prevent storm surges from reaching inland areas or block the water flow; one the other, it must open to let ships through, enable fish migration between river and sea, or regulate salinity. As a result, there is no single, fixed set of desired states: they vary with changing environmental conditions. A storm surge barrier controller must reflect this multi-purpose nature.

2 Multi-purpose as modes

A straightforward way to implement this multi-purpose nature for supervisor synthesis is to model each purpose as a 'mode' automaton. Each mode is either active or inactive and can be (de)activated by the controller based on environmental conditions. In Figure 2, a model is shown that combines a 'storm' mode with a simple gate that can open and close. In this example, the only marked (desired) state is when the gate is open and storm mode inactive.

This choice of marked states purposefully contains an issue. In both states where storm mode is active, all sequences of events to a marked state include `storm_deactivate`,



Figure 1: Picture of the barrier at *Hvide Sande*, from [2].

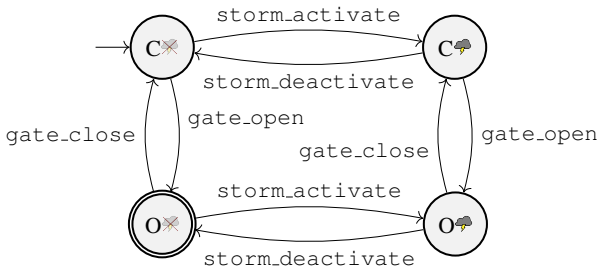

Figure 2: Model with an unfavorable choice of marked states.

which implies waiting for the storm conditions to end. It could be argued that the supervisor should not rely on `storm_deactivate` to reach a marked state. However, supervisor synthesis allows both modes, even if one of them is not actually able to reach a desired state.

Thus, the events representing the (de)activation of a mode should not be considered for reaching a marked state. One method is to label the mode (de)activation events as non-progressive and do synthesis with only progressive events, as described in [3]. Another method is to consider the two mode states (storm and no storm) separately. After all, it is desired that each mode is nonblocking in isolation. The latter method is preferred, as it is more directly linked with the intention of the modeler.

This novel method is an addition to synthesis. In the given example, states with storm mode active would become bad states during synthesis, because they have no path to a marked state while storm mode remains active. This causes the modeler to rethink their definition of marked states and modes, contributing to a more reliable controller.

Acknowledgment This work was supported as part of STORM_SAFE, an Interreg North Sea project co-funded by the European Union.

References

[1] M. Walraven, K. Vrolijk, and B. B. Kothuis, "Chapter 20 - Design, maintain and operate movable storm surge barriers for flood risk reduction," in *Coastal Flood Risk Reduction*, Elsevier, pp. 271–286, 2022.

[2] Kystdirektoratet. (2025). *Hvide Sande og Thorsminde sluser* [Online]. Available: https://kyst.dk/hav-og-anlaeg/vedligehold-af-havneanlaeg-sluser-daemning-og-sejlloeb/hvide-sande-og-thorsminde-sluser

[3] S. Ware and R. Malik, "Progressive events in supervisory control and compositional verification", *Control Theory and Technology*, vol. 12, no. 3, pp. 317–329, 2014

Supervisory Control of Timed Discrete Event Systems with Partially Uncontrollable Time-interval Events [1]

Mariana Guimarães Marques [2,3] , Martijn Goorden [2], João Carlos Basilio [3], Michel Reniers [2]

1 Introduction

The supervisory control theory (SCT) [1] applied to Discrete Event Systems (DES) [2] aims to build a supervisor that, given a system plant P and a set of requirements R, restricts the behavior of P to the desired behavior of R. In this context, the events in a system can be either controllable or uncontrollable. The supervisor, when controlling the plant, can only disable controllable events in order to meet the requirements. Ideally, the supervisor is guaranteed to be controllable, nonblocking and maximally permissive. SCT was originally introduced for applications to systems modeled by finite automata, where there is no time information on the occurrence of events [2].

In order to include real-time applications, a class of timed DES is proposed in [3]. In this work, we use a class of timed DES called Time-Interval Automata (TIA) [4] (Fig. 1), which is a subclass of the timed automata formalism presented in [3]. In this class, even though events are still assumed to occur asynchronously, they must occur within time intervals, and each transition is associated with an interval. There exists a unique real-time clock that resets at every event occurrence.

Regarding SCT, we introduce partially uncontrollable time-interval events, which are both controllable and uncontrollable, depending on the time interval. For such systems, we define controllability and propose an algorithm to synthesize a maximally permissive controllable supervisor.

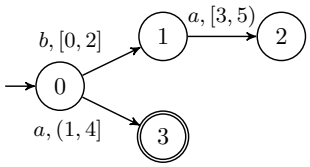

Figure 1: Example of a TIA.

[1]This work was supported in part by the Coordenação de Aperfeioamento de Pessoal de Nível Superior (CAPES), Finance Code 001. The work of João C. Basilio was supported in part by the Brazilian Research Council (CNPq) under Grant 316881/2021-0, and in part by the Carlos Chagas Foundation, Fundação de Amparo à Pesquisa do Estado do Rio de Janeiro (FAPERJ), under Grant E-26/201.144-2022.

[2]The authors are with Department of Mechanical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands m.guimaraes.marques@tue.nl

[3]The authors are with COPPE - Electrical Engineering Program, Federal University of Rio de Janeiro, 21949-900, Rio de Janeiro, Brazil mariana.marques@coppe.ufrj.br

2 Partially Uncontrollable Events

Since the occurrence of an event is constrained to a time interval, it might be the case that some event is both controllable and uncontrollable depending on the time it occurs within the interval. For a practical example, consider a simple system consisting of a tram that arrives at station S. Among all the events that can possibly occur after the tram arrives at the station, consider the event that models the *closing of doors*, denoted as d. Event d occurs in the interval $[0, 60]$ seconds. The tram must stay at the station for at least 20 seconds, and thus, within the time interval $[0, 20)$, event d is disabled since the conductor cannot close the doors. The conductor gains control of the doors in the next 30 seconds, which means that they can close the door as soon as no passengers are seen exiting or entering the tram; therefore, during the time interval $[20, 50)$, the event d is enabled. However, in the last 10 seconds, the closing of the doors, although it can happen, will close without the conductor intervention, and so, in the interval $[50, 60]$, the event d is uncontrollable. Consequently, event d is both controllable and uncontrollable, depending on the time interval, i.e. partially uncontrollable.

3 Contributions

Considering systems modeled with TIA and that events can be partially uncontrollable, we propose in this work: *(i)* definition of controllability, in terms of time-interval language; *(ii)* definition of supremal controllable time-interval sublanguage and *(iii)* and algorithm that synthesizes the supervisor, which is nonblocking and marks the supremal controllable time-interval sublanguage.

References

[1] Ramadge, Peter JG, and Walter Murray Wonham. "The control of discrete event systems." Proceedings of the IEEE 77.1 (2002): 81-98.

[2] C. G. Cassandras and S. Lafortune, "Introduction to Discrete Events Systems", Springer, 2008.

[3] R. Alur and D. Dill, "Automata for modeling real-time systems," in 17th ICALP, Warwick University, England, Proceedings 17. Springer, 1990, pp. 322–335.

[4] M. G. Marques, R. J. Barcelos, and J. C. Basilio, "The use of time-interval automata in the modeling of timed discrete event systems and its application to opacity," IFAC-PapersOnLine, vol. 56, no. 2, pp. 8654–8659, 2023.

Fast Tilt-Based Electron Microscope Alignment Using A-Optimal Sensor Scheduling

J.S. van Hulst[1], E.M. Franken[2], B.J. Janssen[2], W.P.M.H. Heemels[1], and D.J. Antunes[1]

[1]Control Systems Technology Section, Eindhoven University of Technology, The Netherlands
[2]Thermo Fisher Scientific, Eindhoven, the Netherlands

1 Motivation

Transmission electron microscopy (TEM) enables atomic-scale imaging with applications in numerous fields, including materials science, semiconductors, and life sciences. The electron beam is shaped by electromagnetic lenses, which suffer from high sensitivity to environmental conditions and manufacturing inconsistencies. The resulting imperfections in the electron beam are known as aberrations, causing reduced contrast and resolution. Aberrations can be compensated by adjusting the electron optical components, but the aberrations are not directly measurable and drift over time. As a result, fast and accurate aberration estimation is highly desirable.

2 Problem Formulation

A key physical insight is that deliberately tilting the electron beam in the presence of aberrations causes an apparent image shift [1]. These tilt-induced image shifts can be measured from TEM images and then used to estimate the underlying aberration coefficients. Crucially, the speed and accuracy of aberration estimates depends on the choice of beam tilt sequence, which motivates the research question: *what sequence of tilts results in the fastest and most accurate aberration estimates?*

3 Approach

We employ a Kalman filter (KF) to estimate aberrations from a sequence of tilt-induced image shifts, which naturally handles temporal drifting. We formulate the tilt sequence selection as a sensor scheduling problem. The tilt sequence is optimized by minimizing the trace of the KF's predicted error covariance (A-optimality). A key result is that this cost is independent of actual measurements [2], thus enabling offline optimization. The resulting non-convex problem is solved using gradient-based receding-horizon optimization with multi-starts. Fig. 1 shows how optimized patterns reduce estimation uncertainty faster than naive approaches.

4 Results

Experimental validation on a real TEM demonstrates that optimized tilt patterns significantly outperform naive approaches, achieving the same estimation performance in approximately 42% less time. The realized variance in aberration estimates closely matches the Kalman filter predictions, indicating accurate modeling. Overall, the proposed method reduces the aberration correction procedure from several minutes to less than a minute. Fig. 2 illustrates the improvement in image quality after applying the proposed aberration estimation and correction method.

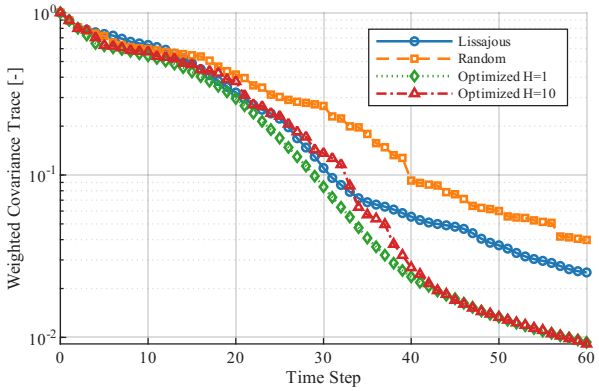

Figure 1: Weighted posterior covariance trace at each time step for different tilt patterns. Optimized patterns (myopic H=1 and receding-horizon H=10) reduce estimation uncertainty significantly compared to Lissajous or random patterns.

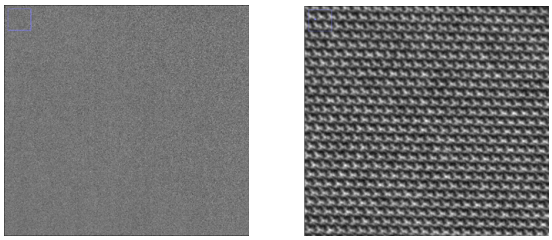

Figure 2: TEM images before (left) and after (right) aberration correction using the proposed method on a SiN sample, showing substantially improved contrast.

References

[1] K. J. Hanszen and L. Trepte, "Die Kontrastübertragung im Elektronenmikroskop bei partiell kohärenter Beleuchtung," 1971.

[2] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.

Informative Experiment Design for Parameter Estimation of the Doyle-Fuller-Newman Battery Model

E. Weldeghebreal, F.A. le Roux, M.C.F. Donkers
Control Systems Group, Eindhoven University of Technology, Netherlands
Email: {e.weldeghebreal,f.a.l.le.roux,m.c.f.donkers}@tue.nl

1 Introduction

Due to the electrification of transport and an increase in the need for energy storage methods, Lithium-ion batteries have received significant attention. Li-ion battery models model the behaviour and internal states of the battery cell, which are used to estimate the State-of-Charge (SoC) and State-of-Health (SoH) of the cell. The SoC and SoH estimates are used in Battery Management Systems (BMS) to improve the lifecycle of battery by reducing harmful usage of the battery. The Doyle-Fuller-Newman (DFN) model is a physics-based battery model capable of simulating the internal state of the Li-ion cell [1]. The DFN model requires a large number of parameters to execute, which are estimated when not known. Estimating the parameters of the DFN model leads to accurate results for a discharge down to 3.4 V [2]. This work aims to improve the accuracy of the parameter estimation procedure by increasing the information available on the parameter estimates in the experiments used in the estimation procedure.

2 Proposed Method

The parameter estimation procedure found in [2] is used to estimate the parameters of the DFN model. To design the experiments used by this parameter estimation procedure, a class of signals developed by [3] is used to generate a number of different input currents. The Fisher Information is used to define the information available on the model parameters.

The Fisher Information matrix is defined as

$$F(\beta, i_{\text{app}}) = \left(\frac{\partial V}{\partial \beta}(\beta, i_{\text{app}})\right)^{\top} \left(\frac{\partial V}{\partial \beta}(\beta, i_{\text{app}})\right), \qquad (1)$$

where β is a vector of the model parameters, i_{app} is the current profile of the provided experiment, and V is the model output. So, the Fisher Information matrix contains information on how much the model parameters influence the output. Since the Fisher Information is dependent on the provided experiment i_{app}, we can optimise over different experiments to maximise the information on the model parameters.

The optimisation is done by generating a number of experiments with different hyper-parameters based on the class of signals defined in [3], calculating the Fisher Information of the experiments, and choosing the optimal experiment based on the A-, D-, and E-optimality criteria. The complete method and results are described in [4].

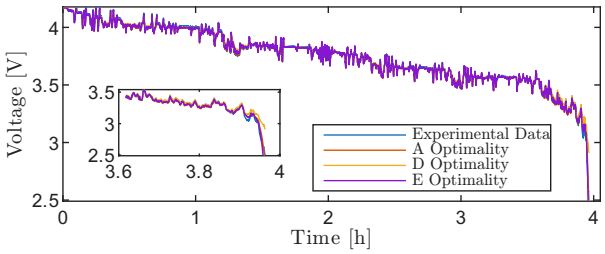

Figure 1: Validation output using parameter sets found with the A-, D-, and E-optimal experiments.

3 Results

Figure 1 displays the voltage outputs of the DFN model using three different sets of parameters. The parameters sets are found using the parameter estimation procedure with the optimal experiments related to the A-, D- and E-optimality. The validation in Figure 1 is executed on a validation set, different from the estimation sets. The RMSEs of the validations are 11.6 mV, 25 mV, and 9.1 mV for the A-, D-, and E-optimal experiments, respectively, confirming that more informative experiments lead to accurate results even in the lower discharge range.

References

[1] Doyle, Fuller, and Newman, "Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell," *JES*, 1993.

[2] Khalik, Donkers, Sturm, and Bergveld, "Parameter Estimation of the Doyle-Fuller-Newman Model for Lithium-Ion Batteries by Parameter Normalization, Grouping, and Sensitivity Analysis," *JPS*, 2021.

[3] Sheikh, Bergveld, and Donkers, "A comprehensive approach to sparse identification of linear parameter-varying models for lithium-ion batteries using improved experimental design," *JES*, 2024.

[4] Weldeghebreal, le Roux, and Donkers, "Informative Experiment Design for Parameter Estimation of the Doyle-Fuller-Newman Battery Model (submitted)," *JPS*, 2026.

Robust dexterous grasping through virtual hand-object interconnection

Julien Vanderheyden[1], Guillaume Drion[1], Pierre Sacré[1,*], Fulvio Forni[2,*]
[1]Dept. of Electrical Engineering and Computer Science, University of Liège, Belgium
[2]Department of Engineering, University of Cambridge, United Kingdom
{julien.vanderheyden,gdrion,p.sacre}@uliege.be, f.forni@eng.cam.ac.uk

Introduction

Dexterous grasping in unstructured environment remains an open problem. In such settings, the central challenge is to achieve a general methodology for grasping across diverse objects while maintaining robustness to high levels of perceptual noise. Relying on detailed object knowledge reduces flexibility and limits applications, as sensor-based reconstruction is often unreliable due to noise, occlusions, and incomplete observations. A promising approach to avoid these limitations is to abstract the object representation, capturing only its stable and invariant aspects while absorbing perceptual uncertainty. We hypothesize that representing objects using simple shape primitives [1] (spheres, cylinders, or boxes) embeds sufficient information to enable adaptive grasping, and we propose a grasping policy built upon this representation.

Approach

Our idea for achieving adaptive grasping under such coarse modeling is to frame the grasp action as a dynamic interconnection between the robot and the object, following the philosophy of Virtual Model Control (VMC) [2, 3]. Conceptually, virtual mechanical components such as springs and dampers are structured in a suitable geometry related to the shape primitive of the object and virtually attached to the robot and to the real object. The controller is then implemented to emulate the forces these elements would exert. To enable the motion to be adaptive to any object dimension, the springs attach points are not fixed but rather free to slide all along the object surface, as illustrated in Fig. 1. The key idea is that the grasp motion should emerge from this coupled dynamics rather than being imposed through explicit motion commands.

Experiments

Experimental validation is performed using a Shadow Dexterous Hand (see Fig. 2). Robustness is evaluated through grasping experiments on 3D-printed shape primitives of different sizes, under progressively increasing levels of controlled noise. This noise affects both the estimated object dimensions and its pose, including position and orientation. Generalization performance is then assessed by applying the same policy to a broad set of everyday objects. The results show that the proposed approach exhibits strong robustness to uncertainty while generalizing effectively beyond the primitive shapes.

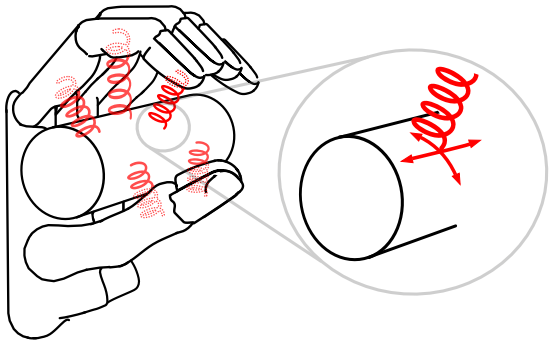

Figure 1: Contact is established by interconnecting the hand and the object while adaptive grasping is achieved by allowing contact points to slide along the object surface



Figure 2: Experimental setup: an anthropomorphic ShadowHand robot grasping various 3D-printed shape primitives

Acknowledgment This work was supported by the Belgian Government through the Federal Public Service Policy and Support. *PS and FF contributed equally to this work as senior authors.

References

[1] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic grasp planning using shape primitives", in *2003 IEEE International Conference on Robotics and Automation (ICRA)*, Sep. 2003, pp. 1824–1829.

[2] J. E. Pratt, "Virtual model control of a biped walking robot", Master thesis, Massachusetts Institute of Technology, 1995.

[3] D. Larby and F. Forni, "Optimal virtual model control for robotics: Design and tuning of passivity-based controllers", *IEEE Transactions on Robotics*, vol. 42, pp. 439–454, 2026.

Energy-based learning in resistor-capacitor circuits

Marieke Heidema, Bart Besselink, Henk van Waarde
Bernoulli Institute, RUG, the Netherlands
h.m.heidema@rug.nl b.besselink@rug.nl h.j.van.waarde@rug.nl

1 Introduction

Analog electronics have gained interest in recent years as they have the potential to perform computations that are substantially faster and more energy-efficient when compared to conventional digital computers. **Energy-based learning** algorithms, such as contrastive learning, can be used to train analog computing systems using local learning rules [1]. Previous research includes analysis of the contrastive learning approach for a resistor network with adjustable resistors [2]. In this study [3], a contrastive learning algorithm is developed for a network of adjustable linear resistors, ideal capacitors, and current sources. Its convergence is proven for any step-size belonging to a specified set.

2 Resistor-capacitor circuit

Consider a resistor-capacitor circuit whose structure is captured by a graph $\mathscr{G}$ with N nodes and $B = n_R + n_C + n_S$ branches. The first $N-1$ nodes are connected through n_R **adjustable resistors**, while node N is the ground. Each of the first $N-1$ nodes are also connected to a **grounded capacitor**, accounting for an additional $n_C = N-1$ branches. Finally, n_S branches correspond to current sources, each connecting two nodes from the set $\{1,2,\dots,N-1\}$. Figure 1 shows an example of such a circuit.

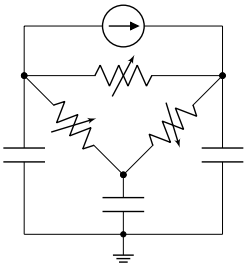

Figure 1: An example of the considered type of circuits.

Let $i_R \in \mathbb{R}^{n_R}$, $i_C \in \mathbb{R}^{n_C}$, and $i_S \in \mathbb{R}^{n_S}$ denote the vectors of currents through the resistor, capacitor, and current source branches, respectively. Moreover, let the vectors of voltages over the resistor, capacitor, and current source branches, respectively, be denoted by $v_R \in \mathbb{R}^{n_R}$, $v_C \in \mathbb{R}^{n_C}$, and $v_S \in \mathbb{R}^{n_S}$. We will assume that the voltages v_R are the measured outputs of the model. Using Ohm's law and Kirchhoff's laws then leads to the dynamics

$$C\tfrac{d}{dt}v_C = -D_R G D_R^\top v_C + D_S i_S, \;\; v_R = D_R^\top v_C, \qquad (1)$$

with state $v_C \in \mathbb{R}^{n_C}$, input $i_S \in \mathbb{R}^{n_S}$, and output $v_R \in \mathbb{R}^{n_R}$. Here, $G > 0$ is a diagonal matrix given by $G = \mathrm{diag}(g)$, with $g \in \mathbb{R}^{n_R}$ the vector of positive conductance values, and $C > 0$ is the diagonal capacitance matrix with positive capacitance values on its diagonal. Moreover, the incidence matrices $D_R \in \mathbb{R}^{n_C \times n_R}$ and $D_S \in \mathbb{R}^{n_C \times n_S}$ capture the orientation of the branches corresponding to resistors and current sources, respectively, and which nodes they connect.

3 Problem statement

For a given vector of conductances $g \in \mathbb{R}^{n_R}$ and constant input $i_S \in \mathbb{R}^{n_S}$, we define the *steady-state output* of (1) as

$$v_R^*(g, i_S) := \lim_{t\to\infty} v_R(t),$$

where $v_R(t)$ is the output of (1) corresponding to the constant input i_S. Now, we want to find a vector of conductances g so that the pair $(i_S, v_R^*(g, i_S))$ matches some given input-output data. In particular, we solve the following problem.

Problem 1 *Given input-output data (i_S^d, v_R^d), find a sequence of positive conductances $(g^k)_{k=0}^\infty$, where each g_j^{k+1} is determined only based on local information known to that resistor, namely, g_j^k and $v_{R,j}^*(g^k, i_S^d)$, such that $\lim_{k\to\infty} g^k = g^*$ for some positive g^* satisfying $v_R^*(g^*, i_S^d) = v_R^d$.*

In Problem 1, we ask for a **local updating rule** since this avoids energy loss and delay caused by transferring data between the resistors and a centralized processor. The problem is solved as follows. We observe that the dynamics of the network are governed by the **"energy" function**

$$E(g, i_S, v_R) = \tfrac{1}{2} v_R^\top G v_R - v_R^\top M i_S \qquad (2)$$

in the sense that, given g and i_S, $v_R^*(g, i_S)$ corresponds to the minimum of the energy function. Based on this energy function (2), a cost function is defined as the difference between two energy states: a free state and a state associated with the training data. The cost function is proven to be convex and its gradient Lipschitz continuous, guaranteeing convergence of an iterative algorithm minimizing the cost function for any step-size belonging to a specified set.

References

[1] B. Scellier, M. Ernoult, J. Kendall, and S. Kumar, "Energy-based learning algorithms for analog computing: a comparative study," *Advances in neural information processing systems*, vol. 36, pp. 52 705–52 731, 2023.

[2] A.-M. Huijzer, T. Chaffey, B. Besselink, and H. J. van Waarde, "Convergence of energy-based learning in linear resistive networks," arXiv preprint arXiv:2503.00349, 2025.

[3] M. Heidema, B. Besselink, and H. van Waarde, "Energy-based learning in resistor-capacitor circuits," manuscript submitted, 2025.

On duality of memristive elements

Brayan M. Shali brayan.shali@kuleuven.be
Rodolphe Sepulchre rodolphe.sepulchre@kuleuven.be

Department of Electrical Engineering (ESAT), KU Leuven, Belgium

1 Introduction

Duality has long played a central role in electrical circuit theory. By interchanging voltage and current, series and parallel interconnections, and resistance and conductance, duality reveals that many circuit laws and synthesis results arise from structural principles rather than from specific component models. It enables results obtained in one formulation to be transferred systematically to its dual and highlights the geometric structure underlying circuit behavior.

While duality has been extensively studied for linear and nonlinear resistive, capacitive, and inductive elements, its systematic treatment in the context of memristive elements remains limited. We aim to address this by exploring the duality properties of memristive elements and examining how memory, dissipation, and constitutive relations transform under voltage-current interchange. As a starting point, we revisit duality for resistive elements, where these ideas admit a clear and well-established formulation.

2 Duality of resistive elements

Duality for resistive elements reflects the basic symmetry between voltage- and current-based descriptions of dissipation [1]. A resistor may be characterized either by specifying the voltage produced in response to an applied current or, equivalently, by specifying the current drawn in response to an applied voltage. These two descriptions correspond to exchanging resistance and conductance and are what we refer to as the duality of resistive elements.

To make this more precise, consider the linear resistor

$$v = Ri \tag{1}$$

with resistance $R > 0$. This element may equivalently be characterized as a linear conductor

$$i = Gv \tag{2}$$

with conductance $G = R^{-1} > 0$. The duality between these descriptions can be understood geometrically by associating scalar potentials with the constitutive relations. In particular, for the linear resistor (1), we define the *content* as

$$\mathcal{C}(I) := \int_0^I v\,di = \tfrac{1}{2}RI^2, \tag{3}$$

which represents the accumulated dissipative effect as a function of the current. Dually, the linear conductor (2) admits the *co-content*

$$\mathcal{C}^*(V) := \int_0^V i\,dv = \tfrac{1}{2}GV^2. \tag{4}$$

The functions $\mathcal{C}$ and $\mathcal{C}^*$ form a convex conjugate pair in the sense of Legendre–Fenchel duality. Moreover, they satisfy $\nabla\mathcal{C}(i) = v$ and $\nabla\mathcal{C}^*(v) = i$, that is, the resistor relation (1) is generated by the potential $\mathcal{C}$, while the dual conductor relation (2) is generated by the dual potential $\mathcal{C}^*$. Duality between resistance and conductance is therefore reflected not only at the level of constitutive maps, but also at the level of the associated scalar potentials. The same viewpoint extends to nonlinear resistors and conductors by replacing the linear constitutive map with a monotone nonlinear relation, for which the content and co-content remain well-defined and continue to form a convex conjugate pair.

3 Duality of memristive elements

A memristive element may be viewed as a natural generalization of a resistor in which the instantaneous voltage-current relation depends on the history of the current. Accordingly, a memristor is described not by a static constitutive relation but by a relation between voltage and current *trajectories*. From this perspective, a memristor may be modeled as a causal operator acting on suitable signal spaces, where causality reflects the physical requirement that the voltage at a given time depends only on past and present values of the current. To retain consistency with dissipative behavior, it is natural to require this operator to be monotone.

The presence of memory generally breaks symmetry and, thus, prevents the existence of a scalar content generating the full relation. This raises the question: can a meaningful notion of content still be recovered? A related issue concerns the definition of a dual description: under what conditions can such a description be regarded as a conjugate or inverse while remaining compatible with causality, monotonicity, and stability? It is worth exploring whether the Riemannian gradient formulation of memristors in [2] can help address these questions and recover a content-based notion of duality.

References

[1] D. Jeltsema and J. M. Scherpen, "Multidomain modeling of nonlinear networks and systems," *IEEE Control Systems Magazine*, vol. 29, no. 4, pp. 28–59, 2009.

[2] F. Forni and R. Sepulchre, "Gradient modelling of memristive systems," *arXiv e-prints*, p. arXiv:2504.10093, 2025.

A hybrid systems approach for rhythmic neuromorphic control of a pendulum

Elena Petri[1], Romain Postoyan[2], W. P. M. H. (Maurice) Heemels[1]

[1]Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands.

[2]Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France.

Email: e.petri@tue.nl

1 Introduction

Motivated by the potential advantages of brain-inspired technologies over classical clock-based technologies, such as energy efficiency, low-latency, adaptability, and robustness, neuromorphic engineering is an emerging research domain. In the field of systems and control, an important gap in the neuromorphic control literature is the lack of systematic modeling, analysis and design tools for spiking systems. Because biological neurons operate in an event-based manner, where information is communicated via the occurrence or absence of spikes, we believe that hybrid dynamical systems are a well-suited and promising theoretical tool for the design, modeling and analysis of spiking-based controllers. To show this, we use the hybrid dynamical systems framework in [1] to model and formally analyze a rhythmic neuromorphic control problem of a pendulum system, recently proposed as a prototypical setup in [2], and also studied in [3]. The control goal in [2] is to induce desirable pendulum oscillations using a neuromorphic controller generating bursts of spikes whenever the pendulum's angular position crosses its resting position. The length of these bursts of spikes directly influences the amplitude of the oscillations.

2 Goal and contributions

Our objective is to provide a formal analysis of the results presented in [2] using hybrid systems tools. In particular, we model the activation of the burst of spikes as a single spike represented by a Dirac delta pulse, whose amplitude I is related to the duration of the burst. The overall closed-loop pendulum system is then modeled as a hybrid dynamical system, where spikes correspond to jumps of amplitude I on the angular velocity, and flows correspond to the continuous-time open-loop (linearized) pendulum dynamics between spikes. To robustify the detection of the angular position crossing its resting position, which corresponds to the jump condition in the hybrid model, an auxiliary logic variable is introduced.

Exploiting the hybrid model, we formally prove that, given any spike amplitude $I > 0$, there exists a unique hybrid limit cycle for the closed-loop system describing the pendulum oscillations. We also show that this hybrid limit cycle exhibits a uniform exponential stability property. These properties are proved by explicitly using the expression of the solution; we therefore exploit the particular structure of the considered problem. Simulations (see Figure 1) illustrate our results. Since for any given spike amplitude $I > 0$ there exists a unique limit cycle describing the pendulum oscillations, and the solution exponentially converges to it, we can adapt the pulse amplitude to generate a hybrid limit cycle of a desired maximum amplitude. This is done by comparing the maximal angular position, obtained when the angular velocity crosses zero, with the desired one. Based on the sign of this error, a spiking adaptation mechanism adds or subtracts a fixed quantity from the spike amplitude I. We propose a new hybrid model in which jumps correspond either to the spiking control action, when the pendulum crosses its resting position, or to the spiking adaptation mechanism, when the angular velocity changes sign. Simulations illustrate the effectiveness of the adaptation mechanism for the convergence to the desired limit cycle. A formal analysis of this new hybrid system is left for future work.

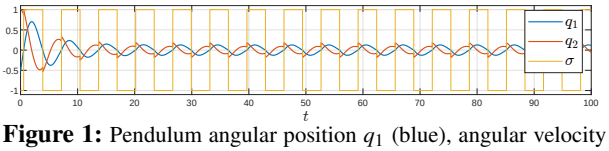

Figure 1: Pendulum angular position q_1 (blue), angular velocity q_2 (red) and auxiliary logic variable σ (yellow).

This abstract is based on the conference paper [4].

References

[1] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. New Jersey, USA, Princeton University Press, 2012.

[2] R. Schmetterling, F. Forni, A. Franci, and R. Sepulchre, "Neuromorphic control of a pendulum," *IEEE Control Systems Letters*, vol. 8, pp. 1235–1240, 2024.

[3] T. Medvedeva, A. Franci, and F. Castaños, "Formalizing neuromorphic control systems: A general proposal and a rhythmic case study," *IEEE Conference on Decision and Control,* Rio de Janeiro, Brazil, 2025.

[4] E. Petri, R. Postoyan, and W. P. M. H. Heemels, "Rhythmic neuromorphic control of a pendulum: A hybrid systems analysis," *IEEE Conference on Decision and Control,* Rio de Janeiro, Brazil, 2025.

Regulating Rhythmic Compliance in Central Pattern Generators

Robin Wroblowski[1], Rodolphe Sepulchre[1,2]

[1]KU Leuven, Department of Electrical Engineering (ESAT), KasteelPark Arenberg 10, 3001 Leuven, Belgium

[2]Department of Engineering, University of Cambridge, TrumpingtonStreet, Cambridge CB2 1PZ, United Kingdom.

Email: *robin.wroblowski@kuleuven.be*, *rodolphe.sepulchre@kuleuven.be*

1 Rhythmic Event-based Interaction

In the context of robotic locomotion, the adoption of central pattern generators has proven to be a useful bio-inspired principle for generating rhythmic actuator commands. Despite this, the process by which such CPGs incorporate sensory feedback from the environment remains a challenging issue in both biology and robotics. Conventionally, this issue has been addressed through the relative contribution of feedforward and feedback mechanisms within the control system, which is modulated using the sensory feedback gain [1].

However, the limited number of analyses of this balance are restricted to systems of coupled mathematical oscillators, in which the rhythmic controller interacts continuously with the environment. This differs from biological reality, in which information is transmitted through the propagation of *events* generated by spiking neurons. Such event-based signals enable robust signaling and naturally provide energy-efficiency through temporal sparsity. Furthermore, regulating events rather than trajectories overcomes the need for calibration[2] and enables decentralized design and control of network properties. This work explores how to regulate the feedforward-feedback balance in an event-based setting by locally tuning the receptor neurons' sensitivity to control the CPG network's *rhythmic compliance*.

2 Neuronal Tuning of Input Compliance

We investigate how a rebound-bursting CPG controller interacts with a rhythmic plant by studying pendulum control using a half-center oscillator (HCO), following [3].

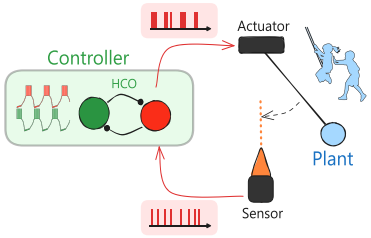

Figure 1: A rhythmic controller (HCO) interacts with a plant (pendulum) purely through event-based signals.

For sufficiently strong coupling, the controller and plant achieve synchronization provided their intrinsic periods are sufficiently similar. The aim is to adjust the sensitivity of the receptor neuron receiving plant feedback to elicit distinct interaction behavior under identical feedback conditions. The fundamental mechanism through which input sensitivity is modulated draws from prior research indicating a direct link between the robustness of network oscillations and the cellular mechanism of rebound excitability [4]. Thus, a cellular-level tuning mechanism provides a decentralized approach for influencing global network behavior.

As a proof of concept, Figure 2 illustrates the distinct behaviors of compliant and stiff rhythms, arising from the absence or presence of a slow positive feedback mechanism. Initially, the two controllers and the plant share the same intrinsic period (T_0), leading to synchronization. A period mismatch ($T_1 - T_0$) is then introduced by adjusting the intrinsic frequency of the pendulum. The compliant controller, being sensitive to sensory feedback, becomes entrained by the plant, whereas the stiff controller remains *stubborn*, imposing its own intrinsic frequency on the plant.

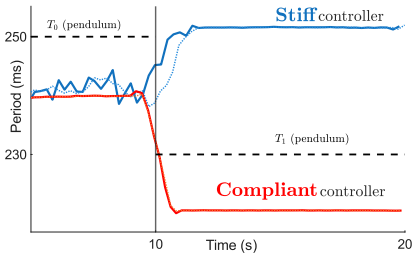

Figure 2: Two controllers show distinct interaction behavior in the case of period mismatch.

Analyzing the input-output mapping of receptor neurons may offer insights for developing a decentralized framework for designing rhythmic, event-based interactions. A systematic study of the input-output properties of event-based systems, as suggested in [5] and [6], could provide a powerful means to understand both the intrinsic and extrinsic dynamics of event-driven and excitable systems.

References

[1] A. Kuo "The relative roles of feedforward and feedback in the control of rhythmic movements, *Motor Control*, 2002

[2] R. Sepulchre, A. Cecconi, M. Bin, L. Marconi "Regulation without calibration", *Arxiv*, 2025

[3] R. Schmetterling, F. Forni, A. Franci, R. Sepulchre, "Neuromorphic Control of a Pendulum", *IEEE Control Syst Lett*, 2024

[4] J. Dethier, G. Drion, A. Franci, R. Sepulchre, "A positive feedback at the cellular level promotes robustness and modulation at the circuit level", *J Neurophysiol.*, 2015

[5] R. Wroblowski & R. Sepulchre "Describing Functions and Phase Response Curves of Excitable Systems", *Arxiv*, 2025

[6] A. Cecconi, M. Bin, L. Marconi, R. Sepulchre, "On the Contraction of Excitable Systems", *Arxiv*, 2025

Acknowledgment This work was supported by European Research Council under the Advanced ERC Grant Agreement SpikyControl 101054323.

Input–Output Modeling of Neuromorphic Systems

Alessandro Cecconi
KU Leuven, Belgium
University of Bologna, Italy
alessandro.cecconi@kuleuven.be

Rodolphe Sepulchre
KU Leuven, Belgium
Cambridge, UK
rodolphe.sepulchre@kuleuven.be

1 Introduction

Neuromorphic systems mix analog dynamics with event-based communication. When information is carried by discrete events rather than continuous trajectories, it is natural to seek an event-level input–output description: an excitable circuit is then characterized as the discrete mapping from input to output event trains.

Such a description requires *reliability*: for a given stimulus, the induced event train should be reproducible and robust. Reliability is a discrete property, although inherited from contraction of the underlying analog circuit: incremental stability makes the steady-state response well-defined for suitable inputs and, through the event readout, induces a well-defined steady-state event map [1]. This twofold representation, a continuous physical model together with its induced event-level map, supports a simplified steady-state analysis, in the spirit of the classical frequency-domain approach: under periodic inputs, the input–output steady-state response is described by a phase map, and interconnections reduce to fixed points of the composition.

2 Problem formulation

2.1 Continuous-time physical model

We consider a conductance-based circuit initialized at rest. The capacitor current balance is

$$C\dot{v}(t) = -i_\ell(t) - \sum_k i_k(t) - i_{\text{ext}}(t). \tag{1}$$

All synaptic currents are outputs of memristive elements

$$\tau_j(v_{\text{in}}(t))\dot{x}_j(t) = -x_j(t) + \mu_j(v_{\text{in}}(t)), \tag{2}$$

$$i_k(t) = \bar{g}_k\, \varphi_k(x(t))\big(v_{\text{out}}(t) - E_k\big), \tag{3}$$

where $x = \text{col}(x_1,\dots,x_m) \in [0,1]^m$ denotes the gating variables, $\tau_j(\cdot) > 0$ and $\mu_j(\cdot) \in (0,1)$ are the time constant and steady-state activation, and $\varphi_k : [0,1]^m \to [0,1]$ is an activation function. The signals v_{in} and v_{out} denote the driving and output voltages, respectively. The leak current $i_\ell(t)$ is the static case with $\varphi_\ell(x) \equiv 1$ and $\bar{g}_\ell > 0$. The constants $E_k \in \mathbb{R}$ and $\bar{g}_k > 0$ are reversal potentials and maximal conductances.

2.2 Event input–output model

An event readout $\mathscr{D}$ decodes spike times from the voltage signal at a prescribed scale, as depicted in Figure 1. It defines a mapping from voltage traces to locally finite sets of

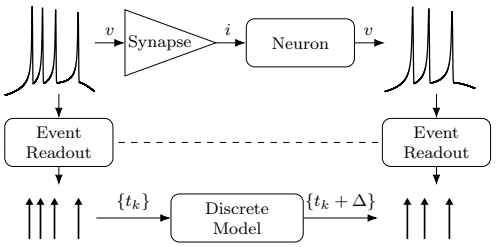

Figure 1: Event readout $\mathscr{D}$ extracts discrete events from continuous voltage.

times. An input event train $\mathscr{I} \in \mathbb{E}$ is encoded as an external voltage $v_{\text{out}}(\cdot;\mathscr{I})$ driving the synapse, which determines $v(\cdot;\mathscr{I})$ through (1)–(3). The induced event-level map is

$$\mathscr{F} : \mathbb{E} \to \mathbb{E}, \qquad \mathscr{F}(\mathscr{I}) = \mathscr{D}\big[v(\cdot;\mathscr{I})\big]. \tag{4}$$

In a contractive excitable regime, the input $\mathscr{I}$ uniquely determines the steady-state voltage trace, hence the event output. If $\mathscr{D}$ is continuous on the set of admissible traces generated by (1)–(3), then $\mathscr{F}$ is a stable event-level input–output description.

3 Phase maps and stability of interconnections

For periodic forcing $\mathscr{I}_T = \{kT\}_{k\in\mathbb{N}}$ in a 1:1 entrainment regime, the steady-state response is described by a phase-delay map between input and output events

$$\mathscr{F}(\mathscr{I}_T) = \{kT + \Delta(T)\}_{k\in\mathbb{N}}, \quad \Delta(T) \in (0,T).$$

For two systems in feedback, a T-periodic solution satisfies

$$T = \Delta_1(T) + \Delta_2(T),$$

and assuming differentiability of Δ_i, local stability at $T^\star$ follows from

$$\big|\Delta_1'(T^\star) + \Delta_2'(T^\star)\big| < 1.$$

Simulations on synaptically coupled Hodgkin–Huxley neurons show that with inhibitory coupling the approach predicts stable rhythms, whereas for excitatory coupling it typically does not.

References

[1] A. Cecconi, M. Bin, L. Marconi, and R. Sepulchre, "On the Contraction of Excitable Systems," *arXiv preprint* arXiv:2511.15477, 2025, submitted to ECC 2026.

Tempered Bayes Filtering for Improved Model-Based Estimation

M.J.T.C. van Zutphen
CST Group, TU/e
m.j.t.c.v.zutphen@tue.nl

D. Herceg
CST Group, TU/e
d.herceg@tue.nl

G. Delimpaltadakis
CST Group, TU/e
i.delimpaltadakis@tue.nl

D.J. Antunes
CST Group, TU/e
d.antunes@tue.nl

1 Introduction

Classical Bayesian estimators such as the Bayes and MAP filters rely heavily on the accuracy of the underlying system model; even moderate model mismatch can significantly degrade performance. Recent work in Bayesian inference has shown that *tempering* likelihoods or posteriors can improve predictive accuracy when models are imperfect. We extend these ideas to filtering and introduce the *tempered Bayes filter*, a three-parameter family that includes the Bayes and MAP filters as special cases and preserves recursive structure. Figure 1 provides the results of a numerical example presented in [1] showing the advantages of tempering.

2 Problem Formulation

Consider a hidden Markov model (HMM) defined by

$$p(x_{k+1}|x_k) := \text{Prob}(x_{k+1}|x_k), \quad p(y_k|x_k) := \text{Prob}(y_k|x_k),$$

with initial state distribution $p_0(x_0)$. Let $\hat{p}$ be an imperfectly identified approximation of the above model. The classical Bayes filter recursively computes the Bayes belief

$$b(x_k|y_{0:k}) \propto p(y_k|x_k)\int p(x_k|x_{k-1})b(x_{k-1}|y_{0:k-1})\,\mathrm{d}x_{k-1}, \quad (1)$$

while the MAP filter instead generates MAP belief vector

$$b_{\text{MAP}}(x_k|y_{0:k}) \propto p(y_k|x_k)\max_{x_{k-1}\in X} p(x_k|x_{k-1})b_{\text{MAP}}(x_{k-1}|y_{0:k-1}),$$

corresponding to the most-likely trajectory ending at x_k.

We temper the Bayesian-update step posterior by introducing likelihood-, and posterior tempering parameters, as

$$p_\lambda(x_{0:k}|y_{0:k}) \propto \left(p(y_{0:k}|x_{0:k})^{\lambda_L}p(x_{0:k})\right)^{\lambda_P}, \quad (2)$$

then add belief-tempering parameter λ_B to yield parameterization $\lambda = (\lambda_L, \lambda_P, \lambda_B)^\top$, defining the tempered belief as

$$b_\lambda(x_k|y_{0:k}) \propto \left(\int p_\lambda(x_{0:k}|y_{0:k})\,\mathrm{d}x_{0:k-1}\right)^{\lambda_B}. \quad (3)$$

We note that $\lambda = (1,1,1)$ recovers the classical Bayes filter.

3 Main Results

Substituting (2)–(3) into (1) yields a recursive implementation of the tempered Bayes filter

$$b_\lambda(x_k|y_{0:k}) \propto \left[p(y_k|x_k)^{\lambda_L\lambda_P}\int p(x_k|x_{k-1})^{\lambda_P} b_\lambda(x_{k-1}|y_{0:k-1})^{1/\lambda_B}\,\mathrm{d}x_{k-1}\right]^{\lambda_B},$$

computationally as complex as the original Bayes filter.

We show that (2) is an optimizer of the ELBO-objective

$$p_\lambda = \arg\max_q \Big\{ \mathbb{E}_q[\log p(y_{0:k}|x_{0:k})] - \frac{1}{\lambda_L}\text{KL}(q\|p) + \frac{1}{\lambda_L}\left(\frac{1}{\lambda_P}-1\right)H(q)\Big\},$$

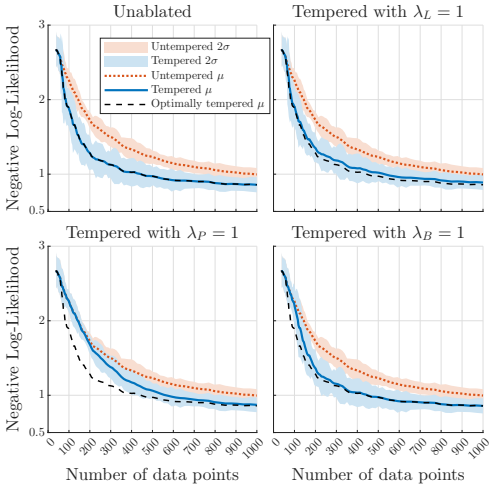

Figure 1: Filtering performance vs. the amount of available data; unablated (standard) and ablated conditions.

providing an interpretation of tempering: λ_L reweights data vs. prior, while λ_P directly regulates posterior entropy.

Furthermore, for $\lambda_P = 1/\lambda_B$ and $\lambda_L = 1$, taking the limit $\lambda_P \to \infty$ recovers the MAP filter, i.e.,

$$\lim_{\lambda_P = 1/\lambda_B \to \infty} b_\lambda(x_k|y_{0:k}) = b_{\text{MAP}}(x_k|y_{0:k}).$$

Thus, part of the framework parameter space can be interpreted as interpolating between Bayes and MAP filtering.

A first-order analysis of the expected negative log-likelihood (NLL) score, a metric minimalized only by a perfect filter

$$N_k = \mathbb{E}_{x_{0:k},y_{0:k}}[-\log b_\lambda(x_k|y_{0:k})],$$

shows that under model mismatch $\hat{p} \neq p$, the gradient $\nabla_\lambda N_k$ at $\lambda = (1,1,1)$ is non-zero, implying that tempering improves performance over the untempered Bayes filter.

Finally, for linear-Gaussian models, "tempered Kalman filter" recursions are derived, where tempering parameters modulate prior/likelihood covariance and belief entropy.

4 Conclusion and Future Work

We present a recursive, entropy-tunable Bayes generalization improving data-based model filtering performance. Future work: adaptive tuning and integration with control.

References

[1] M. J. T. C. van Zutphen, D. Herceg, G. Delimpaltadakis, and D. Antunes, "Tempering the Bayes Filter Towards Improved Model-Based Estimation," arXiv:2512.02823, 2025.

Accelerating Diffusion Models for Adaptive Motion Planning in Dynamic Environments

Matthias Van Eysendeyk[1,2] Edward Sandra[1,2] Dries Dirckx[1,2] Lander Vanroye[1,2,3]
Flavia Sofia Acerbo[1,2] Jan Swevers[1,2] Wilm Decré[1,2]

[1]MECO Research Team, Department of Mechanical Engineering, KU Leuven, Leuven, Belgium
[2]Flanders Make@KU Leuven, Leuven, Belgium
[3]Intermodalics, 3000 Leuven, Belgium
e-mail: matthias.vaneysendeyk@kuleuven.be

1 Introduction

Diffusion models have recently emerged as a powerful deep learning approach for robot motion planning, capable of generating multi-modal, high-quality trajectories in complex environments. However, their slow sampling speed poses a significant challenge for real-time applications, particularly in dynamic settings that require replanning at high frequencies. This challenge is addressed by implementing a knowledge distillation approach. Leveraging the resulting speedups, replanning can be executed up to 100 Hz, utilizing a cost-based trajectory selection mechanism that balances goal proximity, predicted collision risk, and trajectory smoothness.

2 Methodology

Diffusion models are inherently slow due to there requirement of many sequential model evaluations in order to completely denoise a noisy sample. To decrease this number of sequential denoising steps, a knowledge distillation approach is implemented consisting of the following steps:

1. train the original diffusion model, which will be considered as the *teacher* [2],
2. initialize a second, *student* model with the weights of the teacher model,
3. teach the student model to mimic the teacher, while combining multiple denoising steps of the teacher into one denoising step of the student model.

Using this approach, sampling multi-modal trajectories becomes possible in as few as 10 ms. This allows diffusion models to be used in dynamic environments, by sampling a new batch of trajectories every replanning step. The only choice left to make, is which trajectory π from the batch $\mathbf{B}$ to follow in each replanning step. For this, a cost function is designed consisting of three terms. The first term $\mathbf{C}_{\text{dur}}(\pi)$ takes into account the duration of the planned trajectory, promoting faster trajectories. The second term $\mathbf{C}_{\text{coll}}(\pi)$ predicts the future locations of the obstacles based on past observations and applies a penalty when the robot is predicted to collide with one of the obstacles. The final term $\mathbf{C}_{\text{smooth}}(\pi)$ reduces the volatility of the replanning algorithm, by applying a small cost when the new trajectory is vastly different from the previous one, preventing the robot from making too many unnecessary mode-switches throughout its motion. With this cost function, the replanning algorithm can utilize the multi-modal trajectory generation capabilities of the diffusion model, while avoiding any moving obstacles.

$$\mathbf{C}(\pi) = \mathbf{C}_{\text{dur}}(\pi) + w_{\text{c}}\,\mathbf{C}_{\text{coll}}(\pi) + w_{\text{s}}\,\mathbf{C}_{\text{smooth}}(\pi), \quad (1)$$

$$\pi^{\star} = \arg\min_{\pi \in \mathbf{B}} \mathbf{C}(\pi) \quad (2)$$

3 Results

The approach is tested on a 7-DoF Franka manipulator, moving from start to goal state in an environment containing several randomly moving spherical obstacles. Figure 1 shows an example, where the robot is able to adapt to the moving obstacles by changing its trajectory when necesary.

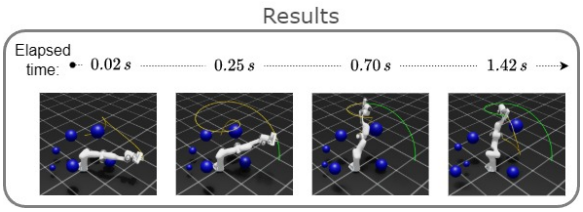

Figure 1: Results from trajectory generation using accelerated diffusion models, on an unseen environment with random, dynamic obstacles.

References

[1] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," 2020.

[2] Sandra, E., Vanroye, L., Dirckx, D., Cartuyvels, R., Swevers, J., and Decré, W. (2025). Accelerated multi-modal motion planning using context-conditioned diffusion models. arXiv preprint arXiv:2510.14615.

Acknowledgment This work was supported by Flanders Make through SBO project LearnOpTra (Learning meets optimisation for robust and multimodal trajectory planning and control).

MPC-Informed Residual Reinforcement Learning for Hybrid Control

Senne Bogaerts
senne.bogaerts@kuleuven.com

Flavia Sofia Acerbo
flaviasofia.acerbo@kuleuven.be

Patrick Scheffe
patrick.scheffe@kuleuven.be

Jan Swevers
jan.swevers@kuleuven.be

Wilm Decré
wilm.decre@kuleuven.be

MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium

1 Introduction

Modern robotics applications demand real-time control strategies that can handle complex environments while maintaining hard safety requirements. Model Predictive Control (MPC) offers a solution as it provides safety guarantees through explicit constraint handling and can handle complex environements. However, it is computationally demanding for real-time deployment, may struggle under substantial model uncertainty, and cannot directly process high-dimensional sensor data. Reinforcement Learning (RL) addresses these limitations by requiring only computationally lightweight neural network forward passes at execution time, enabling real-time deployment. It can also leverage domain randomization to improve robustness to model uncertainty and can naturally accommodate diverse input modalities. Unfortunately, standard RL methods cannot provide hard safety guarantees and suffer from poor data efficiency. The reward design process further complicates training: dense rewards can lead to suboptimal policies despite intensive shaping efforts, while sparse rewards worsen data efficiency. To combine the strengths of MPC and RL, this work proposes a hybrid residual structure [1] where MPC provides the nominal action and RL learns a corrective residual. In addition, the RL policy is informed by the MPC's planned actions and predicted trajectories.

2 Methodology

Our hybrid control scheme integrates MPC and RL through a residual architecture, as shown in Figure 1. The MPC and RL policies operate on state spaces $\mathbf{S}_{\text{MPC}}$ and $\mathbf{S}_{\text{RL}}$, which are not required to be identical. At each time step, the MPC generates a nominal action $\mathbf{a_0}$, the planned sequence $\mathbf{a_{0:T-1}}$, and the predicted trajectory $\hat{\mathbf{s}}_{\mathbf{0:T}}$ over its horizon. These outputs are added to the observation space of the RL policy, which outputs a residual $\mathbf{a}_{\text{RL}}$, yielding the final action $\mathbf{a_r} = \mathbf{a_0} + \mathbf{a}_{\text{RL}}$. Fully informing the RL network of the MPC plan is the key difference from standard residual RL. Training the residual policy is time-consuming because the MPC is computationally expensive. Future work will address this.

3 Results

We validate the approach using a simulated Furuta pendulum with an additional mass at the second arm's end (see Figure 1). By randomly sampling the mass and not explicitly

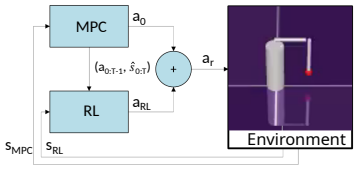

Figure 1: Hybrid control structure

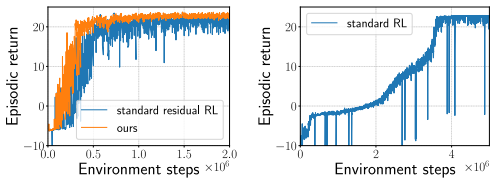

Figure 2: Left: training curves for the proposed method and standard residual RL. Right: training curve for standard RL. All curves use exponential smoothing (factor 0.6).

providing its value to either controller, we introduce uncertainty into the system via domain randomization. The MPC is nonlinear and uses a quadratic objective to swing the pendulum up and stabilize it in the upright position. Figure 2 illustrates the training performance of our extended observation space approach compared to standard residual RL and standard RL, while Table 1 summarizes data efficiency and mean return. The return increases with the time the pendulum remains near the upright position. These results show that our method demonstrates superior data efficiency, lower variance after convergence, and higher mean return.

Table 1: Data efficiency (steps to convergence) and mean return over 100 episodes (identical seeds).

	Ours	Residual RL	Standard RL	MPC
Steps	400k	600k	3.5M	–
Return	22.93	21.62	22.86	14.81

References

[1] T. Johannink et al., "Residual Reinforcement Learning for Robot Control," 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 2019, pp. 6023-6029, doi: 10.1109/ICRA.2019.8794127.

Acknowledgement: This work has been carried out within the framework of Flanders Make's SBO project LearnOpTra (Learning Meets Optimisation for Robust and Multimodal Trajectory Planning and Control).

Can Timeseries Foundational Models be Used for Control?

Lucian Cristian Iacob, Bruno Depraetere, Branimir Mrak
cristian.iacob, bruno.depraetere, branimir.mrak @flandersmake.be
Flanders Make, Belgium

1 Introduction

This work investigates the available state-of-the-art *Foundational Models* (FMs) for timeseries forecasting, specifically, their use within a control loop, to forecast disturbances or references, use as a plant emulator during MIL development of controllers, or, in the future, to even include them explicitly within an MPC solution as a control model. FMs have been the topic of intense development and research in the past few years. However, while large language models and multimodal (audio and video) models have seen incredible leaps in their capabilities in multiple use cases, timeseries forecasting has only recently started to be addressed using FMs. As they are pretrained on massive amounts of data, a possible benefit for engineering applications is that they could provide good predictions, using off-the-shelf models, without extensive fine-tuning and, in an ideal case, even offer good zero-shot predictions.

Industry could thus potentially drastically reduce the need for costly data gathering and downtime by employing off-the-shelf pretrained FMs with minimal finetuning.

2 Methodology and Results

For this present study, we investigated leading FMs such as Chronos [1] and TimesFM [2]. First, we focused on evaluating performance for forecasting. To do so, we generated test signals including increasing amplitude sines, multisines, and exponential signals. We then evaluated the FMs' prediction performance in zero-shot mode, only providing unseen context. The results were observed to generally be on par or close to classical methods like (S)ARIMA and NARX that had been finetuned. Moreover, in real-world periodic data sets (e.g. energy demand), FMs have proven well-suited to capture the periodicity of the data, although with deviations in the amplitude, see Fig. 1.

We then attempted to use FMs to model dynamical systems. However, most FMs at the time of testing did not allow external inputs to be used and only considered one-dimensional signals (or with limited covariate support). Moreover, the finetuning capability may not be available off-the-shelf and complex knowledge from the user may be required to implement it. We found that Chronos 2 did allow for external inputs and, through the AutoGluon package, also offered finetuning using simple API calls. To test Chronos and its ability to model dynamics, we considered linear and nonlinear systems, and a real-world dataset of a complex biochemical process. When the input excited periodic dynamics, the FM could capture well the resulting periodic nature of the output signal. Performance degrades slightly for non-periodic inputs, but can still manage to predict the dynamics to some extent, as shown in Fig. 2 for a linear system using Chronos 2 in zero-shot mode (only unseen context data is given). We note that finetuning on the context data can further improve the result.

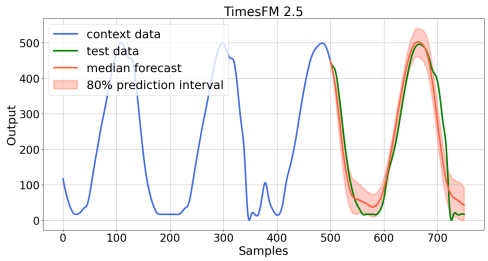

Figure 1: Zero-shot forecasting using TimesFM.

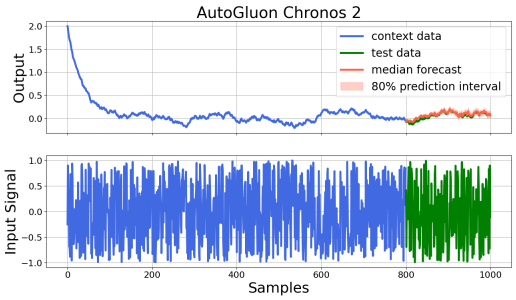

Figure 2: Zero-shot forecasting using Chronos.

Overall, FMs showed promising capabilities for timeseries forecasting in zero-shot mode or with limited finetuning, as well as mixed but encouraging results for modeling of dynamic systems. While several examples performed particularly well, Chronos 2 struggled in more complex scenarios. For now, better performance is expected for forecasting references or predicting disturbances, which are of periodic nature (e.g. periodic demands), however, with proper prior testing, FMs may also be suitable as cheap to train and deploy surrogate models of more complex systems. Do note this is a rapidly evolving field: this work was carried out in 2025 and new versions of FMs are periodically released with improved functionalities.

References

[1] A. F. Ansari, "Chronos-2: From Univariate to Universal Forecasting," arXiv:2510.15821, 2025.

[2] A. Das, "A decoder-only foundation model for timeseries forecasting," arXiv:2310.10688, 2024.

Acknowledgements: This work was carried out in the framework of Flanders Make's CTO_Action 2025-22 - Advanced Control Systems project. It was also supported by the Flemish Government in the framework of the Flanders AI Research Program (https://www.flandersairesearch.be/en) that is financed by EWI (Economie Wetenschap & Innovatie).

Path-Centric Graph Contrastive Representation Learning for Directed Acyclic Graphs

Jorn van Kampen, Mauro Salazar, Theo Hofmans
Control Systems Technology
Eindhoven University of Technology
Email: {j.h.e.v.kampen, m.r.u.salazar, t.hofman}@tue.nl

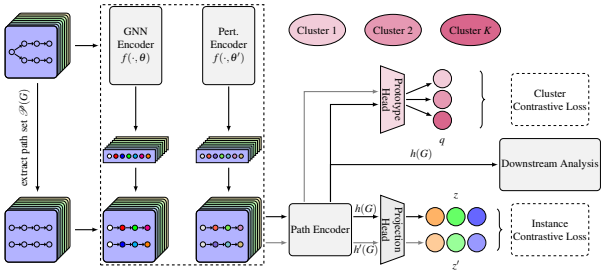

Figure 1: The proposed graph contrastive learning framework of path representations for Directed Acyclic Graphs (DAGs). Paths consisting of node embeddings are used as input to the encoder. Projection and prototype heads are used only during training, downstream tasks operate on path embeddings prior to projection.

Modern engineering systems are increasingly designed through large-scale combinatorial exploration of concept designs [1], producing thousands of architectures that must be analyzed, compared, and clustered to gain insight into the underlying design space. Many such architectures, particularly in sequential process systems, manufacturing pipelines, and energy or material flow systems, are naturally represented as graphs, where system components form nodes and edges encode functional dependencies.

Graph representation learning enables mapping these structured designs into continuous embedding spaces, allowing similarity analysis, clustering, and downstream evaluation without manual (biased) feature extraction. However, most existing unsupervised methods primarily focus on node connectivity. While effective for capturing local structure, these methods are limited in their ability to represent directed flow characteristics, which are often the key discriminator between candidate designs.

In this work, we propose a path-centric contrastive representation learning framework for directed acyclic graphs (DAGs) that explicitly captures flow semantics, as shown in Figure 1. Our approach first learns node embeddings using a graph neural network encoder, from which contrastive views are generated via encoder perturbations [2]. Leveraging the acyclic structure of the graphs, we then enumerate source-to-sink paths and construct ordered sequences of node embeddings along each path. These sequences are then encoded using a recurrent neural network to obtain path-level representations that reflect directed functional behavior of the original system. The resulting path embeddings serve as expressive graph-level representations that complement and extend traditional structure-based encodings. Figure 2 shows the results for a downstream clustering task, whereby we obtain clearly defined and separated clusters that would not have been obtained with traditional graph structure-based representation learning methods.

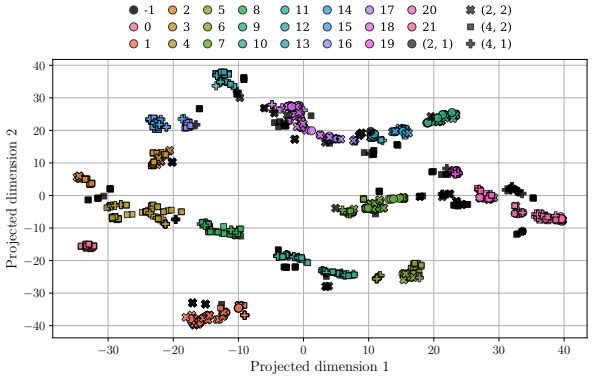

Figure 2: Two-dimensional projection of the learned graph embeddings together with their cluster assignment. Colors indicate the learned cluster, while marker symbols denote the manually assigned cluster based on engineering experience ($N_{\text{sinks}}, N_{\text{sources}}$). The mismatch between learned and heuristic-based clusters reveals structural similarities beyond those captured through engineering experience.

References

[1] J. Wijkniet and T. Hofman, "Modified computational design synthesis using simulation-based evaluation and constraint consistency for vehicle powertrain systems," *IEEE Trans. on Vehicular Tech.* , 2018.

[2] J. Xia, L. Wu, J. Chen, B. Hu, and S. Z. Li, "Simgrace: A simple framework for graph contrastive learning without data augmentation," in *Proc. of the ACM Web Conference*, 2022.

Non-linear Model Predictive Control of Anaerobic Codigestion using Long Short-Term Memory Networks

Andrés Ignacio Pino-Santana
University of Mons / Otto von Guericke University
Systems, Estimation, Control & Optimization (SECO)
Email: AndresIgnacio.PINOSANTANA@umons.ac.be

Achim Kienle
Otto von Guericke Univ. / MPI
Process Synthesis & Dynamics
Email: achim.kienle@ovgu.de

Alain Vande Wouwer
Univ. of Mons
Systems, Estimation, Control & Optimization
Email: alain.vandewouwer@umons.ac.be

1 Introduction

Anaerobic Codigestion is a complex, non-linear biological process that treats multiple substrates for organic waste valorization. However, its advanced control is often constrained by the lack of affordable online sensors for key state variables (e.g., Volatile Fatty Acids, alkalinity).
This work explores a control strategy that relies solely on readily available measurements (inlet flows and methane production), minimizing instrumentation complexity [3]. It proposes a data-driven Non-linear Model Predictive Control (NMPC) architecture integrating a Long Short-Term Memory (LSTM) network [1] as the internal predictive model.

2 Methodology

The strategy is evaluated using a calibrated anaerobic codigestion model based on ADM1 [2], which serves as a virtual pilot-scale plant treating for several months different mixtures of glycerine, gelatine, and pig manure.
Experimental Conditions: A 125-day stabilization buffer with constant inputs (50% of operating range) ensures all process variables, including slow dynamics like pH and biomass, reach steady state, preventing transient effects from corrupting the identification data.
System Identification: An LSTM network was trained to predict methane production using data from step changes in the three inlet flows ($Q_{str1}, Q_{str2}, Q_{str3}$). A 14-step sequence that explores the vertices of this operating space captures system dynamics. First-order autoregressive noise (AR(1), $\sigma = 1.5\%$, $\phi = 0.95$) added to inlet concentrations propagated disturbances, mimicking reality. The network acts as a recursive one-step-ahead predictor for the prediction horizon.
Controller Tuning: The NMPC minimizes a cost function penalizing tracking error and control effort (Δu). Bayesian Optimization determined penalty weights (Q, R) by minimizing a weighted Root Mean Square Error (RMSE) across closed-loop simulations covering the full operating range.

3 Results and Conclusions

The LSTM network accounts for noise in input concentrations, yielding smooth multi-step predictions that approximate the system dynamics.
Closed-loop simulations show successful tracking. Bayesian Optimization identified weights ($R_1 \approx 2.54, R_2 \approx 0.12, R_3 \approx 0.001$) revealing that Q_{str2} acts as the primary manipulated variable for dynamic corrections, with Q_{str1} providing moderate support. The system achieved a weighted tracking RMSE of 1.82 L/h (range 30–80 L/h), confirming the proposal's effectiveness.

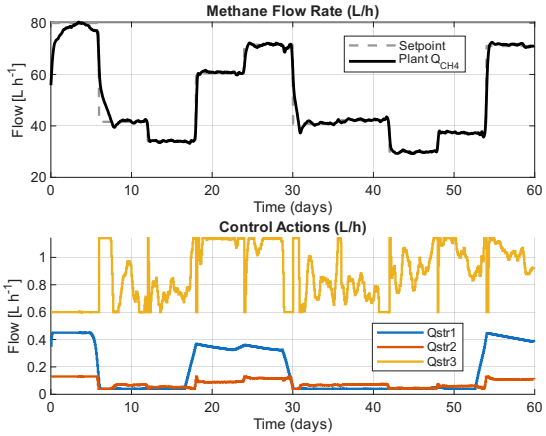

Figure 1: Closed-loop performance: Methane tracking (top) and control actions (bottom) after the 125-day stabilization.

The results indicate that the data-driven NMPC handles process non-linearities and input noise without requiring intermediate state measurements.

References

[1] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735-1780.
[2] Garcia-Gen, S., et al. (2014). Optimisation of substrate blends in anaerobic co-digestion using adaptive linear programming. *Bioresource Technology*, 173, 159-167.
[3] Pino-Santana, A., et al. (2025). Model predictive control of anaerobic digestion processes using a long short-term memory network predictor. *Water Science & Technology*, 92(7), 1063.

Improved Initialization for Port-Hamiltonian Neural Network Models[1]

Gé van Otterdijk, Siep Weiland, Maarten Schoukens
Control Systems, dept. of Electrical Engineering
Eindhoven University of Technology
g.j.e.v.otterdijk@tue.nl

1 Introduction

Port-Hamiltonian neural networks (PHNNs) have shown promising results in the identification of nonlinear dynamics of complex systems, as their combination of physical principles with data-driven learning allows for accurate modelling [1]. However, due to the non-convex optimization problem inherent in learning the correct network parameters, the training procedure is prone to converging to local minima, potentially leading to poor performance. In order to avoid this issue, this paper proposes an improved initialization for PHNNs. The core idea is to first estimate a linear port-Hamiltonian (PH) system to be used as an initialization for the network, after which the neural network adapts to the system nonlinearities, improving convergence and obtained model accuracy.

2 Proposed Method

A PHNN aims to identify a PH model of a nonlinear dynamical system based on input-output data. Such a parametrization is given by:

$$\begin{aligned} \dot{x} &= \big[J_\theta(x) - R_\theta(x)\big]\nabla H_\theta(x) + \big[G_\theta(x) - P_\theta(x)\big], \\ y &= \big[G_\theta(x) + P_\theta(x)\big]^\top \nabla H_\theta(x), \end{aligned} \tag{1}$$

where the matrices $J_\theta(x)$, $R_\theta(x)$, $H_\theta(x)$, $G_\theta(x)$ and $P_\theta(x)$ have to satisfy specific structural requirements, see [2]. In the existing method [1], these nonlinear mappings are parametrized solely through neural networks, which can lead to the optimization procedure converging to local optima. As an alternative, this work proposes to initialize the model with a linear PH estimate to serve as the starting point for the optimization, after which a similar neural network approach is utilised to learn the system nonlinearities. In practice, this comes down to a parametrization of the matrices from (1) according to $J_\theta(x) = J_{lin} + J_{NN}(x)$, etc., with the exception $H_\theta(x) = x^\top Q_{lin} x + H_{NN}(x)$, where all state-dependent nonlinear terms are initialized to yield 0 at the start of the optimization. As such, the model starts as the linear PH estimate but has the capacity to learn any nonlinearities present in the system.

[1]This work is funded by the European Union (ERC, COMPLETE, 101075836). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the grantng authority can be held responsible for them.

Two methods to obtain a linear PH estimate are considered. The first makes use of established system identification techniques to find a linear state-space estimate of the system [3], from which a linear PH estimate can then be derived through the Kalman-Yakubovich-Popov inequality. The second simply employs a linear PHNN to estimate a linear PH model.

3 Results

To validate the performance of the proposed method, a numerical implementation is studied based on data generated from a chained mass-spring-damper system. The nonlinear PHNN with linear initialization is compared against the standard nonlinear PHNN and a linear PHNN. 10 models are trained for each method, for which the convergence during model training is shown in Figure 1. From this figure, it becomes clear that the linear initialization improves over the other methods both by reducing the spread between models as well as achieving a lower normalized root mean squared error (NRMSE) within the fixed training duration.

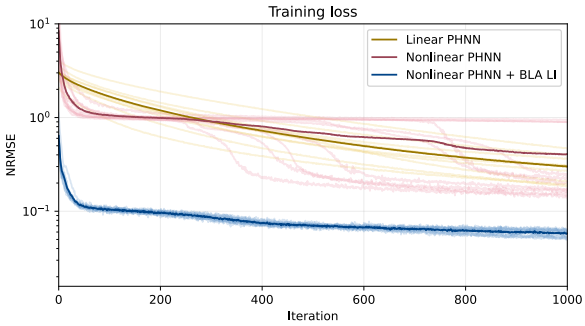

Figure 1: Model convergence during training. Individual model instances are opaque, with the method average in bold.

References

[1] S. Moradi, G. I. Beintema, N. Jaensson, R. Tóth, and M. Schoukens, "Port-hamiltonian neural networks with output error noise models," *arXiv:2502.14432*, 2025.

[2] A. J. van der Schaft and D. Jeltsema, *Port-Hamiltonian systems theory: an introductory overview*. Foundations and trends in systems and control, Now, 2014.

[3] J. Schoukens, M. Vaes, and R. Pintelon, "Linear system identification in a nonlinear setting: Nonparametric analysis of the nonlinear distortions and their impact on the best linear approximation," *IEEE Control Systems Magazine*, vol. 36, no. 3, pp. 38–69, 2016.

Encoder initialisation methods in the model augmentation setting

Jan H. Hoekstra[1], Bendegúz Györök[2], Roland Tóth[1,2], and Maarten Schoukens[1]

[1] Control Systems Group, Eindhoven University of Technology, The Netherlands

[2] Systems and Control Laboratory, Institute for Computer Science and Control, Budapest, Hungary

j.h.hoekstra@tue.nl, gyorokbende@sztaki.hu, r.toth@tue.nl, m.schoukens@tue.nl

1 Introduction

Recent encoder-based methods for *artificial neural network state-space* (ANN-SS) models have shown state-of-the-art performance for *Nonlinear system identification* (NL-SI) with improved computational efficiency, where the encoder is used to estimate the initial state allowing for batch optimisation methods. To address the lack of interpretability of these black-box ANN models, model augmentation approaches can be used. These combine prior available baseline models with the ANN learning components, resulting in faster convergence and more interpretable models. The combination of the encoder-based method with model augmentation has shown potential. However, so far the encoder has still been treated as a black-box function in the overall estimation process, while additional information in the form of the baseline model is available to predict the model state from past input-output data. We propose novel encoder initialisation approaches based on the available baseline model, resulting in improved noise robustness and faster convergence compared to black-box initialisation.

2 Method

We consider a true system and baseline model in a *discrete-time* (DT) NL-SS form with innovation noise. The general model augmentation formulation is given as in [1]

$$\begin{bmatrix} \hat{x}_{k+1} \\ \hat{y}_k \\ z_{1,k} \\ z_{2,k} \end{bmatrix} = \underbrace{\left[\begin{array}{ccc|cc} S_{xx} & S_{xu} & S_{xe} & S_{xw_1} & S_{xw_2} \\ \hline S_{z_1x} & S_{z_1u} & S_{z_1e} & S_{z_1w_1} & S_{z_1w_2} \\ S_{z_2x} & S_{z_2u} & S_{z_2e} & S_{z_2w_1} & S_{z_2w_2} \end{array}\right]}_{\mathbf{S}} \begin{bmatrix} \hat{x}_k \\ u_k \\ e_k \\ w_{1,k} \\ w_{2,k} \end{bmatrix} \tag{1a}$$

$$w_{1,k} = \phi_{\text{base}}\left(z_{1,k}\right) = \begin{bmatrix} f_{\text{base}}\left(z_{1,k}\right) \\ h_{\text{base}}\left(z_{1,k}\right) \end{bmatrix} \tag{1b}$$

$$w_{2,k} = \phi_{\text{aug}}\left(\theta_{\text{aug}}, z_{2,k}\right), \tag{1c}$$

where $\hat{y}_k$ is the model output, e_k an i.i.d. white noise process with finite variance, $\mathbf{S}$ is the interconnection matrix, S are selection matrices of appropriate size, ϕ_{base} is the baseline model, ϕ_{aug} are the learning components, parameterised by θ_{aug}, and $\hat{x}_k = [x_{\text{b},k}^\top \; x_{\text{a},k}^\top]^\top$ is the model state with $x_{\text{b},k}$ the baseline model states and $x_{\text{a},k}$ the augmented states introduced to model dynamics missing from the baseline model.

This work is funded by the European Union (Horizon Europe, ERC, COMPLETE, 101075836) and by the Air Force Office of Scientific Research under award number FA8655-23-1-7061. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

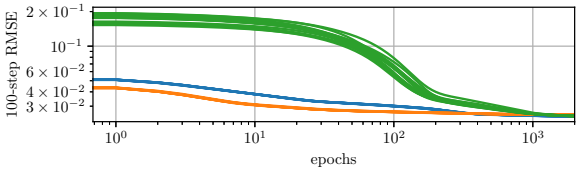

Figure 1: Validation loss for 100 step ahead RMSE over the fitting epochs for the three encoder initialisation methods: model-based (—), data-based (—) and randomly initialised (—).

The multiple-shooting-based truncated objective function [2], where the input-output dataset $\mathscr{D}_N$ is divided into N segments of length T, is given as:

$$V_{\text{trunc}}(\theta) = \frac{1}{2N(T+1)} \sum_{i=1}^{N} \sum_{\ell=0}^{T-1} \left\| \hat{y}_{k_i+\ell|k_i} - y_{k_i+\ell} \right\|_2^2 \tag{2a}$$

$$\begin{bmatrix} \hat{x}_{k_i+\ell+1|k_i} \\ \hat{y}_{k_i+\ell|k_i} \end{bmatrix} := \phi_\theta\left(\theta_{\text{aug}}, \hat{x}_{k_i+\ell|k_i}, u_{k_i+\ell}, e_{k_i+\ell}\right) \tag{2b}$$

$$\hat{x}_{k|k} := \psi_\theta\left(\theta_{\text{enc}}, y_{k-n_a}^{k-1}, u_{k-n_b}^{k-1}\right), \tag{2c}$$

where $\theta = \text{vec}(\theta_{\text{aug}}, \theta_{\text{enc}})$ is the joint parameter vector, ϕ_θ is the LFR-based augmentation structure (1) constructed from ϕ_{base}, ϕ_{aug} and $\mathbf{S}$, and $k+\ell|k$ indicates the state $\hat{x}_k$ or the output $\hat{y}_k$ at time $k+\ell$ simulated from the initial state $\hat{x}_{k|k}$ at time k. The initial state of these subsections is estimated by an encoder function ψ_θ from past input-output data (2c) where $u_{k-n_b}^{k-1} = [\, u_{k-n_b}^\top \; \cdots \; u_{k-1}^\top]^\top$ and $y_{k-n_a}^{k-1}$, $e_{k-n_a}^{k-1}$ is defined similarly.

The following initialization methods are proposed for the encoder ψ_θ: A **model-based** method, where the encoder is initialized as the analytically computed reconstructabiltiy map of the baseline model. Second, a **data-based** method which initialises the encoder as a pre-trained black-box approximation fitted to a simulated state trajectory of the baseline model for the input in the data set.

3 Result

We consider a nonlinear 2 *degrees-of-freedom* (DOF) *mass-spring-damper* (MSD) for the simulation study. The generated dataset $\mathscr{D}_N$ has additive noise resulting in a 20 dB signal-to-noise ratio. The baseline model is augmented with static parallel augemntation [1]. Through a Monte Carlo analysis the impact of the initialisation method is shown on the loss function convergence in Fig. 1.

References

[1] J. H. Hoekstra *et al.*, *Learning-based model augmentation with LFRs*, European Journal of Control, p. 101304, 2025.

[2] G. I. Beintema, *Data–driven Learning of Nonlinear Dynamic Systems: A Deep Neural State–Space Approach*, Phd Thesis, 2024.

Cable hysteresis identification and compensation using frequency-domain techniques

Dirk Alferink, Rob Fey, Nathan van de Wouw, Marcel Heertjes
Eindhoven University of Technology
d.w.t.alferink@tue.nl

1 Introduction

Stage positioning systems in lithography equipment, operate under demanding, high-speed motion trajectories while maintaining nanometer-scale precision. Therefore, it is critical to investigate all potential disturbance sources within wafer scanner machines to ensure future precision targets being met. One notable disturbance source is the cabling attached to the wafer stage delivering electrical power, cooling fluids, and immersion liquids to the stage. Moving cables introduce hysteretic dynamic disturbance forces deteriorating stage position accuracy. To enable feedforward compensation of these disturbances, a nonlinear frequency-domain identification method is introduced.

2 Nonlinear frequency-domain identification

Nonlinear friction-induced hysteretic disturbance forces w can be studied in the frequency-domain higher-order sinusoidal input describing functions (HOSIDFs), yielding remarkably accurate results for tracking error estimation [1]. In the current paper, HOSIDFs will be used for two-step identification of time-domain models for hysteresis. First, we experimentally estimate the HOSIDF column matrix for the hysteretic nonlinearity making use of a Lure' system description (Figure 1):

$$\underline{\hat{\mathbf{D}}}_{\phi,n} = \frac{\mathbf{H}_{ew}^{-1}}{\mathbf{H}_{xr}^{(1)}}\left(\frac{\hat{\mathbf{e}}_n}{A} - \mathbf{H}_{er}\right) \tag{1}$$

with $\mathbf{H}_{ew}, \mathbf{H}_{er}, \mathbf{H}_{xr}^{(1)}$ including FRFs of the linear part, A the amplitude of sinusoidal reference r, and column $\hat{\mathbf{e}}_n$ containing tracking error higher harmonics up to order n. Subsequently, after deriving a hysteresis model-based HOSIDF column $\underline{\mathbf{D}}_{\phi,n}(\theta)$, the model parameters stored $\hat{\theta}$ can be identified via

$$\hat{\theta} = \arg\min_{\theta} \|\underline{\hat{\mathbf{D}}}_{\phi,n} - \underline{\mathbf{D}}_{\phi,n}(\theta)\|_2^2. \tag{2}$$

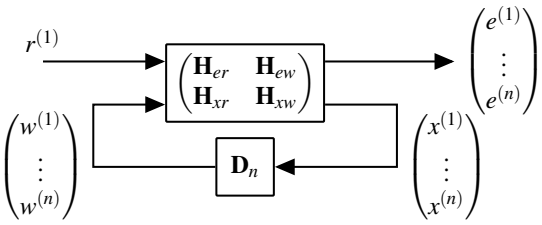

Figure 1: Lur'e system decomposed in harmonic components. Reference $r^{(1)}$, hysteretic disturbance $\mathbf{w}_n$, tracking error $\mathbf{e}_n$, and stage displacement $\mathbf{x}_n$. Matrix $\mathbf{H}$ is a frequency-domain description of the motion system, and $\mathbf{D}_n$ the HOSIDF matrix of nonlinearity ϕ.

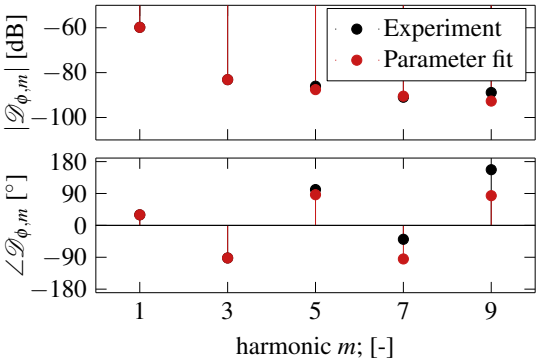

Figure 2: Parameter fit vs. experimentally obtained HOSIDF.

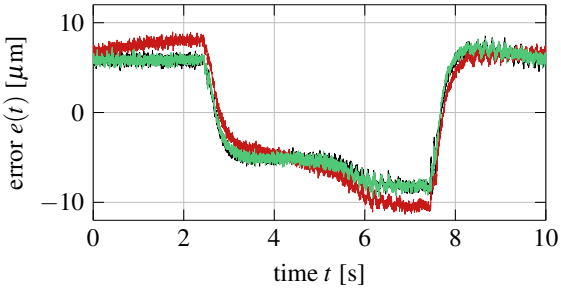

Figure 3: Tracking error without [——], with [——] cable, and with cable and feedforward control using model parameters $\hat{\theta}$ [——].

3 Experimental demonstration

An experimental motion system with a cable is used for demonstration of the method. Spectral analysis of the describing function shows that (2) admits an unique solution $\hat{\theta}$ with zero argument if and only if harmonics $m = 1, 3$ are included. The identified and experimental HOSIDFs are compared in Figure 2. Figure 3 shows that application of feedforward hysteresis compensation results in resemblance of the situation without cable as desired.

4 Conclusions and future work

The nonlinear frequency-domain method enables identification and compensation of cable-induced hysteresis forces. Future work will include cable inertia modeling, and higher harmonics parameter optimization with higher harmonics $m > 3$ for better model accuracy.

References

[1] D. W. T. Alferink, et al. Hysteresis in motion control systems: a frequency-domain analysis on higher harmonics, in *ACC, 2025*, pp. 5073–5078.

System identification of power exhaust dynamics in the X-Point Target divertor at Tokamak à Configuration Variable

M. Winkel[1,2], K. Verhaegh[1], B. Kool[2], G. Derks[2], K. Lee[3], M. Carpita[3], A. Perek[3], R.I. Morgan[3], M. Zurita[3], C. Theiler[3], D. Brida[4], M. van Berkel[2], the TCV Team[a], and the EUROfusion Tokamak Exploitation team[b]

[1]Eindhoven University of Technology (TU/e), Eindhoven, Netherlands
[2]Dutch Institute for Fundamental Energy Research (DIFFER), Eindhoven, Netherlands
[3]École Polytechnique Fédérale de Lausanne (EPFL) – Swiss Plasma Center (SPC), Lausanne, Switzerland
[4]Max Planck Institute for Plasma Physics, Garching, Germany
[a]see author list of B.P. Duval et al 2024, Nucl. Fusion, 64 112023
[b]see author list of E. Joffrin et al 2024, Nucl. Fusion 64 112019

1 Introduction

Managing power exhaust is a critical hurdle in the realization of a nuclear fusion reactor as a sustainable energy source. In a fusion reactor, strong magnetic fields are used to form a plasma with temperatures exceeding 100 million °C. Consequently, the exhaust of the reactor (the 'divertor') must handle extreme heat fluxes that would otherwise damage the walls [1]. To protect the divertor surfaces, fuel and impurity gas injection is used to induce 'plasma detachment', a state that significantly reduces heat flux to the walls. Maintaining this state within a narrow operational window remains an open control challenge. Effective regulation is complicated by a significant time-scale separation as the rejection of fast plasma disturbances is hindered by relatively slow gas valve actuators and significant transport delays in the gas system [2].

An alternative divertor magnetic geometry may alleviate this challenge, with the X-point Target (XPT) divertor as a promising candidate. The XPT divertor relaxes the operational constraints on maintaining the detached state with respect to conventional-like (SN) divertor designs and is hypothesized to feature a built-in disturbance rejection capacity in quasi-static experiments [3]. However the dynamic advantage of the XPT over conventional designs has not yet been experimentally quantified. Addressing this knowledge gap is the main focus of this contribution.

2 System identification

A system identification approach is used to quantify the XPT's inherent disturbance-rejection capacity compared to a SN divertor. Experiments on the experimental fusion device TCV [4] are performed, applying sinusoidal perturbations across three primary input channels; gas fuelling injection, gas impurity injection and core heating. The response of the divertor plasma is observed across various diagnostics. We characterize the system's frequency response, employing the Local Polynomial Method (LPM) to correct for slow transients and drifts in the experiments [5, 6]. We show that the detached state in the XPT divertor features an inherent disturbance-rejection capacity to all perturbation types compared to the SN. The disturbance rejection of the XPT divertor could be highly beneficial for future reactors to maintain the detached state within operational constraints.

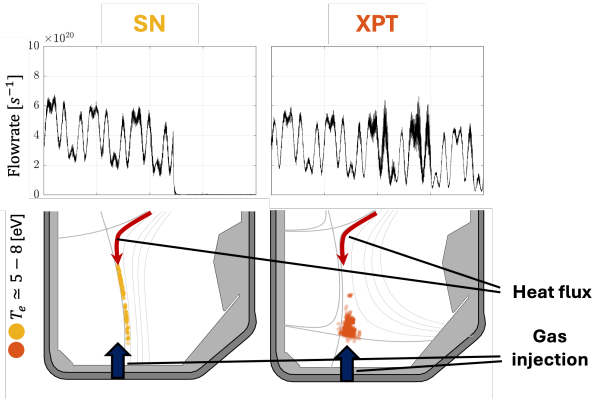

Figure 1: Similar sinusoidal gas fuelling perturbations (top) induce less movement in a cold temperature front (bottom) of $T_e \simeq 5-8$ eV for an XPT compared to SN divertor in TCV, indicative of heat flux buffering by the XPT.

References

[1] H. Zohm *et al.* 2021 *Fusion Eng. Des.* **166** 112307
[2] M. van Berkel *et al.* 2026 *Nucl. Fusion* **66** 036017
[3] K. Lee *et al.* 2025 *Phys. Rev. Lett.* **134** 185102
[4] B. Duval *et al.* 2024 *Nucl. Fusion* **64** 112023
[5] R. Pintelon and J. Schoukens 2012 *System Identification: A Frequency Domain Approach* 2nd edn (Hoboken, NJ: Wiley-IEEE Press)
[6] M. van Berkel *et al.* 2020 *Plasma Phys. Control. Fusion* **62** 094001

Acknowledgements

DIFFER is part of the institutes organisation of NWO. This work has been carried out within the framework of the EUROfusion Consortium, partially funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 - EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union, the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Least Costly Space-Filling Experiment Design for the Identification of a Nonlinear System

Máté Kiss[1], Maarten Schoukens[1], Roland Tóth[1,2],
[1] Control Systems, Eindhoven University of Technology, Eindhoven, Netherlands (m.kiss@tue.nl)
[2] Institute for Computer Science and Control, Budapest, Hungary

1 Introduction

The quality of an estimated nonlinear model highly depends on the data quality that was used for the system identification. By using a Gaussian Process (GP)-based optimal input design approach [1], a so-called space-filling dataset can be generated in the feature space of the system model. However, the resulting input design can be costly to apply to the real system. Therefore we propose a space-filling input design strategy that can minimize the experimentation cost in terms of a user defined measure, while still guaranteeing a prescribed level of space-fillingness. Through a Monte Carlo simulation study we demonstrate that the proposed method can appropriately shape the excitation signal to significantly reduce the experimental cost compared to a classical optimality criteria [2] while the identified model performance remains adequate.

2 Method

We address the space-filling input design problem for a controllable, deterministic system model:

$$\begin{aligned} y(k) &= f(u(k-1),\dots,u(k-m_{\mathrm{u}}),y(k-1),\dots,y(k-m_{\mathrm{y}})) \\ &= f(x(k)), \end{aligned} \tag{1}$$

where $u(k) \in \mathscr{U} \subseteq \mathbb{R}^{n_{\mathrm{u}}}$ and $y(k) \in \mathscr{Y} \subseteq \mathbb{R}^{n_{\mathrm{y}}}$ are the input and the output of the system at time instant $k \in \mathbb{I}_1^N$. The maximum time lags for the respective signals are defined by m_{y}, m_{u}. The feature space of the model is defined by the vector $x(k) \in \mathscr{X} \subseteq \mathbb{R}^{n_{\mathrm{x}}}$, where $n_{\mathrm{x}} = n_{\mathrm{u}} m_{\mathrm{u}} + n_{\mathrm{y}} m_{\mathrm{y}}$, and $\mathscr{X} : \mathscr{U} \times \mathscr{Y}$ denotes the joint delayed input-output space. The nonlinear function $f(\cdot) : \mathscr{X} \to \mathscr{Y}$, $f \in \mathscr{C}_1$ represents the system model.

The least costly space-filling input design aims to design an input sequence $\{u_k\}_{k=1}^N$ parametrized by $\theta \in \Theta$ being the set of parameters to be optimized. The goal is to obtain an input signal with the lowest experiment cost $\mathscr{W}(\theta) : \Theta \to \mathbb{R}$ which, when applied to the system model (1) results in the data set $\mathscr{D}_N = \{x_k\}_{k=1}^N$ with $x_k = \mathrm{vec}(u_i(k-1),\dots,u_i(k-m_{\mathrm{u}}),y_i(k-1),\dots,y_i(k-m_{\mathrm{y}}))$, such that it guarantees a user-chosen level of space-fillingness γ over the region of interest $\mathscr{X}$.

Based on the work of [1], the space-filling accuracy constraint is expressed by the classical V-optimal cost function:

Funded by the European Union (ERC, COMPLETE, 101075836). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

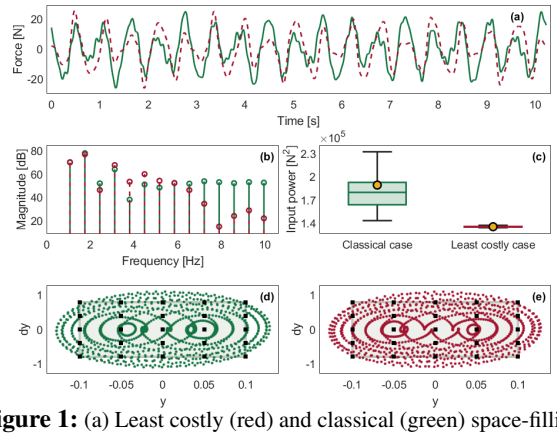

Figure 1: (a) Least costly (red) and classical (green) space-filling signal, (b) excited frequencies of signals, (c) box plot of input powers with their mean value (yellow circle), (d,e) classical and least costly (red) space-filling dataset.

$$\mathscr{V}(\mathscr{D}_N(\theta)) := \frac{1}{M}\sum_{i=1}^{M} \kappa(\tilde{x}_i,\tilde{x}_i) - \kappa(\tilde{x}_i,X) K_N^{-1} \kappa(\tilde{x}_i,X)^\top, \tag{2}$$

where M denotes the number of discrete grid points $\tilde{x}_i$ in the region of interest and κ is a kernel function which is chosen to be an RBF kernel.

3 Simulation Study

In our example, the system represented by (1) corresponds to a second order nonlinear mass-spring-damper system. The excitation signal is a multisine parametrzied by its phases and amplitudes $\theta = \mathrm{vec}(A_1,\dots,A_l,\phi_1,\dots,\phi_l)$ with length $N = 1024$. The least costly input design approach is tested in a Monte Carlo study with 50 realizations, where we compare it to the classical space-filling approach (see Fig.1).
Next, nonlinear output error models were identified with analytical-gradient based optimization. Results show that by including the experiment cost in the space-filling design problem, the experimental cost could be reduced significantly without compromising the input quality or the quality of the estimated models.

References

[1] Y. Liu, M. Kiss, M. Schoukens, R. Tóth, On Space-Filling Input Design for Nonlinear Dynamic Model Learning: A Gaussian Process Approach, IEEE Control System Letters, 2024
[2] A. Atkinson, A. Donev, Optimum experimen- tal designs, clarendon press, Oxford, 1992
[3] X. Bombois, G. Scorletti, Design of least costly identification experiments: The main philosophy accomanied by illustrative examples, Journal Européen des Systemes Automatisés, 2012

A MATLAB toolbox for OMSE-based optimal experimental design in nonlinear bioprocess models

Carlos Segura Piña[a] carlosroberto.segurapina@umons.ac.be
Pierre Denis[a,b] pierre.denis444@gmail.com
Constantinos Theodoropoulos[b] c.theodoropoulos@manchester.ac.uk
Laurent Dewasme[a] laurent.dewasme@umons.ac.be
Alain Vande Wouwer[a] alain.vandewouwer@umons.ac.be

[a]Systems, Estimation, Control and Optimization (SECO), University of Mons, Belgium
[b]Department of Chemical Engineering, The University of Manchester, United Kingdom

1 Abstract

Optimal Experimental Design (OED) plays a key role in improving parameter identifiability in nonlinear dynamic models by maximizing the statistical information extracted from limited experimental resources. However, in bioprocess applications, its practical adoption often remains limited by computational complexity, lack of accessible software tools, or limited awareness among practitioners. [1]

A novel OED criterion, the Overall Mean Squared Error (OMSE), was recently proposed by the authors [2], in which the OMSE matrix formulates the expected squared estimation error before running the experiment:

$$\mathrm{OMSE}(\mathcal{D}) \;=\; \mathbb{E}_{\theta,Y}\Big[(\hat{\theta}(Y)-\theta)(\hat{\theta}(Y)-\theta)^{\top}\Big], \qquad (1)$$

where $\mathcal{D}$ denotes the vector of design variables, including initial conditions, input profiles, and measurement times. An accurate and practical approximation of (1) is obtained via an expected Cramér–Rao bound (ECRB), and the matrix-valued target is converted into a scalar via a standard optimality criterion:

$$\Phi(\mathcal{D}) \;=\; \phi\big(\mathrm{OMSE}(\mathcal{D})\big). \qquad (2)$$

Hence, the robust OMSE-based OED problem is

$$\mathcal{D}_{\mathrm{opt}} = \arg\min_{\mathcal{D}} \; \Phi\big(\mathcal{D}; \mathcal{F}, p_{\theta}^{-}, R(\cdot)\big) \quad \text{s.t.} \quad \mathcal{D}\in\mathcal{C}, \qquad (3)$$

where $\mathcal{F}$ denotes the nonlinear state-space model, p_{θ}^{-} the prior parameter distribution, $R(\cdot)$ the measurement-noise covariance, and $\mathcal{C}$ encodes experiment-specific constraints.

This contribution presents OEDLab, a MATLAB toolbox for robust OED based on the OMSE criterion, which requires a mathematical model in state-space form together with prior parameter information. Based on these inputs, OEDLab optimizes the experimental design variables and provides *a priori* predictions of parameter estimation error after model calibration. The overall structure and workflow of the toolbox are illustrated in Fig. 1.

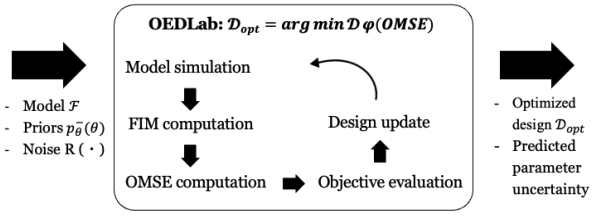

Figure 1: OEDLab workflow.

OEDLab is demonstrated on the identification of a variety of nonlinear bioprocess models under batch and fed-batch operating conditions. The toolbox is applied to multistate systems with varying numbers of parameters, where prior parameter distributions are either specified or constructed from previously collected experimental data and incorporated into the design procedure.

The performance of the OMSE-based uncertainty prediction is assessed through a Monte Carlo validation study, in which synthetic datasets are generated under the optimized experimental conditions and parameters are re-estimated. The empirical estimation errors show very close agreement with the OMSE-based predictions across nominal and inflated prior uncertainty levels, demonstrating that OEDLab delivers consistent *a priori* uncertainty forecasts for nonlinear bioprocess models.

References

[1] V. Kozachynskyi, M. Illner, E. Esche, and J.-U. Repke, "The optimal experiment? Influence of solution strategies on model-based optimal experimental design," *Computers & Chemical Engineering*, 2024.

[2] P. Denis, C. Theodoropoulos, and A. Vande Wouwer, "Robust optimal experimental design for identifiability using the overall mean squared estimation error," submitted to *Computers & Chemical Engineering*, 2025.

Acknowledgment This work is supported by the Win4Excellence project GT4Health (#2310145).

Contraction-preserving unconstrained augmentation of Lur'e-type systems*

Adarsh Subrahamanian Moosath, Fahim Shakib, Rob Fey, and Nathan van de Wouw
Eindhoven University of Technology, The Netherlands
Emails: {a.subrahamanian.moosath, m.f.shakib, r.h.b.fey, n.v.d.wouw}@tue.nl

1 Introduction

In many practical applications, first-principles (FP) models are formulated as linear time-invariant (LTI) systems by neglecting complex nonlinearities, which leads to systematic discrepancies between the model and the true plant behaviour. To address this, we propose a novel model augmentation framework for discrete-time Lur'e-type systems, where an existing LTI FP model is augmented with unmodelled static nonlinearities captured using a neural network (NN). The resulting hybrid model is guaranteed to be contractive and convergent. Contraction ensures that trajectories converge towards each other irrespective of initial conditions, while convergence guarantees the existence of a unique, bounded, periodic steady-state response for any bounded periodic input. Together, these properties provide stability guarantees analogous to global exponential stability (GES) of the LTI FP model. The proposed framework enables unconstrained learning, promoting scalability to large-scale systems, and allows training using steady-state data.

2 Lur'e-type hybrid models

The Lur'e-type hybrid model takes the discrete-time state-space form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + B_3\psi_k, \quad y_k = Cx_k + Du_k, \\ \mu_k &= C_3x_k + D_3u_k, \quad \psi_k = \varphi(\mu_k), \end{aligned} \tag{1}$$

where $u_k \in \mathbb{R}^{n_u}$, $x_k \in \mathbb{R}^{n_x}$ and $y_k \in \mathbb{R}^{n_y}$ denote the input, state, and output vector of the hybrid model, respectively. Matrices A, B, C, and D are inherited from an existing GES LTI FP model. The static nonlinear mapping $\varphi : \mathbb{R}^{n_\mu} \to \mathbb{R}^{n_\psi}$ captures unmodelled nonlinearities and is parameterised using a neural network. The interconnection structure between the linear and nonlinear parts of the hybrid model is defined by user-specified matrices B_3, C_3, and D_3, which are selected based on prior knowledge of the system's nonlinear structure. Consequently, all learnable hybrid model parameters are those characterising the nonlinear mapping φ. To guarantee contraction and convergence of the hybrid model, the neural network parametrising φ must satisfy a prescribed Lipschitz bound. While several approaches exist to estimate Lipschitz constants for neural networks, scalable methods are required for large-scale models. To this end, we employ a Lipschitz-bounded equilibrium network (LBEN) [1], which enforces the required bound by construction and enables unconstrained learning of a contractive and convergent hybrid model. Moreover, the unconstrained parameterisation of the equilibrium network enables the use of generic NN optimisation algorithms such as Adam.

*This work is part of the MASC project co-financed by Holland High Tech, top sector High-Tech Systems and Materials, with a PPP innovation subsidy for public-private partnerships for research and development.

3 Application

To demonstrate the proposed model augmentation framework, we consider a benchmark problem based on the Duffing oscillator. The system comprises a linear mass-spring-damper FP model augmented with a cubic stiffness nonlinearity captured using a LBEN. The dataset is generated by harmonically exciting the system on a frequency and amplitude grid, with transients removed prior to training. Owing to the contraction and convergence properties of the true plant, trajectories converge to a unique steady-state solution regardless of initial conditions. Table 1 compares the mean squared error (MSE) of the FP and hybrid model predictions across different datasets. The hybrid model improves accuracy by more than two orders of magnitude over the FP model, demonstrating its ability to reconstruct the true nonlinear plant behaviour. Figure 1 shows that the learned nonlinear force–displacement relationship φ resembles the true static nonlinearity φ^* well in the visited region of μ. Figure 2 compares the time domain output response of the true plant $\acute{y}$, the FP model $\tilde{y}$, and the hybrid model y.

Tab. 1: MSE between FP and hybrid model predictions

Model / MSE	Training	Validation	Test 1	Test 2
FP model	0.2182	0.2342	0.2488	0.1091
Hybrid model	4.27×10^{-4}	4.22×10^{-4}	5.05×10^{-4}	1.82×10^{-4}

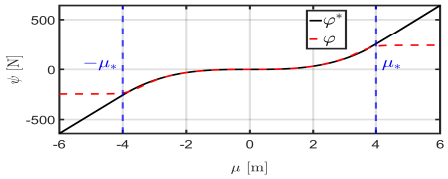

Fig. 1: The learned cubic nonlinearity (red dashed line) is compared against the true cubic nonlinearity (black line).

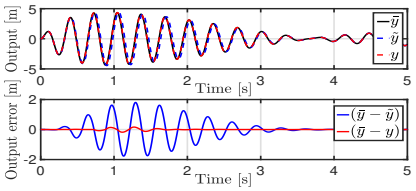

Fig. 2: (Top) Time-domain output of the true plant, the FP model, and the hybrid model. (Bottom) Corresponding output errors, between the hybrid model and the FP model.

4 Conclusions

The proposed framework combines LTI FP models with NNs to capture unmodelled nonlinearities, ensuring contraction and convergence while demonstrating improved predictive accuracy on a Duffing oscillator use case.

References

[1] R. Wang, et. al., "Direct parameterization of Lipschitz-bounded deep networks," *ICML*, 2023.

Learning Symbolic Control Barrier Functions from Data via Transformer-Based Generation

Lufeng Xu
DTPA, ENTEG, University of Groningen
Email: lufeng.xu@rug.nl

Ming Cao
DTPA, ENTEG, University of Groningen
Email: m.cao@rug.nl

1 Introduction

This paper addresses the problem of learning explicit symbolic expressions of control barrier functions (CBFs) from data. While CBFs provide an effective and widely adopted mechanism for enforcing safety, their explicit symbolic forms are often unavailable in practical applications. In contrast, collecting data that reflect safe operational behaviours is comparatively straightforward. Motivated by this observation, a transformer-based generative framework is proposed to synthesise candidate symbolic CBF expressions from data. Compared with existing learning-based CBF approaches, the key novelty of the proposed method lies in its ability to produce explicit and interpretable symbolic CBF expressions, rather than implicit or black-box representations.

2 The CBF Generation Method

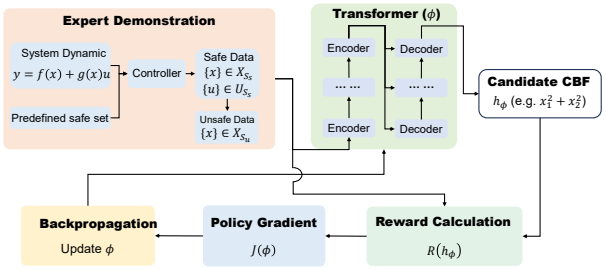

Figure 1: Overall structure of the proposed framework.

$$\begin{aligned} \min_h \quad & \|h\| \\ \text{s.t.} \quad & h(x_i) \geq \gamma_1, \forall x_i \in X_{S_s}, \\ & \Delta h(x_i, u_i) + \alpha(h(x_i)) \geq \gamma_2, \forall (x_i, u_i) \in (X_{S_s}, U_{S_s}), \\ & h(x_i) \leq -\gamma_3, \forall x_i \in X_{S_u}. \end{aligned} \tag{1}$$

The constraints imposed on the generated CBF h are given in (1). Satisfaction of these constraints guarantees that the learned CBF h ensures safety over the continuous region inferred from the discrete data points.

The overall framework is divided into the following components:

Data collection: Safe datasets X_{S_s} and U_{S_s} are collected from expert demonstrations. Unsafe data X_{S_u} are synthesized from neighborhoods of the estimated safety boundary.

Transformer generation: The model ingests datasets and decodes a symbolic candidate $h_\phi(x)$ as a sequence of tokens.

Reinforcement: A composite reward based on constraints in Eq. (1) is computed to evaluate the quality of the generated expression, $R = w_1 \ell_{\text{safe}} + w_2 \ell_{\text{unsafe}} + w_3 \ell_{\text{dyn}} + w_4 \ell_{\text{var}}$.

Backpropagation: The reward is used to update the Transformer parameters via backpropagation. $J(\phi) = \mathbb{E}[R]$, $\phi \leftarrow \phi + \eta \nabla_\phi J(\phi)$.

3 Simulation Results

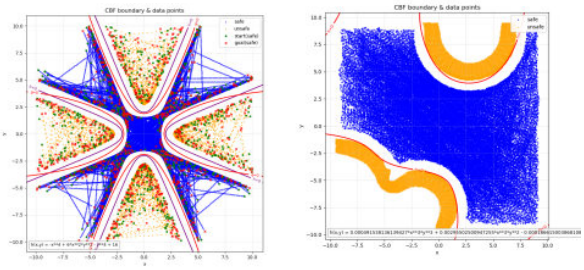

Figure 2: Simulation results.

In Fig. 2, the blue points correspond to the safe trajectories, while the yellow points represent the unsafe dataset. The red curve denotes the generated function. It can be observed that the learned function effectively distinguishes between the safe and unsafe data.

4 Conclusion

This work presented a data-driven approach for learning explicit symbolic control barrier functions. The proposed method captures the relationship between observed data and the underlying CBF structure, and, by exploiting the expressive power of transformer architectures, generates interpretable symbolic CBF expressions.

References

[1] A. Robey, H. Hu, L. Lindemann, et al., "Learning control barrier functions from expert demonstrations". *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 3717–3724, 2020.

Feedback Stabilization of Polynomial Systems: From Model-based to Data-driven Methods

Huayuan Huang, Kanat Camlibel, Raffaella Carloni, Henk van Waarde
Bernoulli Institute, University of Groningen, The Netherlands
huayuan.huang@rug.nl, m.k.camlibel@rug.nl, r.carloni@rug.nl, h.j.van.waarde@rug.nl

1 Introduction

The control of polynomial systems has gained considerable attention. In this study, we propose new global stabilization approaches for a class of polynomial systems in both model-based and data-driven settings [1].

2 Problem statement

Consider the polynomial input-affine system

$$\begin{bmatrix}\dot{x}_1(t)\\ \dot{x}_2(t)\end{bmatrix} = \begin{bmatrix}A_1\\ A_2\end{bmatrix} F(x(t)) + \begin{bmatrix}0\\ B_2\end{bmatrix} G(x(t))u(t), \qquad (1)$$

where $x_1(t) \in \mathbb{R}^{n_1}$, $x_2(t) \in \mathbb{R}^{n_2}$, $x(t) = \begin{bmatrix}x_1^\top(t) & x_2^\top(t)\end{bmatrix}^\top \in \mathbb{R}^n$ is the state, and $u(t) \in \mathbb{R}^m$ is the input. The parameter matrices are given by $A_1 \in \mathbb{R}^{n_1 \times f}$, $A_2 \in \mathbb{R}^{n_2 \times f}$ and $B_2 \in \mathbb{R}^{n_2 \times g}$. The functions $F \in \mathbb{R}^f[x]$ and $G \in \mathbb{R}^{g\times m}[x]$ represent polynomial mappings of the state. We first focus on model-based control, and our objective is to design a rational state feedback controller $u = K(x)$ that stabilizes the system (1).

Second, we aim to find a globally stabilizing controller for polynomial systems using data. Given F and G, we denote the system of the form (1) by (A_1, A_2, B_2). We consider a true system (A_{s1}, A_{s2}, B_{s2}), where the parameter matrices A_{s1}, A_{s2} and B_{s2} are assumed to be unknown. Instead, we are given the input data $\mathscr{U} := \begin{bmatrix}u(t_0) & \cdots & u(t_{T-1})\end{bmatrix}$, the state data $\mathscr{X} := \begin{bmatrix}x(t_0) & \cdots & x(t_{T-1})\end{bmatrix}$ and the state derivative data $\dot{\mathscr{X}} := \begin{bmatrix}\dot{x}(t_0) & \cdots & \dot{x}(t_{T-1})\end{bmatrix}$, where T is a positive integer and $t_0 < \cdots < t_{T-1}$. These data are generated by the true system influenced by unknown noise $\mathscr{W} := \begin{bmatrix}w(t_0) & \cdots & w(t_{T-1})\end{bmatrix}$, such that

$$\dot{\mathscr{X}} = \begin{bmatrix}A_{s1}\\ A_{s2}\end{bmatrix} \mathscr{F} + \begin{bmatrix}0\\ B_{s2}\end{bmatrix} \mathscr{G}_{\mathscr{U}} + \mathscr{W}, \qquad (2)$$

where $\mathscr{F} := \begin{bmatrix}F(x(t_0)) & \cdots & F(x(t_{T-1}))\end{bmatrix}$ and $\mathscr{G}_{\mathscr{U}} := \begin{bmatrix}G(x(t_0))u(t_0) & \cdots & G(x(t_{T-1}))u(t_{T-1})\end{bmatrix}$. Although the noise matrix $\mathscr{W}$ is unknown, we assume that it satisfies the following quadratic inequality:

$$\begin{bmatrix}1\\ \mathrm{vec}(\mathscr{W}^\top)\end{bmatrix}^\top \begin{bmatrix}\varphi & 0\\ 0 & -I\end{bmatrix} \begin{bmatrix}1\\ \mathrm{vec}(\mathscr{W}^\top)\end{bmatrix} \geq 0, \qquad (3)$$

where $\varphi \geq 0$ is given. A system (A_1, A_2, B_2) satisfying

$$\dot{\mathscr{X}} = \begin{bmatrix}A_1\\ A_2\end{bmatrix} \mathscr{F} + \begin{bmatrix}0\\ B_2\end{bmatrix} \mathscr{G}_{\mathscr{U}} + \mathscr{W}, \qquad (4)$$

for some $\mathscr{W}$ as in (3) is called *compatible* with the data. We denote the set of all systems compatible with the data $(\dot{\mathscr{X}}, \mathscr{X}, \mathscr{U})$ by Σ, i.e.,

$$\Sigma := \{(A_1, A_2, B_2) \mid (4) \text{ holds for some } \mathscr{W} \text{ satisfying } (3)\}.$$

It follows from (2) that $(A_{s1}, A_{s2}, B_{s2}) \in \Sigma$, but in general Σ contains other systems. On the basis of the data, the true system cannot be distinguished from other systems in Σ. We can only guarantee that a controller stabilizes the true system if it stabilizes all systems in Σ. Therefore, the goal is to find a single controller that stabilizes the origin of all systems compatible with the data.

3 Results

We first derive a new feedback stabilization approach for polynomial systems in the model-based setting. Compared to [2], the proposed approach imposes different sum-of-squares constraints [3] and allows a broader class of Lyapunov function candidates. Second, we extend this approach to the data-driven setting by employing a specialized S-lemma [4]. Using data corrupted by bounded noise, we derive design conditions for globally stabilizing controllers for unknown polynomial systems. Our data-driven approach incorporates prior knowledge of system parameters, which is demonstrated to reduce conservatism in numerical examples. Both approaches are based on a technical lemma that provides new conditions for global asymptotic stability analysis, offering an alternative to radially unbounded Lyapunov functions.

References

[1] H. Huang, M. K. Camlibel, R. Carloni, and H. J. van Waarde, "Feedback stabilization of polynomial systems: From model-based to data-driven methods," 2025, *arXiv:2505.14457*.

[2] S. Prajna, A. Papachristodoulou, and F. Wu, "Nonlinear control synthesis by sum of squares optimization: A Lyapunov-based approach," in *Proc. Asian Control Conf.*, 2004, pp. 157–165.

[3] P. A. Parrilo, "Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization," Ph.D. dissertation, California Institute of Technology, Pasadena, California, 2000.

[4] H. Huang, M. K. Camlibel, R. Carloni, and H. J. van Waarde, "Data-driven stabilization of polynomial systems using density functions," 2025, *arXiv:2503.07092*.

Scalable Nonlinear DeePC: Bridging Direct and Indirect Methods and Basis Reduction

Thomas de Jong
Control Systems, TU Eindhoven, NL
t.o.d.jong@tue.nl

Mircea Lazar
Control Systems, TU Eindhoven, NL
m.lazar@tue.nl

Siep Weiland
Control Systems, TU Eindhoven, NL
s.weiland@tue.nl

Florian Dörfler
Automatic Control, ETH Zürich, CH
dorfler@ethz.ch

1 Introduction

Willems' fundamental lemma underlies the recently develop data-enabled predictive control (DeePC) method [1]. For linear systems, DeePC, model predictive control (MPC), and subspace predictive control (SPC) are equivalent, but this equivalence generally does not hold for nonlinear DeePC or in the presence of noise. The Π-regularization links DeePC to SPC and improves performance [2]. Nonlinear DeePC extensions using basis functions exist [3] but typically lead to large-scale problems. In this work, we mitigate this via SVD and sparse LASSO selection to obtain a scalable framework for nonlinear DeePC. Furthermore, we extend the Π-regularization to basis functions DeePC and establish theoretical links to non-reduced basis functions SPC.

2 Scalable nonlinear DeePC

We consider nonlinear discrete-time systems:

$$\begin{aligned} x_{k+1} &= f(x_k, u_k), \quad k \in \mathbb{N}, \\ y_k &= h(x_k), \end{aligned} \tag{1}$$

where f and h are *unknown*, and $\mathbb{U}, \mathbb{Y}$ are are sets that specify input and output constraints. To identify a predictor, the system is excited with a persistently exciting input, collecting $\mathscr{D} = \{u_i, y_i\}_{i=0}^{T_{\text{ini}}+T+N-1}$. Hankel matrices $\mathbf{U}_p, \mathbf{Y}_p, \mathbf{U}_f, \mathbf{Y}_f$ are built from this trajectory, and the lifted data matrix $\Phi = \bar{\phi}(\mathbf{U}_p, \mathbf{Y}_p, \mathbf{U}_f) \in \mathbb{R}^{L \times T}$ encodes the basis-function representation of all data [3]. Assuming $T \gg Np + L$, the SVD of the stacked matrix

$$\begin{bmatrix} \Phi \\ \mathbf{Y}_f \end{bmatrix} = U \begin{bmatrix} \Sigma_1 & 0 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix} = U \Sigma_1 V_1^\top, \quad \begin{bmatrix} \tilde{\Phi} \\ \tilde{\mathbf{Y}}_f \end{bmatrix} = \begin{bmatrix} \Phi \\ \mathbf{Y}_f \end{bmatrix} V_1,$$

with $\Sigma_1 \in \mathbb{R}^{(L+Np)\times(L+Np)}$, can be used to define the following reduced DeePC problem.

Problem 2.1 ($\tilde{\Phi}$-DeePC-$\tilde{\Pi}$)**.**

$$\min_{\Xi_k} \ell\left(\mathbf{y}_{[1,N]}(k), \mathbf{u}_{[0,N-1]}(k)\right) + \lambda \left\|(I - \tilde{\Phi}^\dagger \tilde{\Phi}) \tilde{\mathbf{g}}_k\right\|_p \tag{2a}$$

$$\text{s.t.} \quad \tilde{\Phi} \tilde{\mathbf{g}}_k = \bar{\phi}\left(\mathbf{x}_{\text{ini}}(k), \mathbf{u}_{[0,N-1]}(k)\right), \tag{2b}$$

$$\tilde{\mathbf{Y}}_f \tilde{\mathbf{g}}_k = \mathbf{y}_{[1,N]}(k), \tag{2c}$$

$$\left(\mathbf{y}_{[1,N]}(k), \mathbf{u}_{[0,N-1]}(k)\right) \in \mathbb{Y}^N \times \mathbb{U}^N. \tag{2d}$$

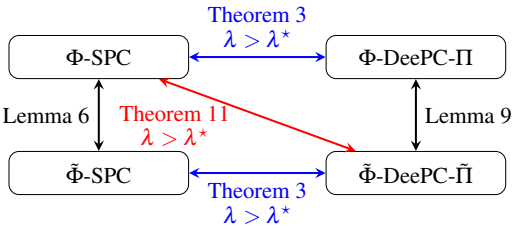

Figure 1: Relation diagram for main results.

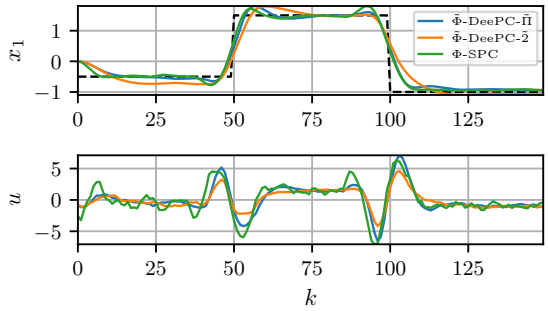

Figure 2: Tracking example for $\tilde{\Phi}$-DeePC-$\tilde{\Pi}$, $\tilde{\Phi}$-DeePC-2 with $\lambda = 10^4$ and Φ-SPC for noisy data.

To extend the Π-regularization from [2] to basis functions DeePC, we show that for $\lambda > \lambda^\star$, with $\lambda^\star$ the Lipschitz constant of ℓ, the feasible set coincides with non-reduced Φ-SPC [3]; for $\lambda \leq \lambda^\star$, the $\tilde{\Phi}$-DeePC-$\tilde{\Pi}$ problem relaxes the non-reduced Φ-SPC corresponding problem. The main results are highlighted in Figure 1, and simulations Figure 2 show $\tilde{\Phi}$-DeePC-$\tilde{\Pi}$ outperforms non-reduced Φ-SPC and 2-norm regularized DeePC under measurement noise.

References

[1] J. Coulson, J. Lygeros, and F. Dörfler, "Data-Enabled Predictive Control: In the Shallows of the DeePC," in *ECC*, 2019.

[2] F. Dörfler, J. Coulson, and I. Markovsky, "Bridging Direct and Indirect Data-Driven Control Formulations," *IEEE Transactions on Automatic Control*, 2023.

[3] M. Lazar, "Basis-Functions Nonlinear Data-Enabled Predictive Control," in *ECC*, 2024.

Direct vs Indirect Approach for Data-Driven Stability Analysis of Switched Linear Systems

Alexis Vuille
alexis.vuille@uclouvain.be

Guillaume O. Berger
guillaume.berger@uclouvain.be

Raphaël M. Jungers
raphael.jungers@uclouvain.be

1 Introduction

A hot topic in data-driven control is when should we embrace direct data-driven control or indirect data-driven control. We study the two approaches for data-driven stability analysis of switched linear systems. We present these two approaches and discuss their strenghts and weaknesses.

2 Problem Statement

We consider a discrete-time *switched linear system*:

$$x(t+1) \in \{Ax(t)+w(t) : A \in \mathscr{A}\}, \tag{1}$$

wherein $\mathscr{A} = \{A_1,\dots,A_m\} \subseteq \mathbb{R}^{n\times n}$ and $w(t)$ is a noise. The *joint spectral radius* of $\mathscr{A}$, denoted by $\rho(\mathscr{A})$, is the infimum of all $r \geq 0$ for which there exists $C \geq 1$ such that every trajectory x of $\mathscr{A}$ satisfies that for all $t \in \mathbb{N}$, $\|x(t)\| \leq Cr\top\|x(0)\|$. Now, we assume we have a collection of noisy measurements of one-step trajcetories of the system:

$$\mathscr{D} := \{(x_i,y_i)\}_{i=1}^N = \{(x_i(0),x_i(1)+w_i)\}_{i=1}^N,$$

where w_i is the measurement noise. Based only on $\mathscr{D}$, how can we approximate the JSR of the system and what stability guarantees can we obtain?

3 Direct Approach

Inspired from [1] the direct approach we directly approximate the JSR with data by solving a scenario aprogram to approximate a quadratic Lyapnuov function:

$$\begin{array}{ll}\min_{\gamma\geq 0, 0\preceq P=P\top\preceq \delta I} & \gamma, \\ \text{s.t} & y_i^\top P y_i \leq \gamma^2 x_i^\top P x_i \quad \forall i \in [N],\end{array} \tag{2}$$

with optimal solutions $P_\star(\mathscr{D},\delta)$ and $\gamma_\star(\mathscr{D},\delta) \approx \rho(\mathscr{A})$.

Theorem 1. Assume that $\{x_i(0)\}_{i=1}^N$ are sampled independently and uniformly on $\mathbb{S}^{n-1}$. Let $\beta \in (0,1)$ and $W > 0$, then it holds under some mild assumptions with probability at least $1-\beta$ on the sampling of $\mathscr{D}$:

$$\rho(\mathscr{A}) \leq \frac{\gamma_\star(\mathscr{D},\delta)+W\delta}{\xi\left(\frac{m\varepsilon(\beta,N)\kappa(P_\star(\mathscr{D},\delta))}{2(1-p(W))}\right)},$$

where the expressions of $\xi(\cdot)$,$\varepsilon(\beta,N)$, $\kappa(P)$ are omitted here due to space limitations and $p(W)$ is a tail probability bound on the norm of the observation noise.

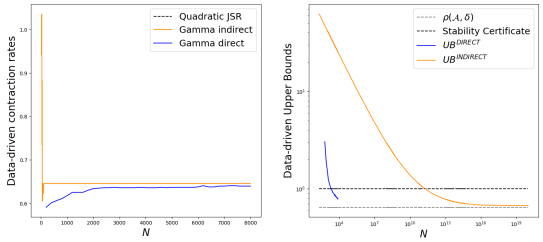

Figure 1: Upper-bounds with confidence $1-\beta = 95\%$ and approximations of the JSR for the consensus network in function of N the total number of samples required.

4 Indirect Approach

In the indirect approach, we first identify a switched linear model $\hat{\mathscr{A}} := \{\hat{A}_1,\dots,\hat{A}_m\}$ using $\mathscr{D}$. Then, we approximate the JSR by solving the following program to compute a quadratic Lyapunov function for the identified model:

$$\begin{array}{ll}\min_{\gamma\geq 0, 0\preceq P=P\top\preceq \delta I} & \gamma, \\ \text{s.t} & \hat{A}_j^\top P \hat{A}_j \leq \gamma^2 P \quad \forall j \in [m],\end{array} \tag{3}$$

with optimal solutions $P_0(\hat{\mathscr{A}},\delta)$ and $\gamma_0(\hat{\mathscr{A}},\delta) \approx \rho(\mathscr{A})$.

Theorem 2. Assume that $\{x_i(0)\}_{i=1}^N$ are sampled independently and uniformly on $\mathbb{S}^{n-1}$. Let $\beta \in (0,1)$, $\varepsilon > 0$ and $W > 0$, then it holds under some mild assumptions with probability at least $1-\beta$ on the sampling of $\mathscr{D}$:

$$\rho(\mathscr{A}) \leq \frac{\gamma_0(\hat{\mathscr{A}},\delta)+(W+\varepsilon)\delta}{\xi\left(\frac{m\varepsilon'(\beta,\hat{\mathscr{A}},\mathscr{D})P_0(\hat{\mathscr{A}},\delta))}{2(1-p(W))}\right)},$$

where the expression of $\varepsilon'(\beta,\hat{\mathscr{A}},\mathscr{D})$ is omitted here due to space limitation.

5 Indirect vs Indirect Approach

We apply the two approaches on a linear switched system of dimension 7 with 3 modes. Fig. 1 shows on the left how well the JSR is approximated with the number of samples used and on the right side how many data points are needed to guarantee stability with confidence $1-\beta = 95\%$. We see that while the indirect approach clearly approximates better the JSR it offers weaker probilistic guarantees.

References

[1] Banse, A., Wang, Z., & Jungers, R. M. (2024). Learning stability guarantees for constrained switching linear systems from noisy observations. Nonlinear Analysis: Hybrid Systems, 51, 101425.3

Frequency-Domain Willems' Fundamental Lemma for Data-Driven Prediction with Noisy Data

Michiel Wind[1,*], Tomas Meijer[2], Victor Dolk[1,3], Maurice Heemels[1]
[1] Eindhoven University of Technology, the Netherlands, [2] Lund University, Sweden
[3] ASML, Veldhoven, *m.wind@tue.nl, date: February 12, 2026

1 Introduction

The frequency-domain alternative to the widely celebrated data-enabled predictive control (DeePC) called frequency-domain data-enabled predictive control (FreePC) has been developed in [1], and is shown to be equivalent to (DeePC) in the deterministic case with steady-state data. In [2] an extension is proposed to deal with non-steady-state data.
Utilizing the frequency-domain formulation has clear advantages for the industry, where decades of classical system identification procedures are adopted to obtain Frequency Response Function (FRF) measurements. With FreePC, predictive control can be adopted by directly using the *non-parametric* FRF measurements *without* the need for a *parametric* state-space model. Furthermore, the data is typically denser in information leading to small datasets, and when periodic excitation is used, non-parametric noise descriptions are easily obtained which gives valuable insight.
Currently, the inclusion of noise in both (DeePC) and (FreePC) proves to be a challenge, since (WFL) only holds in the deterministic case. Applying (WFL) naively would result into bias in the predicted outputs. In [3] it is shown that by adopting instrumental variables (IVs) in (DeePC), the predicted outputs become asymptotically unbiased. For (FD-WFL), a singular value decomposition (SVD) is shown to mitigate the effects of the noise in [2].
In this work, we aim to explore classical system identification techniques, e.g., using (IVs) in (FD-WFL) to obtain (asymptotically) unbiased predictions.

2 FD-WFL predictions

Consider a discrete-time LTI system with an input spectrum $\mathscr{U}_{[0,M-1]}$ that is persistently exciting (PE) of order $L+n_x$, with L prediction length and n_x the state dimension, then the second statement of the frequency-domain version of Willems' fundamental lemma states that the input-output pair $(u_{[0,L-1]}, y_{[0,L-1]})$ is a trajectory of the underlying process, if and only if there exists a $g \in \mathbb{R}^{2M-1}$ such that

$$\begin{bmatrix} u_{[0,L-1]} \\ y_{[0,L-1]} \end{bmatrix} = \begin{bmatrix} \Re F_L(\mathscr{U}_{[0,M-1]}) & \Im F_L(\mathscr{U}_{[1,M-1]}) \\ \Re F_L(\mathscr{Y}_{[0,M-1]}) & \Im F_L(\mathscr{Y}_{[1,M-1]}) \end{bmatrix} g, \quad (1)$$

with $\Re(\cdot)$ and $\Im(\cdot)$ the real and imaginary component of $F_L(V_{[m,n]}) := [W_L(e^{j\hat{\omega}_m}) \otimes V_m \dots W_L(e^{j\hat{\omega}_n}) \otimes V_n]$, and $W_L(z) := [1\ z\ \dots\ z^{L-1}]^\top$, see [1] for more details.

3 DeePC with instrumental variables

In [3], the specific choice of instrumental variable $Z = \begin{bmatrix} U_p & Y_p & U_f \end{bmatrix}^\top$, with U_p and Y_p off-line past input and outputs Hankel matrices, and off-line future input U_f Hankel matrix, ensures correlation with the data matrix, while being uncorrelated with Hankel matrix containing the noise E_f, i.e., asymptotically

$$\lim_{N\to\infty} \mathbb{E}[E_f Z^\top] = 0, \quad (2)$$

with N the total data length. By defining $\hat{g}$ through the linear system

$$ZZ^\top \hat{g} = \begin{bmatrix} \hat{U}_p & \hat{Y}_p & \hat{U}_f \end{bmatrix}^\top, \quad (3)$$

we obtain the vector $\hat{g}$ that is used for predictions of the future output given by

$$\hat{Y}_f = Y_f Z^\top (ZZ^\top)^{-1} \begin{bmatrix} \hat{U}_p & \hat{Y}_p & \hat{U}_f \end{bmatrix}^\top, \quad (4)$$

with $(\hat{U}_p, \hat{Y}_p, \hat{U}_f, \hat{Y}_f)$, the estimated input-output trajectories which are asymptotically unbiased.

4 Open Issue

The choice of Z must comply with the 2 conditions to be an instrumental variable for frequency-domain data where the role of $N \to \infty$ changes with respect to time-domain data, while the matrices that appear in the data-equation in the frequency-domain do not have the Hankel structure.
In [3] it is assumed that there exists a finite p such that $||\tilde{A}^p||_F \approx 0$ and that $\tilde{A}$ is stable. In the frequency-domain, we may disregard both of these assumptions by leveraging non-steady-state data as shown in [2] and using closed-loop measurements for unstable plants as shown in [1].

References

[1] T. J. Meijer et al., *From a Frequency-Domain Willems' Lemma to Data-Driven Predictive Control*, arXiv:2501.19390, 2025.

[2] T. J. Meijer et al., *Leveraging Non-Steady-State Frequency-Domain Data in Willems' Fundamental Lemma*, arXiv:2504.06403, 2025.

[3] J. van Wingerden et al., *Data-Enabled Predictive Control with Instrumental Variables: The Direct Equivalence with Subspace Predictive Control*, arXiv:2209.05210, 2022.

Interpolation Conditions for Data Consistency and Prediction in Noisy Linear Systems

Martina Vanelli
UCLouvain, Belgium
martina.vanelli@uclouvain.be

Nima Monshizadeh
University of Groningen
n.monshizadeh@rug.nl

Julien M. Hendrickx
UCLouvain, Belgium
julien.hendrickx@uclouvain.be

1 Introduction and model description

Data-driven control has become a crucial aspect of modern control theory, offering powerful tools for system analysis and design. Many existing approaches rely on the assumption of persistence of excitation, which, in the absence of noise, ensures that the underlying dynamical system can be uniquely identified from the data. In practice, data are often affected by noise, preventing unique identification of the system parameters. Furthermore, the persistence of excitation condition is not always satisfied or even necessary.

In this work [1], we propose a general framework that exploits available structure and prior knowledge to extract the maximum possible information from available data, regardless of its quantity or quality. In this initial study, we consider noisy linear time-invariant

$$x_{k+1} = Ax_k + Bu_k + v_k\,, \tag{1}$$

where $x_k \in \mathbb{R}_n$ is the state vector, $u_k \in \mathbb{R}_m$ is the input vector, and v_k is the process noise. We assume that the input matrix B is known, while the state-transition matrix A is unknown. We also assume to have access to input-state measurements, that is, $x_0,\dots,x_t$ and $u_0,\dots,u_{t-1}$ for some $t > 0$. Importantly, we make no assumption on the number or richness of the data and the number of samples t may even be smaller than the system dimension n. We assume a bound on the system norm, which corresponds to a bound on energy increase. In particular, when this bound equals one, the system exhibits non-decreasing energy and thus is stable. Additionally, the process noise is assumed to have a bounded energy.

Our goal is to address data-driven questions such as determining whether the observed data are compatible with given noise and energy bounds, identifying which parameter combinations explain the trajectories, and characterizing the set of feasible future states and optimizing over it. Unlike the set membership estimation which seeks all models consistent with data and bounded uncertainty, our aim is to provide a systematic way to address such queries directly from data without explicitly identifying all admissible system matrices or noise realizations. To achieve this goal, we rely on *interpolation conditions*, that is, necessary and sufficient conditions on a set of points that guarantee the existence of an interpolating function (or operator) belonging to a certain class (e.g. L-smooth convex functions). The use of interpolation conditions has led to a novel type of analysis in optimization, enabling to derive exact worst-case performances of algorithms.This work is a preliminary step in the direction of exploiting interpolation conditions in data-driven control to characterize the set of trajectories that are consistent with a dynamical system in a given class.

2 Main result

Our main contributions are the following. First, we derive suitable interpolation conditions for our setting. By means of a matrix-completion argument, we extend a classical algebraic result for bounded linear operators to the case of linear operators composed of multiple bounded components. Then, we exploit its application for LTI systems. Specifically, using the generalized interpolation conditions, we formulate semidefinite programs (SDP) that verify whether a dataset is compatible with assumed noise and energy amplification (or decrease) bounds and that compute minimal admissible bounds. We apply the algebraic result to characterize the set of all next states consistent with past data and prior assumptions. This characterization can be used for safety verification and worst-case (energy) minimization. Since our formulation is not an SDP in general, we derive an explicit description of the feasible set: in the noise-free case it reduces to a single ellipsoid determined by the data and prior bounds; with process noise it becomes a union of ellipsoids, whose parameters we explicitly derive. An important advantage of our preliminary results is that they do not require any assumptions on the amount of data. Indeed, our aim is to exploit the maximal information we can obtain from any amount of data and investigate how this information can be combined with the physical knowledge of the system to perform effective control.

3 Acknowledgements

This work was supported by ARC via the Siddarta project.

References

[1] M. Vanelli, N. Monshizadeh and J. M. Hendrickx, "Interpolation Conditions for Data Consistency and Prediction in Noisy Linear Systems", https://arxiv.org/abs/2504.08484, 2025

Contraction estimator: a global optimal solution for robotics optimization problems

Xinghua Liu
Engineering and Technology Institute Groningen
University of Groningen
P.O. Box 72, 9700 AB Groningen
The Netherlands
Email: xinghua.liu@rug.nl

Ming Cao
Engineering and Technology Institute Groningen
University of Groningen
P.O. Box 72, 9700 AB Groningen
The Netherlands
Email: m.cao@rug.nl

1 Introduction

The success of intelligent robots in dealing with practical tasks relies on reliable information fusion technologies, whose goal is to solve optimization problems constructed using data from multiple sensors [1]. Optimization problems in robotics are sometimes non-convex, non-smooth, discrete, and not even formulated in the Euclidean space. We propose a novel optimization method, termed the contraction estimator, to address these challenges of optimization problems. We start from the original contraction theory, the Banach Contraction Principle, and then propose a novel framework to solve the optimization problems by leveraging contraction theories and a contraction operator. The contraction operator approximates a generating function to search for lower function values directly in the range of the objective function, rather than updating the variable directly using the gradient of the objective function and assuming that leads to a lower function value.

2 Contraction estimator

2.1 Contraction theory

Proposition 1 *(Banach Fixed Point Theorem [2]) Let (X,d) be a complete metric space and W be a contraction operator on X with the Lipschitz constant $\alpha \in [0,1)$ such that $d(W(x),W(y)) \leq \alpha d(x,y)$ for all $x,y \in X$. Then W has a unique fixed point x^* in X; for an arbitrary point $x_0 \in X$, the sequece $(x_t)_{t\in\mathbb{N}}$ generated by the Picard iterarion process(defined by $x_{t+1} = W(x_t), k \in \mathbb{N}\cup\{0\}$) converges to x^*; and $d(x_t,x^*) \leq \frac{\alpha^t}{1-\alpha} d(x_0,x_1)$ for all $t \in \mathbb{N}$.*

2.2 Contraction estimator

For the optimization problem $\min_x f(r(x))$ with the residual function $r(x_t)$, we design the following expression

$$r(x_{t+1}) = (I+S)^{-1} r(x_t) = \sum_{i=0}^{\infty} (-1)^i S^i r(x_t), \rho(S) < 1, \quad (1)$$

where $S = AA^T(AA^T + \mu I)^{-1}$. This construction is analogous to the Taylor expansion, which can be regarded as a special case of a generating function. Consider the following contraction operator

$$W = I - AQ = I - AA^T \left(AA^T + \mu I\right)^{-1}, \quad (2)$$

where $Q = A^T \left(AA^T + \mu I\right)^{-1}$, A is any operator on X, $\mu > 0$, and I is the identity operator with a proper dimension. For a given optimization problem, the operator $A \in \mathbb{R}^{m\times n}$ connecting the residual error space and the variable space in robotic problems, and $n = dim(M)$ is the dimension of M. One can design a contraction estimator with the operator W. Take $A = J_t$ at the time instant t, a new operator constructed from the objective function at its current estimation x_t, and the corresponding preprocessing operator Q becomes $Q_t = J_t^T \left(J_t J_t^T + \mu I\right)^{-1}$. A potential selection of A is the Jacobian of the residual function $J_t = \frac{\partial r(x)}{\partial x}\Big|_{x_t} \in \mathbb{R}^{m\times n}$ and then the iterative estimation is given by

$$x_{t+1} = x_t \ominus A Q_t r_t, \quad (3)$$

where the operation $\ominus$ represents algebraic subtraction in the Euclidean space and can be a group action on Lie groups.

3 Conclusion

In this work, we develop an estimator based on the contraction of the optimization objective. The proposed contraction estimator guarantees that the estimation is globally convergent. The contraction property of the proposed method originates from a contraction operator. This contraction operator is derived from the approximation of the Taylor expansion of the optimization objective. The analysis of the proposed methods highlights the broad applicability of our proposed framework.

References

[1] Fei Zhao, Chengcui Zhang, and Baocheng Geng. Deep multimodal data fusion. *ACM computing surveys*, 56(9):1–36, 2024.

[2] Stefan Banach. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta mathematicae*, 3(1):133–181, 1922.

A Two-point Random Gradient-free Method for Model-free Feedback Optimization

Amir Mehrnoosh
ICTEAM institute, UCLouvain
amir.mehrnoosh@uclouvain.be

Gianluca Bianchin
ICTEAM institute, UCLouvain
gianluca.bianchin@uclouvain.be

1 Abstract

Feedback optimization (FO) [1] has been established as a powerful paradigm for the optimal steady-state operation of dynamical systems. By integrating optimization algorithms directly into the control loop, FO allows systems to steer toward optimal operating points. However, a significant limitation of traditional FO methods is their dependency on explicit knowledge of the plant model or gradients. In many applications, such as power electronic converters, energy markets, or industrial processes, accurate models are unavailable due to switching nonlinearities, unknown disturbances, or parameter drift. Moreover, the performance objectives (e.g., thermal stress, user satisfaction) are often only accessible via real-time measurements, whereby analytical gradient computation becomes impossible.

To overcome this, recent works have explored *model-free* approaches. Notably, one-point residual feedback methods [2] have been proposed; however, they suffer from slower convergence rates compared to first-order model-based methods. In this work, we address this limitation by proposing a novel two-point random gradient-free method for feedback optimization. Our approach adepts zeroth-order optimization techniques [3] to dynamical systems that leads to efficient model-free operation.

This work features three main contributions. *First*, we introduce a randomized algorithm that estimates descent directions using two function evaluations. Unlike static optimization, applying two-point methods to dynamical systems is non-trivial because the system state evolves between measurements. Our algorithm overcomes this by alternating between randomized *exploration steps* (perturbing inputs to probe the system and estimate the gradients) and *update steps* (adjusting controls based on estimated gradients). This design allows the controller to operate using only real-time performance measurements, without requiring state resets or model knowledge.

Second, we provide a rigorous convergence analysis. We prove that for smooth, nonconvex cost functions and the proper choice of parameters, our method drives the system to an ε-neighborhood of a stationary point. Crucially, we show that our method achieves a convergence rate of $\mathcal{O}(p\varepsilon^{-1})$, where p is the dimension of the input. This result significantly outperforms existing one-point feedback optimization methods (which typically scale as $\mathcal{O}(\varepsilon^{-2})$ or $\mathcal{O}(\varepsilon^{-1.5})$) and recovers the theoretical efficiency of two-point methods used in static optimization.

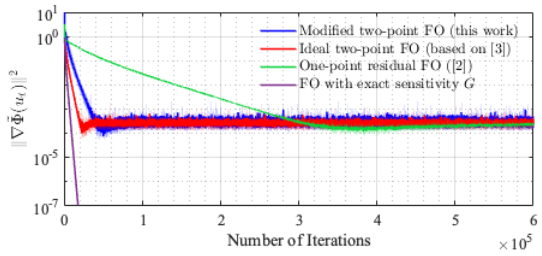

Figure 1: Comparison of convergence rates. The proposed two-point method (blue) outperforms the one-point residual method (green) and approaches the performance of the idealized baseline (red).

Third, we provide explicit guidelines for parameter selection, meaning the conditions for the stepsize η and smoothing parameter δ that guarantee stability and convergence.

Figure 1 illustrates the performance of the proposed method through numerical simulations. The plot compares the squared norm of the gradient (a measure of optimality) over time. Our method strikes a favorable balance between the fast but unimplementable "idealized" two-point baseline and the slower one-point residual feedback methods.

Future research directions include extending this framework to constrained optimization problems and exploring adaptive stepsize methods to further enhance transient performance.

References

[1] M. Colombino, E. Dall'Anese, and A. Bernstein, "Online optimization as a feedback controller: Stability and tracking," vol. 7, no. 1, pp. 422–432, 2020.

[2] Z. He, S. Bolognani, J. He, F. Dörfler and X. Guan, "Model-Free Nonlinear Feedback Optimization," in *IEEE Transactions on Automatic Control*, vol. 69, no. 7, pp. 4554-4569, July 2024, doi: 10.1109/TAC.2023.3341752.

[3] Y. Nesterov and V. Spokoiny, "Random gradient-free minimization of convex functions." *Foundations of Computational Mathematics* 17, no. 2 (2017): 527-566.

Optimization Models for Hydrogen-based Hybrid Electric Aero Engines[1]

Uto Perra, Faezeh Pak, Theo Hofman, Frank Willems, Mauro Salazar
Control Systems Technology
Eindhoven University of Technology (TU/e)
Email: {u.perra, f.pak, t.hofman, f.p.t.willems, m.r.u.salazar}@tue.nl

Abstract

Due to the significant increase in aerospace carbon-related emissions, several hybrid propulsion systems have been studied. Among the various proposed solutions, hybrid systems composed of conventional gas turbine (GT) engines and solid oxide fuel cells (SOFC) appear to be a promising configuration. Such systems have been extensively studied in recent years for stationary applications, and more recently, they have also been explored for aerospace applications. Figure 1 presents a schematic of a hybrid propulsion system composed of a conventional gas turbine, solid oxide fuel cell (SOFC), electric motor (EM), gearbox (GB), and propeller. To fully exploit the benefits of a hybrid configuration, it is of utmost importance to apply the proper allocation of power among the various propulsion sources within the vehicle.

This paper pertains to two research streams: Optimal control of hybrid automotive powertrains, and models of (hybrid) aero engines. In the first stream, given the progressive development of hybrid technology over the years, power management between drive systems within the vehicle is well-known. Sundström et al. [1] and Robuschi et al. [2] presented models and optimization algorithms to compute the minimum fuel power management strategies for given driving cycles. Although the power management strategies discussed here are largely derived from automotive applications, there is a notable gap in their direct application to aerospace. Specifically, in hybrid aviation systems, the challenge comes from the strong thermodynamic connection between the power sources and their complex characteristics. In the second stream, Doff-Sotta et al. [3] proposed a convex formulation for optimal energy management for aircrafts with hybrid propulsion composed of gas turbine and electric motor components. However, the study is based on a given power profile demand considering a constant mass throughout the duration of the flight. In addition, it is applied to hybrid battery-powered powertrains. To conclude, to the best of the authors' knowledge, no studies have explored the use of convex models to determine the optimal operation of hybrid GT-SOFC propulsion systems.

This work presents convex models for a hybrid GT-SOFC integrated aircraft propulsion system, accounting for the strong thermodynamic couplings between the subsystems. Next, a minimum fuel operation problem is formulated and solved by leveraging the aforementioned models.

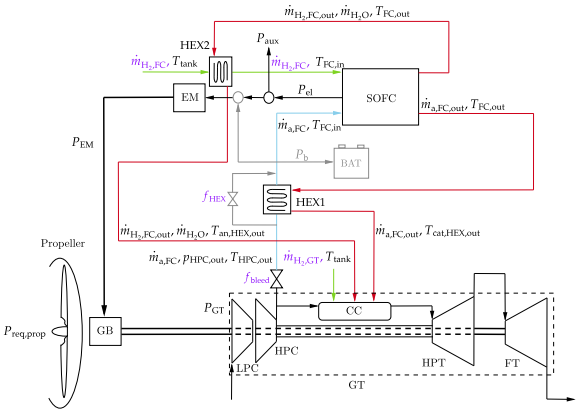

Figure 1: Scheme of the examined hybrid propulsion that includes: the compressor (C), the combustion chamber (CC), the high pressure turbine (HPT), the free turbine (FT), the propeller, the gearbox (GB), the SOFC and the electric motor (EM). The bold black lines correspond to the mechanical connections, while the grey lines to different analyzed configurations.

[1]Research funded by the European Union (ID 101138488) and by the UK Research and Innovation (ID 10106893). Views expressed are those of the author(s) only and do not necessarily reflect those of the funding entities. Neither the European Union nor the granting authority can be held responsible for them.

References

[1] O. Sundström, L. Guzzella, and P. Soltic, "Optimal hybridization in two parallel hybrid electric vehicles using dynamic programming," in *IFAC World Congress*, vol. 41. IFAC, 2008, pp. 4642–4647.

[2] N. Robuschi, M. Salazar, N. Viscera, F. Braghin, and C. H. Onder, "Minimum-fuel energy management of a hybrid electric vehicle via iterative linear programming," *IEEE Trans. on Vehicular Tech.* , 2020.

[3] M. Doff-Sotta, M. Cannon, and M. Bacic, "Optimal energy management for hybrid electric aircraft," *IFAC-Papers Online*, 2020.

Bandwidth Trade-offs in High-Performance, Safe Sensorless Actuators

Léon Borremans
Federated Labs AI and Robotics (FLAIR), Vrije Universiteit Brussel (VUB), Belgium
Leon.Borremans@vub.be

1 Abstract

Future robotic systems increasingly rely on **transparent actuators**, as they are intrinsically **safer** and can operate **without dedicated torque sensors** by exploiting motor-side measurements for torque estimation [1]. Typical examples include quasi-direct-drive actuators and intrinsically backdrivable transmissions. At the same time, robots are expected to execute tasks rapidly and efficiently, which requires high dynamic performance actuators [2].

These requirements give rise to **two essential but competing objectives**. Fast task execution requires a high bandwidth transfer from motor torque to motor response [3], whereas safe physical human–robot interaction requires a high bandwidth transfer from load torque to motor response to reliably detect external interactions and impacts [2]. This creates a fundamental trade-off between an actuator's **dynamic performance and interaction safety**.

This work presents a **mathematical framework** that formally **characterizes** the **bandwidth of forward and backward drive transfers** and clarifies how key actuator parameters—such as gear ratio and motor inertia—govern this trade-off. *Forward drive transfer* is defined as the transfer function from motor torque to motor acceleration, while *Backward drive transfer* is defined as the transfer function from load torque to motor acceleration. The framework is intended to support informed actuator design, including motor and gearbox selection for transparent, sensorless robotic systems. The bandwidth of these transfer functions provides an intuitive and quantitative metric for **dynamic performance** and **interaction safety**, **respectively**.

Assuming a given load inertia and stiffness, the **framework evaluates how gear ratio and motor inertia influence the bandwidth of both metrics**. To capture this relationship in a scalable manner, the inertia ratio is introduced as $\alpha^2 = J_L/J_M$, the ratio between motor and load inertia. The results are summarized in **Figure 1** and reveal two trends. The bandwidth of forward drive increases with higher gear ratios and lower inertia ratios, whereas the bandwidth of backward drive increases with lower gear ratios and higher inertia ratios. Moreover, practical actuator design introduces an additional constraint: higher gear ratios typically require smaller motors, leading to a **positive correlation between gear ratio and inertia ratio**. **Consequently**, improving forward drive inherently degrades backward drive, and vice versa, making the **simultaneous maximization of both metrics impossible in practice**, thus **necessitating** an **optimization framework**.

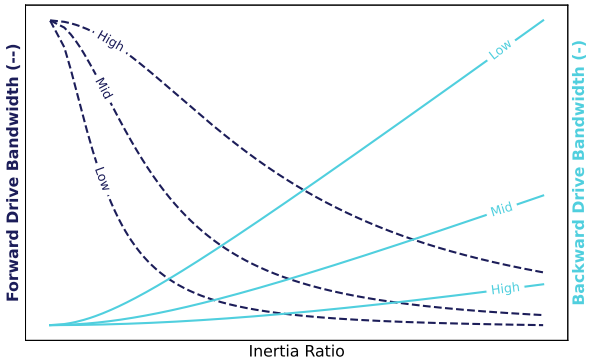

Figure 1: Bandwidth of forward drive (left) and backward drive (right) for low, moderate and high gear ratio.

Future work will build upon these insights by developing a multi-objective optimization framework for robotic actuator design. In addition to forward and backward drive bandwidth, complementary metrics—such as unit-torque load-side acceleration—will be incorporated. The **future framework will provide quantitative design guidelines for parameters such as gear ratio and motor inertia**, offering a systematic first step toward the design of highly dynamic, safe, and sensorless robotic actuators.

References

[1] P. M. Wensing, A. Wang, S. Seok, D. Otten, J. Lang, and S. Kim, "Proprioceptive actuator design in the MIT Cheetah: Impact mitigation and high-bandwidth physical interaction for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 509–522, Jun. 2017.

[2] S. Haddadin, A. De Luca, and A. Albu-Schäffer, "Robot collisions: A survey on detection, isolation, and identification," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.

[3] F. Roos and J. Wikander, "The influence of gear ratio on performance of electromechanical servo systems," *IFAC Proceedings Volumes*, vol. 39, no. 16, pp. 884–889, 2006.

Inertial Primal Dual Dynamics for Smooth Saddle Point Problems with Bilinear Coupling

Zepeng Wang and J.G. Peypouquet
Bernoulli Institute, University of Groningen
Email: zepeng.wang@rug.nl, j.g.peypouquet@rug.nl

1 Introduction

Consider the following smooth saddle point problem with bilinear coupling, which is given by

$$\min_{x\in\mathbb{R}^n}\max_{y\in\mathbb{R}^m}\mathscr{L}(x,y)=f(x)+\langle Ax,y\rangle-g(y), \qquad (1)$$

where $f:\mathbb{R}^n\to\mathbb{R}$ and $g:\mathbb{R}^m\to\mathbb{R}$ are continuously differentiable convex functions, and $A\in\mathbb{R}^{m\times n}$ is a coupling matrix. This problem can be solved by primal dual algorithms [1, 2], or their continuous time counterparts, primal dual dynamics [3, 4]. In this study, we limit our attention to the latter approach, and focus on those inertial primal dual dynamical systems, which are inspired by Nesterov's accelerated gradient method.

There have been some inertial primal dual systems proposed to solve (1), but these systems do *not* contain full Hessian-driven damping, and as a consequence, may suffer from trajectory oscillations. To alleviate the oscillations, we develop two inertial primal dual dynamical systems with Hessian-driven damping. We establish a fast convergence rate for the primal dual gap. We also extend our results to the affinely constrained convex optimization problem, which is a special case of (1).

2 Inertial primal dual dynamics

The proposed inertial primal dual dynamical systems are as follows.

(i) When f and g are convex, the system is given by

$$\begin{cases}\ddot{x}+\frac{\alpha}{t}\dot{x}+\beta_f\nabla_{xx}\mathscr{L}(x,y)\dot{x}+\left(\gamma+\frac{r}{t}\right)\nabla_x\mathscr{L}(x,y+\theta t\dot{y}) &= 0,\\ \ddot{y}+\frac{\alpha}{t}\dot{y}-\beta_g\nabla_{yy}\mathscr{L}(x,y)\dot{y}-\left(\gamma+\frac{r}{t}\right)\nabla_y\mathscr{L}(x+\theta t\dot{x},y) &= 0,\end{cases} \qquad \text{(PD-AVD-H)}$$

where $\alpha\geq 3$ and $\beta_f,\beta_g,\gamma,r,\theta>0$.

(ii) When f and g are strongly convex with parameters $\mu_f,\mu_g>0$, the system is given by

$$\begin{cases}\ddot{x}+\alpha\dot{x}+\beta_f\nabla_{xx}\mathscr{L}(x,y)\dot{x}+\gamma\nabla_x\mathscr{L}(x,y+\theta\dot{y}) &= 0,\\ \ddot{y}+\alpha\dot{y}-\beta_g\nabla_{yy}\mathscr{L}(x,y)\dot{y}-\gamma\nabla_y\mathscr{L}(x+\theta\dot{x},y) &= 0,\end{cases} \qquad \text{(PD-HBF-H)}$$

where $\alpha,\beta_f,\beta_g,\gamma,\theta>0$.

3 Main results

Let (x^*,y^*) be the saddle point of $\mathscr{L}(x,y)$ in (1). When f and g are convex, we have the following results for system (PD-AVD-H):

Theorem 1 *Let $f:\mathbb{R}^n\to\mathbb{R}$ and $g:\mathbb{R}^m\to\mathbb{R}$ be convex. Let $(x,y):[t_0,\infty)\to\mathbb{R}^n\times\mathbb{R}^m$ be a solution of the system* (PD-AVD-H)*, where $\alpha\geq 3$, $\beta_f,\beta_g,\gamma>0$, $r\geq\max(\beta_f,\beta_g)$ and $\theta=\frac{1}{\alpha-1}$. Then, we have, for every $t\geq t_0$,*

$$\mathscr{L}(x(t),y^*)-\mathscr{L}(x^*,y(t))\leq O\left(\frac{1}{\gamma t^2}\right).$$

When f and g are strongly convex, we have the following results for system (PD-HBF-H):

Theorem 2 *Let $f:\mathbb{R}^n\to\mathbb{R}$ be μ_f-strongly convex and $g:\mathbb{R}^m\to\mathbb{R}$ be μ_g-strongly convex. Let $(x,y):[t_0,\infty)\to\mathbb{R}^n\times\mathbb{R}^m$ be a solution of the system* (PD-HBF-H)*, where $\alpha\leq 2\sqrt{\min(\mu_f,\mu_g)\gamma_0}$, $\min(\beta_f,\beta_g)>0$, $\gamma=\gamma_0+\frac{\alpha}{2}\max(\beta_f,\beta_g)$ and $\theta=\frac{2}{\alpha}$ with $\gamma_0>0$. Then, we have, for every $t\geq t_0$,*

$$\mathscr{L}(x(t),y^*)-\mathscr{L}(x^*,y(t))\leq O\left(e^{-\frac{1}{2}\alpha t}\right).$$

References

[1] A. Chambolle, T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," Journal of Mathematical Imaging and Vision, 40 (2011) 120-145.

[2] K. Ding, J. Fliege, P.T. Vuong, "Fast convergence of the primal-dual dynamical system and corresponding algorithms for a nonsmooth bilinearly coupled saddle point problem," Computational Optimization and Applications, 90 (2025) 151-192.

[3] X. Zeng, L. Dou, J. Chen, "Accelerated first-order continuous-time algorithm for solving convex-concave bilinear saddle point problem," IFAC-PapersOnLine, 53 (2020) 7362-7367.

[4] X. He, Y.P. Fang, "Nesterov acceleration for strongly convex-strongly concave bilinear saddle point problems: discrete and continuous-time approaches," arXiv:2509.08258, (2025).

Fast Optimal Control for Implicitly Constrained Dynamics

Louis Callens, Jan Swevers, Wilm Decré
MECO Research Team, Department of Mechanical Engineering, KU Leuven
Flanders Make @ KU Leuven, Leuven, Belgium
louis.callens@kuleuven.be

1 Introduction

Optimal control has been proven effective in many applications such as robotics, motion planning or process control. Direct methods that transcribe the Optimal Control Problem (OCP) into a Nonlinear Program (NLP) are commonly used to solve these OCPs. In the general case where both equality and inequality constraints are present, interior-point NLP-solvers such as IPOPT have been proven robust. In these methods, a relaxed version of the KKT-system is solved using newton iterations. Every such iteration, requires solving a linear system. However, NLPs resulting from transcribing OCPs contain some problem structure. More specifically, the coefficient matrix of the linear system has a predictable structure which can be exploited to more efficiently factorize the matrix.
FATROP is an interior-point NLP-solver that exploits this structure [2]. It uses a generalization of the Riccati recursion that explicitly takes into account equality constraints and hence outperforms IPOPT [1]. However, the FATROP algorithm does not support implicit integration schemes, which are known to be numerically more stable for stiff dynamics. In this work, we propose a generalization of the recursion algorithm used in FATROP to make it applicable to a broader class of problems with more general dynamics constraints.

2 Methodology

FATROP expects gap-closing constraints of the form

$$f_{\text{expl}}(x_k, u_k) - x_{k+1} = 0 \tag{1}$$

whereas a more general formulation would be

$$f_{\text{impl}}(x_k, u_k, x_{k+1}) = 0. \tag{2}$$

However, this constraint changes the structure of the coefficient matrix of the primal-dual system to be solved each iteration which requires a modified Ricatti recursion. Note that (2) is equivalent to imposing the constraints

$$z_k - x_{k+1} = 0 \tag{3}$$
$$f_{\text{impl}}(x_k, u_k, z_k) = 0 \tag{4}$$

where z_k are newly introduced controls, (3) represent explicit dynamics and (4) is now a stage-wise constraint, fitting the structure required by FATROP. However, this reformulation adds one additional constraint and control for each state variable, which can significantly increase problem size and hence the computation time.
We propose a modified version of the Riccati recursion that explicitly takes into account the structure arising from using (2).

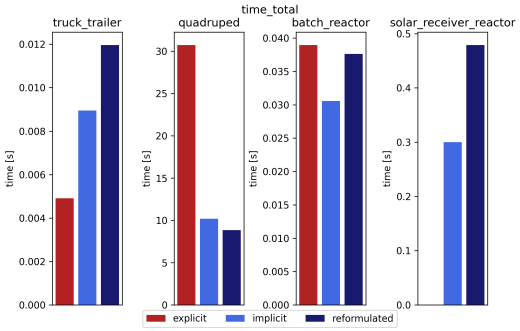

Figure 1: Comparison of total time required to solve example OCPs using an explicit integrator, implicit modification of the Riccati recursion and reformulated dynamics approach (see (3)-(4)).

3 Results

Preliminary results are shown in Figure 1, which shows computation times required by FATROP to solve some example OCPs. The dynamics in the truck-trailer example are not stiff and hence the explicit integrator outperforms the implicit integrators. However, for the other (stiff) problems, the explicit integrator takes much longer or does not converge at all. Additionally, reformulating the dynamics appears to increase computation time compared to the modified version of the Ricatti recursion.

Acknowledgment This work has been carried out within the framework of Flanders Make SBO project LearnOptra. Flanders Make is the Flemish strategic research center for the manufacturing industry.

References

[1] Lander Vanroye, Joris De Schutter, and Wilm Decré. "A generalization of the Riccati recursion for equality-constrained linear quadratic optimal control". In: *Optimal Control Applications and Methods* 45.1 (2024), pp. 436–454. DOI: https://doi.org/10.1002/oca.3064.

[2] Lander Vanroye et al. "FATROP: A Fast Constrained Optimal Control Problem Solver for Robot Trajectory Optimization and Control". In: Oct. 2023, pp. 10036–10043. DOI: 10.1109/IROS55552.2023.10342336.

Numerical Design of Online Optimization Algorithms

Erwan Meunier
ICTEAM, UCLouvain
erwan.meunier@uclouvain.be

Julien M. Hendrickx
ICTEAM, UCLouvain
julien.hendrickx@uclouvain.be

This abstract presents a line of research on the numerical Design of Online Optimization algorithms using the Performance Estimation Problem framework.

1 Online Optimization

Online Convex Optimization (OCO) is a prominent area of research where the objective is to minimize functions f_t (usually L-Lipschitz) received sequentially over a finite-time horizon $[T]$, where the case $f_t = f, \forall t \in [T]$ corresponds to the classical offline setting. Facing a different objective function at each time-step can be useful to capture some dynamic problems where the optimal solution of the problem (e.g. the physic of the problem) is not fixed once and for all. This kind of problems arises in many applications such as in Control or in Machine Learning [2] to mention a few.

A paradigmatic First-Order Method (FOM) to solve online optimization is the so-called *Online Gradient Descent* (OGD) which updates the current minimizer estimate x_t by doing an *online* gradient step (i.e. the subgradient $\partial f_t(x_t)$ is indexed by the time) controlled by a step-size strategy $\eta_t \in \mathbb{R}^T$, and projecting back onto the convex bounded feasible set $\mathscr{X}$:

$$x_{t+1} \leftarrow P_{\mathscr{X}}(x_t - \eta_t \partial f_t(x_t)).$$

The main Performance Metric for OCO algorithms is the so-called (*Static*) *Regret*:

$$Regret(T) := \sum_{t=1}^{T} f_t(x_t) - \min_{x^* \in \mathscr{X}} \sum_{t=1}^{T} f_t(x^*)$$

which compares the total cost incurred at each time-step in the Online scenario to the total cost incurred if the sequence of functions was known in advance.

The classical aim is to reach a sublinear regret and for a proper choice of step-sizes, one can show that the regret for OGD is indeed $O(\sqrt{T})$ [2]. The degree of conservatism of analytical bounds has been recently addressed by computer-aided methods in [3], [6] and [5].

2 Design via Performance Estimation Problems

To obtain an accurate performance analysis of online optimization methods, we exploit the *Performance Estimation Problem (PEP)* framework [1], which allows computing exact worst-case bound of first-order optimization algorithms. In principle, a tight performance bound on an algorithm could be obtained by running it on every single instance (i.e. functions and initial conditions) allowed by the setting considered and selecting the worst performance obtained.

The PEP formulates this abstract idea as a real optimization problem (an SDP in most cases) that maximizes the error measure of the algorithm result, over all possible functions and initial conditions allowed, e.g.:

$$w(\eta_t) := \sup_{x_1, f_t(x_t), \partial f_t(x_t)} Regret(T)$$

$$\text{s.t} \begin{cases} \texttt{Class of functions} \\ x_t \sim \texttt{Method w.r.t } \eta_t \end{cases}$$

Our **goal** is to improve OCO algorithms (e.g. OGD) by minimizing the worst-case performance w.r.t. η_t, i.e.:

$$\min_{\eta_t \in \mathbb{R}^T} w(\eta_t)$$

by using the method introduced in [4] aimed at improving (offline) FOMs. In the latter, they use the first-order information $\nabla_{\eta_t} w(\eta_t)$ of the inner PEP problem, to perform a gradient descent over the step-sizes policy η_t.

3 Contributions

1. We have incorporated the second-order information $\nabla^2 w$ in the approach of [4] by relying on a trust-region method.

2. We demonstrate that, in the online setting, the sparsity of $\nabla^2 w$ can be leveraged to get faster convergence while maintaining a reasonable update complexity.

3. We reduce the dimension of the search space by constraining the class of step-sizes (e.g. finding $a, b, c \in \mathbb{R}$ giving the best performance for the scheme $\eta_t = a(t^b + c)^{-1}$) while maintaining a good improvement of the method.

Acknowledgment E. Meunier is supported by the French Community of Belgium through a FRIA fellowship (F.R.S.-FNRS). J. M. Hendrickx is supported by the "SIDDARTA" Concerted Research Action (ARC) granted by the UCLouvain.

References

[1] A. B. Taylor, J. M. Hendrickx, and F. Glineur, "Smooth strongly convex interpolation and exact worst-case performance of first-order methods," *Mathematical Programming*, vol. 161, no. 1, pp. 307–345, Jan. 2017.

[2] S. C. Hoi, D. Sahoo, J. Lu, and P. Zhao, "Online learning: A comprehensive survey," *Neurocomputing*, vol. 459, pp. 249–289, 2021.

[3] F. Jakob and A. Iannelli, *Online convex optimization and integral quadratic constraints: An automated approach to regret analysis*, 2025.

[4] Y. Kamri, J. M. Hendrickx, and F. Glineur, *Numerical design of optimized first-order algorithms*, 2025.

[5] E. Meunier and J. M. Hendrickx, *Several performance bounds on decentralized online optimization are highly conservative and potentially misleading*, 2025.

[6] J. Weibel, P. Gaillard, W. M. Koolen, and A. Taylor, *Optimized projection-free algorithms for online learning: Construction and worst-case analysis*, 2025.

Loop-Shaping for Narrow-Band Active Damping Control

Aditya Natu, Hassan HosseinNia
Precision and Microsystems Engineering, Delft University of Technology, The Netherlands
Email: a.m.natu@tudelft.nl, s.h.hosseinniakani@tudelft.nl

1 Research Contribution

Lightly damped structural resonances limit the performance of high-precision mechatronic systems and flexible structures. Conventional active damping schemes, such as positive and negative position feedback, provide moderate damping at the targeted mode, but often suffer from spillover at low and high frequencies, especially in systems with closely spaced modes or high noise sensitivity. This work presents a generalized higher-order band-pass active damping framework that enhances modal selectivity and stability for narrow-band damping. A non-minimum-phase filter is incorporated to compensate for delay-induced phase degradation, while a tunable damping coefficient increases the design flexibility. Simulation studies on identified collocated and non-collocated systems demonstrate that the proposed controller achieves sharp resonance suppression, improved noise attenuation, and reduced spillover compared with conventional active damping schemes.

2 Loop Shaping via Band-Pass Controller

For an active damping controller $C_d(s)$ that is implemented within a negative feedback loop, the open-loop transfer function $\mathcal{L}(s)$ is defined as $\mathcal{L}(s) = G(s)C_d(s)$. The closed-loop transfer function $\mathcal{PS}_{xd}(s)$, which maps the input disturbance d to the performance output x, is given by:

$$\mathcal{PS}_{xd}(s) = \frac{G(s)}{1 + G(s)C_d(s)} = G(s)\left(\mathcal{S}_{xn}(s) + 1\right). \quad (1)$$

To achieve the desired closed-loop resonance peak attenuation at the system's natural frequency ω_n of $\mathcal{PS}_{xd}(s)$ and suppression of noise propagation to the performance output in $\mathcal{S}_{xn}(s)$, the open-loop transfer function $\mathcal{L}(s)$ must satisfy the following gain conditions:

$$\begin{array}{ll|ll} |\mathcal{L}(s)| \gg 1 & | & |\mathcal{PS}_{xd}(s)| \ll |G(s)| & \text{at } \omega \approx \omega_n, \\ |\mathcal{L}(s)| \ll 1 & | & |\mathcal{S}_{xn}(s)| \approx 0 & \forall\, \omega \neq \omega_n. \end{array} \quad (2)$$

These conditions correspond to an idealized triangular loop-gain profile, where the gain exceeds unity within a finite frequency band. This is rendered by a generalized delay-compensated band-pass damping controller, expressed as:

$$C_d(s) = \underbrace{g(n)\left(\frac{s}{(s+\omega_c)^2}\right)^n \cdot h(m)(s+\omega_c)^m}_{\text{Band-Pass Damping Controller}} \cdot \underbrace{\sigma\left(\frac{s-\omega_d}{s+\omega_d}\right)}_{\text{NMP Filter}}, \quad (3)$$

where ω_c and ω_d represent the corner frequency of band-pass and NMP filter respectively, with $g(n) = M(2\omega_c)^n$ and $h(m) = 2^{-m/2}\omega_c^{-m}$ ensure a fixed maximum gain $M > 0$ for any combination of slope orders (n, m).

3 Illustrative Case Studies

The band-pass controller is designed to provide narrow-band damping at the targeted mode while minimizing low- and high-frequency spillover and noise amplification. Performance is compared with benchmark Positive Position Feedback (PPF) and Negative Position Feedback (NPF) controllers.

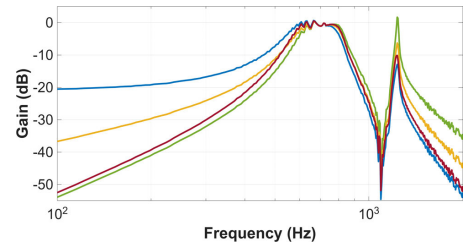

(a)

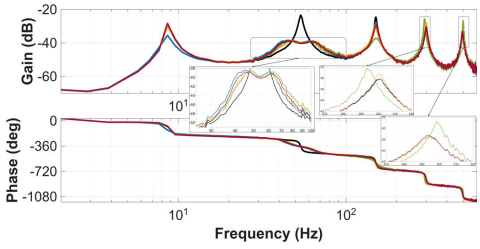

(b)

Figure 1: Simulated frequency response of (a) collocated nanopositioning system $G(s)$ (—) and closed-loop response $|\mathcal{S}_{xn}(s)|$, and (b) non-collocated flexural beam $G(s)$ (—) and closed-loop response $\mathcal{PS}_{xd}(s)$, obtained via implementing active damping controller $C_d(s)$ as PPF controller (—), NPF controller (—), proposed band-pass controller with $n = 1$ (—), and with $n = 2$ (—).

On the Price of Constraints in Linear Tracking Model Predictive Control*

Anastasios Vlachos[1] a.vlachos@rug.nl
Aren Karapetyan[2] aren.karapetyan@belimo.ch
Anastasios Tsiamis[3] atsiamis@control.ee.ethz.ch
Efe C. Balta[3,4] efe.balta@inspire.ch
John Lygeros[3] jlygeros@ethz.ch

1 Introduction

Trajectory tracking under constraints is a central problem in control, arising in applications where actuator limits and safety requirements must be enforced. While unconstrained linear quadratic tracking admits closed-form optimal solutions, the presence of input constraints fundamentally alters achievable performance. Model Predictive Control (MPC) is the standard framework for addressing constrained tracking problems; however, a quantitative characterization of the performance degradation induced by constraints remains largely unexplored.

In this work, given a reference described by the pair $\mathbf{r} \coloneqq (x_{\mathrm{ref}_t}, u_{\mathrm{ref}_t})_{t=0}^{T}$, we are concerned with the following constrained optimization problem

$$\min \underbrace{\|x_T - x_{\mathrm{ref}_T}\|_P^2 + \sum_{t=0}^{T-1} \|x_t - x_{\mathrm{ref}_t}\|_Q^2 + \|u_t - u_{\mathrm{ref}_t}\|_R^2}_{J(u;x_0,\mathbf{r})}$$
$$\text{s.t.} : x_{t+1} = Ax_t + Bu_t,$$
$$u_t \in \mathcal{U} = \{u \in \mathbb{R}^m | Hu \leq h\}$$

Our goal is to quantify how much optimality is sacrificed due to input constraints and to understand how this degradation depends on both the constraint set and the structure of the reference trajectory.

2 Regret of Constraints

We compare the optimal constrained MPC policy u^c with the unconstrained optimal controller $u^\star$ and define the resulting performance gap as the Regret of Constraints, formally as

$$\mathcal{R}^c(u^c) = J(u^c; x_0, \mathbf{r}) - J(u^\star; x_0, \mathbf{r}). \tag{1}$$

This notion captures the intrinsic cost of enforcing constraints in tracking control.

Our analysis applies to general reference trajectories, which may be unsafe, i.e. violate the constraints or unrealizable with respect to the system dynamics. To address this, we introduce a two-step reference transformation that enforces realizability and constraint satisfaction, enabling a separation between constraint-induced and reference-induced performance losses.

Our main technical result provides explicit complexity bounds on the Regret of Constraints. Informally, for the tracking problem, the Regret of Constraints is bounded by

$$\mathcal{R}^c \leq \underbrace{\max\left\{0, \mathcal{O}\left(\log \frac{1}{\min_i h_i}\right)\right\}}_{\text{constraints influence}} + \underbrace{\mathcal{O}(\|E\|^2) + \mathcal{O}(\|W\|^2)}_{\text{reference structure influence}}, \tag{2}$$

where E and W quantify the unsafety and the unrealizability of the reference trajectory, and become 0 when the reference is safe and realizable, respectively.

For more details on the derivation of this result, as well as numerical simulations that validate the theoretical bound, see [1].

3 Conclusion

This work provides a transparent characterization of the trade-offs inherent in constrained tracking MPC. By explicitly quantifying the suboptimality of constraints, the results offer insight into performance limitations of constrained tracking controllers and lay the groundwork for extensions to online tracking control.

References

[1] A. Vlachos, "The price of constraints in linear tracking model predictive control," Master thesis, ETH Zürich, 2025.

*This work is based on the first author's Master's thesis completed at the Automatic Control Laboratory, ETH Zürich.

[1]Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands.
[2]BELIMO Automation AG, 8340 Hinwil, Switzerland.
[3]Automatic Control Laboratory, Department of Information Technology and Electrical Engineering, ETH Zürich, 8092 Zürich, Switzerland.
[4]Control and Automation Group, inspire AG, 8005 Zürich, Switzerland.

Learning Cost-to-Go for Myopic Model Predictive Control using Sobolev training

Tom Minten[1], Koen de Vos[1], Dinesh Krishnamoorthy[2], Elena Torta[1]
[1] Robotics group, Department of Mechanical Engineering, Eindhoven University of Technology
[2] Department of Engineering Cybernetics, Norwegian University of Science and Technology
P.O. Box 513, 5600 MB, Eindhoven, The Netherlands
Email: t.p.minten@tue.nl

1 Introduction

Model Predictive Control (MPC) is widely used to compute optimal trajectories for robotic systems with nonlinear models. However, solving nonlinear problems is computationally demanding, limiting real-time performance. To address this, we shorten the prediction horizon and add a learned cost-to-go function, using a gradient-informed deep neural network. Unlike prior work [1] that learns only the cost-to-go value, we also aim to learn the local curvature using a Sobolev framework with available gradient data, enabling efficient learning.

2 Expert and Myopic MPC

The trajectory generating MPC solves:

$$V^* = \min_{\mathbf{x}_k, \mathbf{u}_k} \sum_{k=0}^{N} l(\mathbf{x}_k, \mathbf{u}_k) \tag{1a}$$

$$\text{s.t. } \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) \tag{1b}$$

$$\mathbf{x}_0 = \mathbf{x}_{\text{init}} \tag{1c}$$

$$\mathbf{x}_k \in \mathscr{X}, \mathbf{u}_k \in \mathscr{U} \tag{1d}$$

With $\mathbf{x}_k$ and $\mathbf{u}_k$ representing the system state and control input at time step k. $l(\mathbf{x}_k, \mathbf{u}_k)$ denotes the stage cost, penalizing distance to the target point and excessive control effort. V^* is the optimal objective value. Function f represents the system's nonlinear dynamical model. This MPC formulation is used as expert MPC, for generating training data. The myopic MPC formulation, adapted from [1], is defined as:

$$\begin{aligned} V^* = \min_{\mathbf{x}_k, \mathbf{u}_k} & \sum_{k=0}^{N_m-1} l(\mathbf{x}_k, \mathbf{u}_k) + \mathscr{V}(\mathbf{x}_{N_m}; \theta) \\ \text{s.t. } & \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) \\ & \mathbf{x}_0 = \mathbf{x}_{\text{init}} \\ & \mathbf{x}_k \in \mathscr{X}, \mathbf{u}_k \in \mathscr{U} \end{aligned}$$

With $N_m << N$ as the myopic horizon length, set to $N_m = 2$ in this work. $\mathscr{V}(\mathbf{x}_{Nm}; \theta)$ represents the approximated cost-to-go function, modeled by a deep neural network. Note that including the curvature information in the cost-to-go using Sobolev training, enables a data efficient learning method.

3 Learning Cost-to-Go via Sobolev Supervised Training

The training dataset is created by solving expert MPC for randomly sampled states. Each solution provides the optimal value V^* and its gradient $\nabla_x \mathscr{V}(x_j; \theta)$. The gradient is essentially given by the Lagrange multiplier λ_j, associated with the initial condition constraint (1c). The resulting dataset $\mathscr{D} = \{(x_j, \lambda_j, V^*(x_j))\}_{j=0}^{n_s}$ is used for Sobolev training [2]. The cost-to-go network $\mathscr{V}(x; \theta)$ is trained by minimizing:

$$\min_{\theta} \sum_j \|V_j^* - \mathscr{V}(x_j; \theta)\|^2 + \alpha \| -\lambda_j - \nabla_x \mathscr{V}(x_j; \theta)\|^2$$

where α balances both loss terms. The first term fits the scalar cost-to-go, while the second enforces gradient consistency, acting as a soft regularization that promotes descent properties similar to [1]. Unlike their hard constraints, this approach is practical for training deep neural networks, ensuring smooth gradients without compromising feasibility.

4 Results and Outlook

As a proof of concept, a simulated ground robot with four system states was controlled using an expert MPC with $N = 10$. The proposed myopic MPC achieved a reduction in computation time of 35.3% and 37.3% across two scenarios, while respecting system constraints and generating trajectories close to those of the expert MPC. The neural network comprises 4 hidden layers with 32 neurons each. These results highlight the potential of gradient-informed neural networks for accelerating MPC without sacrificing feasibility. Future work will extend this approach to higher-dimensional problems related to aerial platforms.

References

[1] T.M.J.T. Baltussen, C.A. Orrico, A. Katriniok, W.P.M.H. Heemels, and D. Krishnamoorthy. Value Function Approximation for Nonlinear MPC: Learning a Terminal Cost Function with a Descent Property. In *IEEE Conference on Decision and Control (CDC)*, Rio de Janeiro, 2025.

[2] Wojciech Marian Czarnecki, Simon Osindero, Max Jaderberg, Grzegorz Swirszcz, and Razvan Pascanu.

Sobolev Training for Neural Networks. *Advances in Neural Information Processing Systems*, 2017-December:4279–4288, 6 2017.

Exploiting patterns for demand response control with Reinforcement Learning

Jeroen Willems*, Denis Steckelmacher†, Jonathan Baake*, Edward Kikken* Bruno Depraetere*
* Flanders Make, Belgium. E-mail: firstname.lastname@flandersmake.be
† Vrije Universiteit Brussels: firstname.lastname@vub.be

1 Motivation

There are many challenging problems in industry that call for controllers beyond PID. One set of examples are controllers or EMS′s that coordinate multiple units (generators, batteries, heaters, chillers, pumps, ...) in the face of variable demands. Such demand response problems are challenging since they typically require anticipation: to provide response when there are changes in demand despite long delays, and to avoid cycling between on/off states of units when the demand fluctuates (as such units often have on/off states, with significant wear caused by switching on).
For such problems MPC is a natural control candidate, but has three downsides: (i) the need for a control model, (ii) the need to have good predictions available at runtime, and (iii) the high CPU load following from online optimization. Rule Based Control (RBC), the de-facto industry standard, avoids these issues, but its performance is much worse. In this work we study Reinforcement Learning (RL), and investigate if it can reduce these MPC drawbacks.

2 Setting up Reinforcement Learning

We used RL to train a policy for a case with 3 chillers, that need to be coordinated to supply cooling for a variable demand, with the set of demands shown in Figure 1.
We provide the agent with inputs consisting of the current system temperatures, the current states of the 3 chillers, and a set of lagged values of these variables. Crucially, we do not give as an input the time of day nor a prediction of the upcoming demand.
The agent produces one real-valued control signal per chiller, in the -1 to 1 range. Values below -0.5 turn the chiller off. Values between -0.5 and 0 set it to minimum power. Values from 0 to 1 interpolate between minimum and maximum power.

3 Results and discussion

After training, we run the resulting policy on the simulation model. We find that the outcomes of the RL agent match those of the MPC, and both easily outperform the RBC. This can be seen in Table 3, where demands are used similar to those of the training set. In this case we find that the RL can learn a single policy for the entire range of (unobserved) demands. MPC does require information about the demand.

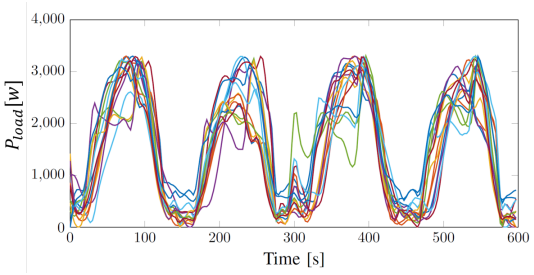

Figure 1: Set of semi-periodic demands considered.

Algorithm	Average cost (train set)	Average cost (test set)
MPC	454 J	450 J
RL	450 J	448 J (≈ MPC cost)
RBC	517 J	512 J

Table 1: Performance for the nominal case.

However, when demands outside of the considered training range are faced, the performance degrades. In Table 3, we show that RL degrades slightly more than the MPC, performing 3% worse. The MPC performance in turn degrades when it is given wrong predictions, as would typically be faced in realistic applications, while the RBC and RL remain unaffected.

Algorithm	Average cost (test set)
MPC	380 J
RL	391 J (+3% compared to MPC of 380 J)
RBC	463 J

Table 2: Impact with scaled down demands.

In conclusion, we can see that RL is able to deliver similar control performance to MPC. It does so without two of the MPC drawbacks: CPU load during operation, and explicit predictions. We still use a model for training the RL agent, so future work will focus on enabling usage of more approximate models, while preserving performance on the real system. We will apply similar ideas to also make RL generalize better to demands outside the training set.

Acknowledgements: This research is done in the framework of the Flanders AI Research Program (https://www.flandersairesearch.be/en) that is financed by EWI (Economie Wetenschap & Innovatie).

Acyclic coupling control: Applied to interconnected mechatronic systems

M. van der Hulst[1,*], N. Dirkx[2], R.A. González[1], K. Tiels [1], J. van de Wijdeven[2], T. Oomen[1,3]
[1] Dept. of Mechanical Engineering, Eindhoven University of Technology, The Netherlands
[2] ASML, Veldhoven, The Netherlands
[3]Delft Center for Systems and Control, Delft University of Technology, The Netherlands
*Email: M.v.d.hulst@tue.nl

1 Background

The performance of high-tech mechatronic systems is often dictated by shared objectives across the many closed-loop controlled subsystems operating in the machine, such as synchronization, coordination, and alignment. Achieving optimal performance therefore requires control methods that explicitly address these system-level objectives. This work aims to improve system-level performance using coupling controllers added to the existing subsystem controllers, where a structural constraint on the added coupling is included to guarantee closed-loop stability.

2 Problem formulation

Consider the system of N closed-loop controlled subsystems in Fig. 1, where $G_0 = \mathrm{diag}(G_1,\dots,G_N)$ and $K_0 = \mathrm{diag}(K_1,\dots,K_N)$ contain the subsystem plants and controllers. Combined performance is determined by $e_{\mathcal{S}} = We$ with subsystem errors $e \in \mathbb{R}^N$ and objective matrix $W \in \mathbb{R}^{n_w \times N}$, $X \in \mathbb{R}^{N\times N}$ is the coupling design parameter, and $Z \in \mathcal{RH}_\infty^{N\times N}$ is chosen such that $G^{-1}Z \in \mathcal{RH}_\infty^{N\times N}$. The coupling X is constrained to be acyclic, i.e., there exists a permutation matrix Π such that $\Pi^\top X \Pi$ is lower triangular, which for a stable X guarantees closed-loop stability. The problem is to compute acyclic X that minimizes

$$\hat{X} = \underset{X\in\mathbb{R}^{N\times N}}{\operatorname{argmin}} \left\|T_{e_{\mathcal{S}}w}(X)\right\|_{\mathcal{H}_2}^2 + \mu\left\|T_{uw}(X)\right\|_{\mathcal{H}_2}^2, \quad \text{s.t. } X \text{ acyclic}, \tag{1}$$

where $T_{e_{\mathcal{S}}w}(X)$ and $T_{uw}(X)$ denote the closed-loop transfer functions from disturbances to $e_{\mathcal{S}}$ and to control input, respectively.

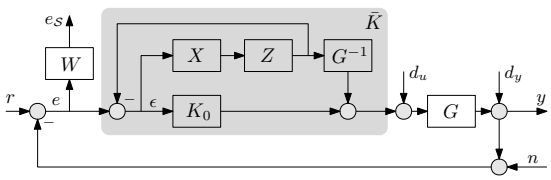

Figure 1: Decentralized feedback control interconnection with add-on coupling controllers.

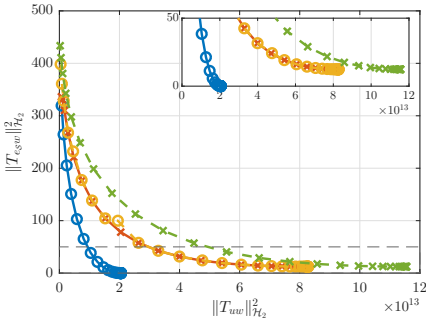

Figure 2: Optimal Pareto front with unconstrained X (-o-), the optimal acyclic X (-o-), the designed acyclic X (-o-), and acyclic solution X obtained through heuristics (-o-).

3 Approach

The main difficulty of problem (1) is the acyclicity constraint, which induces a discrete, combinatorial search space that grows exponentially with N. As result, direct enumeration or mixed integer formulations become intractable for moderate to large N. Instead, acyclicity is enforced using smooth analytic constraints

$$h(X) = \log\det\big(sI - X\circ X\big) - N\log s, \tag{2}$$

where $s > \rho(X\circ X)$ and $h(X) = 0$ if and only if X is acyclic [2]. Hence, problem (1) can be formulated as a smooth, nonconvex optimization problem, which can be solved using standard gradient-based optimization methods.

4 Results

Consider a system of $N = 5$ closed-loop controlled free-moving masses. Figure 2 shows the optimal Pareto curve, obtained by solving Problem (1) for $\mu \in [0, 10^{-8}]$. It can be observed that imposing acyclicity on X leads to only a minor degradation in achievable performance, while guaranteeing that the closed-loop system can never be destabilized by the added coupling.

References

[1] E. Evers, M. Van de Wal, and T. Oomen, "Beyond decentralized wafer/reticle stage control design: A double-Youla approach for enhancing synchronized motion," Control Engineering Practice, vol. 83, pp. 21–32, 2019.

[2] K. Bello, B. Aragam, and P. Ravikumar, "DAGMA: Learning DAGs via M-matrices and a log-determinant acyclicity characterization," in Proc. 36th Conf. Neural Inf. Process. Syst. (NeurIPS), pp. 8226–8239, 2022.

Model predictive control of the plasma density profile with pellet fueling in the Large Helical Device nuclear fusion stellarator

Christopher Orrico[1,2,†], Naoki Kenmochi[3], Yuya Morishita[4], Sadayoshi Murakami[4], Maaike Hoven[1,2], Martijn de Beij[1,2], Lennard Ceelen[1,2], W.M.P.H Heemels[2], Matthijs van Berkel[1], Dinesh Krishnamoorthy[5]

[1] DIFFER - Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands

[2] Eindhoven University of Technology, Control Systems Technology, Eindhoven, the Netherlands

[3] National Institute for Fusion Sciences, Toki, Japan

[4] Kyoto University, Department of Nuclear Engineering, Kyoto, Japan

[5] Norwegian University of Science and Technology, Department of Engineering Cybernetics, Trondheim, Norway

[†] c.a.orrico@tue.nl

1 Introduction

To achieve the plasma conditions necessary to produce net energy in a fusion reactor, we must control the plasma density using frozen hydrogen fuel pellets fired into the fusion plasma core [1]. Yet, many challenges in plasma density profile control with pellets and gas fueling remains unsolved. One problem faced in developing density control for the Large Helical Device (LHD) stellarator is that only approximately 50% of requested pellets fire successfully. The plasma density also must never exceed the Sudo limit, or the plasma will disrupt [2]. Complicating matters further, feedback control for pellets is also unavailable currently on LHD. To overcome these challenges (and others discussed in [1]), we implement a scheme combining system identification, predictive feedforward optimization of pellet injection, and feedback control of gas injection in experiments on LHD. The experimental results are the first successful demonstration of predictive density control with pellet injection on a nuclear fusion reactor.

2 System identification

Perturbative experiments using pellets over multiple injection frequencies and periodic multisine gas fueling signals are performed on over a number of LHD discharges. 19 pellet-perturbed discharges and 7 gas-perturbed discharged are used for system identification. We regress a state space model of dimension $n_x = 4$, $n_u = 4$, $n_y = 101$ using the N4SID algorithm [3]. The output of the state-space is a 1D profile of the plasma density from the plasma core to plasma edge. The model is validated against 6 discharges, showing an average prediction error of 3.88% (evalauted as RRMSE).

3 Offline optimization

We solve a quadratic program for the sequence of pellet decisions that minimizes the tracking error of the predicted density evolution over the entire discharge w.r.t a desired output density profile while never violating the Sudo limit. However, this process must also account for the high pellet uncertainty. We implement a scenario approach [4], whereby we optimize a common input sequence for two scenarios: a *performance* scenario, where only 50% of pellets fire and a *safety* scenario where *all* pellets fire. The pellet sequence must satisfy the constraints of *both* scenarios, ensuring the density limit is not violated.

4 Online MPC

We implement tracking MPC [5] based on the contingency MPC (CMPC) approach [4], which uses gas injection to correct for the pellet uncertainty in real-time. The CMPC predicts the sequence of gas inputs up to a prediction horizon that minimizes the error w.r.t to the desired density profile while satisfying density constraints for both the performance and safety scenarios introduced for offline optimization. The pellet sequence computed during offline optimization is supplied as a parameter for the MPC, along with the initial state computed using a Kalman-filter based observer.

5 Results

We implement the offline optimization and online CMPC for three references on LHD: a low density, mid density, and high density plasma. For each reference, we see that the pellet timings computed using offline optimization are able to drive the plasma density near the desired value, while the online CMPC corrects for missed pellets with gas injection.

References

[1] C.A. Orrico et al 2025 Nucl. Fusion 65 076041

[2] S. Sudo et al 1990 Nucl. Fusion 30 11

[3] van Overschee, P., and B. De Moor. "Subspace Identification of Linear Systems: Theory, Implementation, Applications". Springer Publishing: 1996

[4] Alsterda, J.P. et al "Contingency model predictive control for automated vehicles". In 2019 American Control Conference (ACC) (pp. 717-722). IEEE.

[5] Krupa, P. et al "Model predictive control for tracking using artificial references: Fundamentals, recent results and practical implementation." In 2024 IEEE 63rd Conference on Decision and Control (CDC) (pp. 2977-2991). IEEE.

Path-complete upper bound on the value function of switched systems

Léa Ninite[1], Adrien Banse, Guillaume O. Berger, Raphaël M. Jungers
ICTEAM, UCLouvain

1 Introduction

We study the problem of estimating the *value function* of discrete-time *switched linear systems* under *arbitrary* switching. Such systems are described by

$$x_{k+1} = A_{\sigma(k)} x_k, \quad k \in \mathbb{N},$$

where $x_k \in \mathbb{R}^n$ is the state at time k, $\sigma(k) \in \langle M \rangle := \{1, 2, \dots, M\}$ denotes the mode, and $A_i \in \mathbb{R}^{n \times n}$ for each $i \in \langle M \rangle$.

Given a nonnegative quadratic stage cost $c(x) = x^\top Q x$, $Q \succeq 0$, the associated *value function* is defined as $J(x) = \sup_{\sigma:\mathbb{N}\to\langle M\rangle} \sum_{k=0}^{\infty} c(x_k)$, representing the worst-case total cost.

For such systems, the exponential growth of possible mode sequences over time renders the exact computation of the value function intractable, motivating the design of tractable bounds that provide accurate approximations.

The case where the mode $\sigma(k)$ is *controlled* has been extensively studied in the literature, e.g., [1], whereas much less is known for the case of *arbitrary* switching.

2 Path-complete upper bound

We use the framework of *path-complete graphs* [2] to encode inequalities and compute an upper bound on the value function J. Let $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ be a directed labeled graph, with nodes $\mathcal{S}$ and edges $\mathcal{E} \subseteq \mathcal{S} \times \mathcal{S} \times \langle M \rangle$, where each edge $(\alpha, \beta, i) \in \mathcal{E}$ represents a possible transition under mode i. A graph is *path-complete* if for every finite sequence of modes there exists a path in the graph whose edge labels follow that sequence.

For switched linear systems with quadratic stage cost $c(x) = x^\top Q x$, we associate to each node $\alpha \in \mathcal{S}$ a quadratic function $V_\alpha(x) = x^\top P_\alpha x$, $P_\alpha \succeq 0$. Each edge (α, β, i) encodes the following LMI, inspired by the Bellman equation:

$$P_\alpha \succeq Q + A_i^\top P_\beta A_i. \tag{1}$$

From the functions $\{V_\alpha\}_{\alpha\in\mathcal{S}}$, one can construct an upper bound on the value function $J(x)$ as explained in the theorem below.

Theorem 1. If the graph $\mathcal{G} = (\mathcal{S}, \mathcal{E})$ is *co-complete*[2], and if $\{P_\alpha\}_{\alpha\in\mathcal{S}}$ satisfy (1), then the function $V(x) := \max_{\alpha\in\mathcal{S}} x^\top P_\alpha x$ satisfies $V(x) \geq J(x)$ for all $x \in \mathbb{R}^n$.

[1]Corresponding author: lea.ninite@uclouvain.be

[2]$\mathcal{G} = (\mathcal{S}, \mathcal{E})$ is complete for $\langle M \rangle$ if, for each node $\beta \in \mathcal{S}$ and each mode $i \in \langle M \rangle$ there exists at least one node $\alpha \in \mathcal{S}$ such that $(\alpha, \beta, i) \in \mathcal{E}$.

Importantly, the path-complete graph $\mathcal{G}$ is a *design parameter* in our procedure, allowing a trade-off between complexity and bound tightness.

3 Approximation guarantee

Given a co-complete graph $\mathcal{G} = (\mathcal{S}, \mathcal{E})$, and the corresponding upper bound $V(x) = \max_{\alpha\in\mathcal{S}} x^\top P_\alpha x$ on $J(x)$, one can compute, via LMIs, a factor $\mu \geq 1$ such that

$$\frac{1}{\mu} V(x) \leq J(x) \leq V(x), \quad \forall x \in \mathbb{R}^n.$$

This μ quantifies the conservativeness of the bound V.

4 Results

To obtain asymptotically tight bounds, we consider *dual De Bruijn graphs* $\mathcal{G}_l = (\mathcal{S}_l, \mathcal{E}_l)$ of increasing order l. We can show that there exist matrices $\{P_\alpha^l\}_{\alpha\in\mathcal{S}_l}$ satisfying (1) such that $V^l(x) := \max_{\alpha\in\mathcal{S}_l} x^\top P_\alpha^l x$ approaches the true value function $J(x)$ as $l \to \infty$. Figure 1 shows that, as the order of the graph increases, the upper bound $V^l(x)$ and the scaled lower bound $\frac{1}{\mu} V^l(x)$ get closer, illustrating the non-conservativeness of our approach.

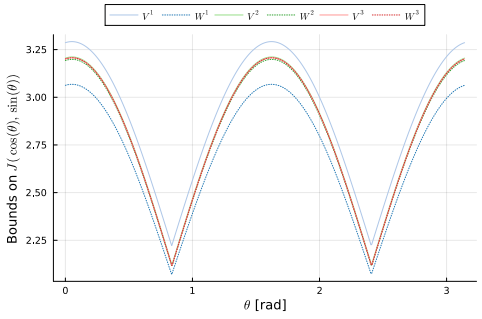

Figure 1: Upper and lower bounds on the value function, $V^l(x)$ and $W^l(x) := V^l(x)/\mu$, for $x = (\cos\theta, \sin\theta)$, obtained using dual De Bruijn graphs of order $l = 1, 2, 3$.

References

[1] A. Rantzer, "On approximate dynamic programming in switching systems," in *Proceedings of the 44th IEEE Conference on Decision and Control*, IEEE, 2005.

[2] A. A. Ahmadi, R. M. Jungers, P. A. Parrilo, and M. Roozbehani, "Joint spectral radius and path-complete graph lyapunov functions," *SIAM Journal on Control and Optimization*, vol. 52, no. 1, 2014.

Security-Aware Planning and Control of Multi-Agent Systems with LTL Tasks

Georgios Mitsos [1] Dimos V. Dimarogonas [2] Siyuan Liu [1]

1 Introduction and Problem Setting

Temporal logic–based planning provides a formal framework for specifying complex tasks and synthesizing correct-by-design controllers. By abstracting agent dynamics into finite-state models, automata-based techniques can generate discrete plans that are subsequently refined into continuous trajectories [1]. In multi-agent systems, such methods provide a formal way to encode team objectives beyond classical cooperative control; however, scalability issues due to joint state spaces have motivated hierarchical and decentralized solutions. Although many autonomous cyber-physical systems require both correctness and confidentiality, information security during task execution has received comparatively little attention. Recent work integrates information-flow security notions such as opacity into robot planning [2], but, to the best of our knowledge, no existing approach addresses multi-agent systems with continuous-time dynamics in which temporal logic specifications and security constraints must be enforced simultaneously.

We consider a multi-agent system operating in a bounded polytopic workspace, where each agent follows continuous-time linear affine dynamics with bounded inputs. A global mission is specified using a linear temporal logic (LTL) formula. The workspace is partitioned into convex regions, and each agent's motion is abstracted as a weighted transition system with states corresponding to regions and transitions representing feasible motions between them. States are labeled with atomic propositions and mapped to observations available to a passive intruder, while a subset of states associated with sensitive tasks is designated as secret. The synchronous product of the individual abstractions yields a global transition system describing the system evolution.

The objective is to synthesize dynamically feasible trajectories such that the system satisfies the LTL specification while preventing the intruder from inferring whether a secret task has been executed and which agent performed it.

2 Planning and Control Approach

To capture information-flow security, we introduce two complementary security notions. *Type-A security* prevents the intruder from inferring whether a secret task occurred, while *Type-B security* ensures the intruder cannot identify the responsible agent. These notions can be imposed independently or jointly. Security is enforced at the discrete level by building secure-by-construction abstractions. The feasibility of each discrete transition is validated by solving a convex trajectory optimization problem subject to the agents' dynamics and input bounds. Control barrier functions ensure that continuous trajectories correctly implement the discrete transition. Standard LTL synthesis is then applied to the resulting secure and dynamically feasible abstraction by synchronizing it with a Büchi automaton corresponding to the LTL specification. An accepting prefix–suffix path is computed using shortest-path algorithms. Finally, the trajectory segments associated with the synthesized discrete path are concatenated. The resulting trajectories are dynamically feasible, satisfy the LTL specification, and enforce the desired security properties by construction. Fig. 1 demonstrates a scenario where Agent 1 and Agent 2 must eventually reach regions A and B, respectively, and always avoid the obstacle, while satisfying Type-A/B security.

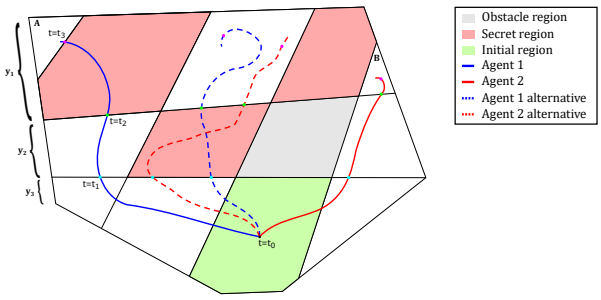

Figure 1: Dashed trajectories are indistinguishable alternatives with respect to observations y_1, y_2, y_3. Colored dots indicate the agents' positions at $t = t_0, t_1, t_2, t_3$.

Compared to existing security-preserving multi-agent planning approaches, the proposed method achieves significantly lower computational complexity.

References

[1] Kloetzer, M. and Belta, C. (2008). *A fully automated framework for control of linear systems from temporal logic specifications*. IEEE Trans. Autom. Control, 53(1), 287-297.

[2] Yu, X., Yin, X., Li, S., and Li, Z. (2022). *Security-preserving multi-agent coordination for complex temporal logic tasks*. Control Eng. Pract., 123, 105130.

[1]Department of Electrical Engineering, Eindhoven University of Technology, Groene Loper 19, 5612 AP Eindhoven, The Netherlands ({g.mitsos, s.liu5}@tue.nl).

[2]Division of Decision and Control Systems, KTH Royal Institute of Technology, Malvinas väg 10, 11428 Stockholm, Sweden (dimos@kth.se).

PCLF extremal lifts are Galois connections, and we sometimes get a full adjunction.

Sébastien Mattenet
UCLouvain
sebastien.mattenet@uclouvain.be

Raphaël Jungers
UCLouvain
raphael.jungers@uclouvain.be

1 Introduction

In the study of arbitrary switching dynamical systems, one is interested in finding a Lyapunov (energy-like) function to prove stability. One method of accelerating the search is through a Path Complete Lyapunov Function (PCLF) [1], that is a graph with candidate Lyapunov function (CLF) at each node and satisfying some conditions on the edges. It has quickly become apparent that, when searching a given class of problems (for example linear problems with quadratic candidates), some graph shapes are strictly worse to others, in the sense that if one finds a solution the other will find one as well. In restricted classes of problems (namely problems with a set of CLF closed under min/max/sum/nothing) this ordering on graph $H \geq_P G$ is equivalent to admitting admitting a graph morphism to a lifted graph $H \to L(G)$. In this paper we investigate the consequence of having an equivalent lift for an arbitrary class of problem.

2 Categorical approach and definitions

We introduce some standard definition of category theory, up to adjunctions. A category is a collection of objects, arrows between them and a composition operation that is associative and such that every object has an identity arrow. Let $\mathbf{Graph}_L$ be the category of L-labeled graphs and graph morphism.
For any category C, let $\mathbf{Ord}\mathsf{C}$ be the category with the same objects but at most one arrow between any two objects. For example $\mathbf{OrdGraph}_L$ is the category of L-labeled graphs, and there is an arrow $G \to H$ if there is a graph morphism $G \to H$.

Given a class of problem P we can define the order $\leq_P$ on graphs by comparing which is better for every problem in P. We can define the category $\mathbf{OrdGraph}_{\leq_P}$, whose objects are L-labeled graphs and whose unique arrow is $G \to H$ iff $G \geq_P H$.

A **lift** is a functor[1] $L : \mathbf{Graph}_L \to \mathbf{Graph}_L$. A lift L is said to be **valid** for a class of problem P if we have the following:
There is a simulation function $H \to L(G) \implies H \geq_P G$.
If the implication goes both way we say the lift is an **extremal valid lift**.

[1] function preserving the categorical structure, see [2].

Lastly, given two functors $R : D \to C$ and $L : C \to D$, a common type of relation is and **adjunction** $R \vdash L$. The precise definition can be found in any book of category, such as [2], but in the case of **Ord**-categories, this correspond to a Galois connection.

$$x \geq R(y) \iff L(x) \geq y, \quad \forall x \in C, y \in D$$

3 Results

A lift is an extremal valid lift if and only if there is an adjunction with $\mathbf{Ord}(R) \circ \mathbf{Ord}(i) = \mathbf{Ord}(L)$

$$\mathbf{Ord}(L) \circlearrowright \mathbf{OrdGraph}_L \underset{\mathbf{Ord}(R)}{\overset{\mathbf{Ord}(i)}{\underset{\longleftarrow}{\overset{\longrightarrow}{\bot}}}} \mathbf{OrdGraph}_{\leq_P}.$$

A Galois connection is the truncated version of an adjunction, the full picture require a category of problem $\mathbf{Prob}_P$ and we get an adjunction

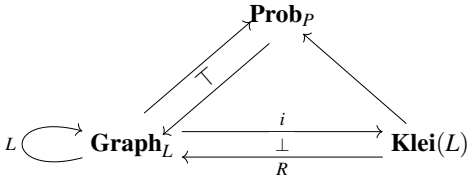

with $\mathbf{OrdKlei}(L) = \mathbf{OrdGraph}_{\leq_P}$. Standard result on adjunction [2] gives us the following corollaries.

Two graphs G, H are said to be order equivalent if there is both a graph morphism $H \to G$ and a graph morphism $G \to H$. If L is an extremal valid lift for P, then the lift is idempotent up to order equivalence. Let G, H be two graphs, $G \times H$ the cartesian product of graphs and assume the supremum $\sup_{\leq_P}(G,H)$ exists for the order, then
1) $L(\sup_{\leq_P}(G,H))$ and $L(G) \times L(H)$ are order equivalent;
2) $\inf_{\leq_P}(G,H)$ must be $\leq_P$-equivalent to the disjoint union of G and H.

References

[1] A. Ahmadi, R. Jungers, P. Parrilo, and M. Roozbehani. "Joint Spectral Radius and Path-Complete Graph Lyapunov Functions". In: *SIAM Journal on Control and Optimization* 52 (2014), pp. 687–717.

[2] S. Mac Lane. *Categories for the Working Mathematician*. Second. Vol. 5. Grad. Texts in Math. Springer, 1998.

Virtual Sensing in Urban Water Quality Monitoring via Kalman Filtering and Machine Learning

Habash, A.[1,2,*], Wagterveld, R.M.[2,3], Cheng, X.[1], Keesman, K.J.[1,2], and van Heijster, P.[1]
[1]Mathematical and Statistical Methods (Biometris), Wageningen University and Research,
PO Box 16, 6700 AA, Wageningen, The Netherlands
[2]Wetsus, European Centre of Excellence for Sustainable Water Technology,
Oostergoweg 9, 8911 MA, Leeuwarden, The Netherlands
[3]Aqa.earth, Tweede Sluisweg 35, 8413 NN Oudehorne, The Netherlands
[*]Email: Abdelrahman.Habash@wur.nl

1 Introduction

With the extent of urbanization we are currently experiencing, alongside anthropogenic interference to natural water bodies, the importance of reliable water quality systems is becoming more important than ever. However, traditional approaches of water quality monitoring face significant limitations. Lab-based analysis (*ex-situ*) is costly and introduces delays in providing results, while sensor-based analysis (*in-situ*) has limited capability to detect complex contaminants.

2 Proposed Framework

This study aims to address these challenges by providing a framework for urban water quality monitoring that can be both cost-effective and provide real-time responses through virtual sensing. Virtual sensors can be utilized to estimate the level/presence of unknown water pollutants based on a limited set of process parameters with measurements from hardware sensors. Drawing upon the methodological foundations established in previous studies [1], [2], the proposed framework (illustrated in Figure 1) integrates Kalman filtering (KF) and machine learning (ML) in a hybrid approach. In this framework, a state-space model is formulated to describe the temporal evolution of measured (e.g., pH, temperature (T), and electrical conductivity (EC)) and unmeasured water-quality variables. The Kalman Filter (KF) operates as a recursive state estimator, combining this dynamical model with real-time sensor measurements to generate minimum-variance estimates of the underlying system states. Subsequently, these filtered state estimates are used as inputs for ML models to predict emerging pollutant concentrations.

3 Expected Outcomes

As urban municipalities need reliable ways to quickly spot water quality problems before they become severe, the foreseen widespread use of these virtual sensors makes it possible for cities to track water quality continuously and immediately detect issues, leading to quicker responses and water quality control. The proposed hybrid KF–ML framework can support accurate real-time estimation and short-term prediction of key water-quality indicators. This capability will improve the monitoring of emerging pollutants and facilitate earlier detection of anomalous conditions.

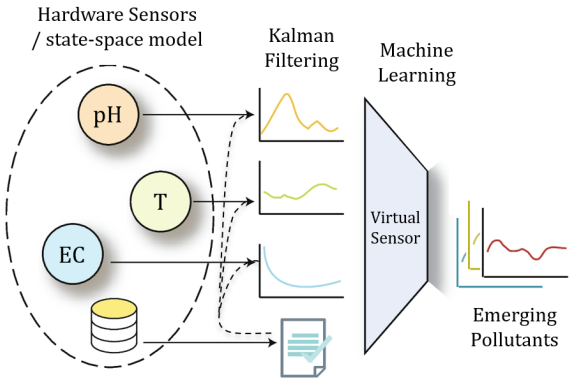

Figure 1: Schematic of proposed framework (modified from [3]).

References

[1] C. L. R. Wu, et al., "Reactive Transport Modeling for Exploring the Potential of Water Quality Sensors to Estimate Hydrocarbon Levels in Groundwater," *Water Resources Research*, Apr. 2024.

[2] E. A. Ross, et al., "Real-Time Monitoring of Chlorate Concentrations, Integrating Electrochemistry Modeling, and Kalman Filtering," *IEEE Trans. Ind. Inf.*, Nov. 2025.

[3] C. Luzon-Alvarez, M. et al., "Acoustic virtual sensors for industrial process monitoring using non-negative matrix factorization," *J AUDIO SPEECH MUSIC PROC.*, Aug. 2025.

Acknowledgement This work was performed in the cooperation framework of Wetsus, European Centre of Excellence for Sustainable Water Technology (www.wetsus.nl). Wetsus is co-funded by the European Union (Horizon Europe, LIFE, Interreg and EDRF), the Province of Fryslân and the Dutch Government: Ministry of Economic Affairs (TTT, SBO & PPS-I/TKI Water Technology), Ministry of Education, Culture and Science (TTT & SBO) and Ministry of Infrastructure and Water Management (National Growth Fund - UPPWATER). This work was carried out within the UrbanM2O project, which has received funding from the European Union under Grant Agreement No. 101180710. The authors like to thank the participants of the research theme "Monitoring & Quality" for the fruitful discussions and their financial support.

Virtual Sensing in Urban Water Quality Monitoring via Kalman Filtering and Machine Learning

Habash, A.[1,2,*], Wagterveld, R.M.[2,3], Cheng, X.[1], Keesman, K.J.[1,2], and van Heijster, P.[1]
[1]Mathematical and Statistical Methods (Biometris), Wageningen University and Research,
PO Box 16, 6700 AA, Wageningen, The Netherlands
[2]Wetsus, European Centre of Excellence for Sustainable Water Technology,
Oostergoweg 9, 8911 MA, Leeuwarden, The Netherlands
[3]Aqa.earth, Tweede Sluisweg 35, 8413 NN Oudehorne, The Netherlands
[*]Email: Abdelrahman.Habash@wur.nl

1 Introduction

With the extent of urbanization we are currently experiencing, alongside anthropogenic interference to natural water bodies, the importance of reliable water quality systems is becoming more important than ever. However, traditional approaches of water quality monitoring face significant limitations. Lab-based analysis (*ex-situ*) is costly and introduces delays in providing results, while sensor-based analysis (*in-situ*) has limited capability to detect complex contaminants.

2 Proposed Framework

This study aims to address these challenges by providing a framework for urban water quality monitoring that can be both cost-effective and provide real-time responses through virtual sensing. Virtual sensors can be utilized to estimate the level/presence of unknown water pollutants based on a limited set of process parameters with measurements from hardware sensors. Drawing upon the methodological foundations established in previous studies [1], [2], the proposed framework (illustrated in Figure 1) integrates Kalman filtering (KF) and machine learning (ML) in a hybrid approach. In this framework, a state-space model is formulated to describe the temporal evolution of measured (e.g., pH, temperature (T), and electrical conductivity (EC)) and unmeasured water-quality variables. The Kalman Filter (KF) operates as a recursive state estimator, combining this dynamical model with real-time sensor measurements to generate minimum-variance estimates of the underlying system states. Subsequently, these filtered state estimates are used as inputs for ML models to predict emerging pollutant concentrations.

3 Expected Outcomes

As urban municipalities need reliable ways to quickly spot water quality problems before they become severe, the foreseen widespread use of these virtual sensors makes it possible for cities to track water quality continuously and immediately detect issues, leading to quicker responses and water quality control. The proposed hybrid KF–ML framework can support accurate real-time estimation and short-term prediction of key water-quality indicators. This capability will improve the monitoring of emerging pollutants and facilitate earlier detection of anomalous conditions.

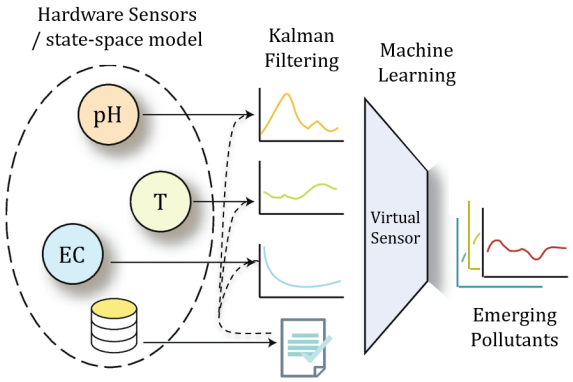

Figure 1: Schematic of proposed framework (modified from [3]).

References

[1] C. L. R. Wu, et al., "Reactive Transport Modeling for Exploring the Potential of Water Quality Sensors to Estimate Hydrocarbon Levels in Groundwater," *Water Resources Research*, Apr. 2024.

[2] E. A. Ross, et al., "Real-Time Monitoring of Chlorate Concentrations, Integrating Electrochemistry Modeling, and Kalman Filtering," *IEEE Trans. Ind. Inf.*, Nov. 2025.

[3] C. Luzon-Alvarez, M. et al., "Acoustic virtual sensors for industrial process monitoring using non-negative matrix factorization," *J AUDIO SPEECH MUSIC PROC.*, Aug. 2025.

Acknowledgement This work was performed in the cooperation framework of Wetsus, European Centre of Excellence for Sustainable Water Technology (www.wetsus.nl). Wetsus is co-funded by the European Union (Horizon Europe, LIFE, Interreg and EDRF), the Province of Fryslân and the Dutch Government: Ministry of Economic Affairs (TTT, SBO & PPS-I/TKI Water Technology), Ministry of Education, Culture and Science (TTT & SBO) and Ministry of Infrastructure and Water Management (National Growth Fund - UPPWATER). This work was carried out within the UrbanM2O project, which has received funding from the European Union under Grant Agreement No. 101180710. The authors like to thank the participants of the research theme "Monitoring & Quality" for the fruitful discussions and their financial support.

Correct-by-Design Control Synthesis of Stochastic Multi-Agent Systems: a Robust Tensor-based Solution

Ruohan Wang
TU Eindhoven
r.wang2@tue.nl

Siyuan Liu
TU Eindhoven
s.liu5@tue.nl

Sofie Haesaert
TU Eindhoven
s.haesaert@tue.nl

1 Abstract

Discrete-time stochastic systems with continuous state spaces are difficult to verify and control due to the curse of dimensionality. We propose an abstraction-based framework with robust dynamic programming that yields control strategies with provable lower bounds on temporal-logic satisfaction via approximate stochastic simulation relations. By exploiting decoupled dynamics, we uncover a Canonical Polyadic Decomposition tensor structure that enables scalable dynamic programming.

2 Problem Formulation

We consider **a multi-agent system** with continuous state space of which each agent i can be written as a general Markov decision process (gMDP) $\mathbf{M}^i = (\mathbb{X}_i, \mathbb{U}_i, \mathbb{Y}_i, \mathbb{T}_i, x_{i,0}, h_i), \quad i = 1, \ldots, N$. We are interested in temporal specifications ϕ that can be automatically translated to a deterministic finite automaton (DFA), as shown in Figure 1. We consider that every atomic proposition can be represented with a conjunction of Boolean predicates of every agent.

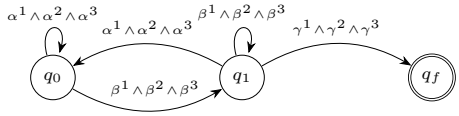

Figure 1: DFA.

Problem. Given a specification ϕ with DFA $\mathcal{A}_\phi$, synthesize control strategies $\mathbf{C}_\pi$ for a multi-agent system $\mathbf{M} = (\mathbf{M}^1, \ldots, \mathbf{M}^N)$ such that the controlled system satisfies ϕ with a lower bound p, i.e., $\mathbb{P}_{\mathbf{M} \times \mathbf{C}_\pi}(\boldsymbol{l} \vDash \phi) \geq p$.

3 Robust Tensor-based Solution

Finite abstractions. We approximate the continuous-state multi-agent system using a finite-state model with abstraction error quantified via approximate simulation relations. Assume there exists a multi-agent-level relation $\mathcal{R} \subset \hat{\mathbb{X}} \times \mathbb{X}$, induced by agent-wise relations $\mathcal{R}^i$ [1], such that $\hat{\mathbf{M}}^i \preceq^{\delta_i}_{\epsilon_i} \mathbf{M}^i$ with respect to metric $d_{\mathbb{Y}}(\hat{y}, y) := \max_i(d_{\mathbb{Y}_i}(\hat{y}_i, y_i))$ for all i. The composite relation is defined as $\mathcal{R} := \{(\hat{x}, x) \in \hat{\mathbb{X}} \times \mathbb{X} | (\hat{x}_i, x_i) \in \mathcal{R}^i, \forall i\}$, for which

$$\hat{\mathbf{M}} \preceq^{\boldsymbol{\delta}}_{\boldsymbol{\epsilon}} \mathbf{M}, \quad \boldsymbol{\delta} := 1 - \prod_{i=1}^{N}(1 - \delta^i), \quad \boldsymbol{\epsilon} := \max_i(\epsilon^i). \quad (1)$$

Value functions. Given a policy $\boldsymbol{\pi}$, define the value function for $\mathbf{M} : \mathbb{X} \times Q \to [0, 1]$ by forward recursion

$$\mathcal{V}^{\boldsymbol{\pi}}_{q,t+1}(x) = \mathbf{1}_{q_f}(q) + \mathbf{1}_{Q \setminus \{q_f\}}(q)\mathbb{E}^{\pi[T-1-t]}_{x,q}[\mathcal{V}^{\boldsymbol{\pi}}_{q',t}(x')], \quad (2)$$

where $\mathbb{E}^{\pi}_{x,q}[f] := \mathbb{E}[f(x', q')|(x, q), u = \pi(x, q)]$ with $q' := \tau(q, \alpha)$ with $\alpha = L(h(x'))$.

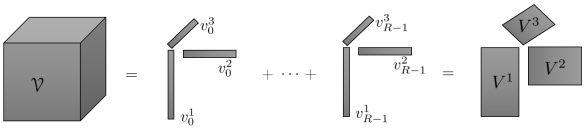

Figure 2: Illustration of an order-3 tensor with its rank-R CPD, v_i^j is a vector.

Canonical Polyadic Decomposition (CPD). There exists a CPD representation of the high-dimensional value function $\mathcal{V} \in \mathbb{R}^{d_1 \times \ldots \times d_N}$ in (2) as

$$\mathcal{V} = \sum_{r=1}^{R} V^{(1)}_{\bullet r} \otimes V^{(2)}_{\bullet r} \otimes \ldots \otimes V^{(N)}_{\bullet r}, \quad (3)$$

where $\otimes$ denotes the outer product, and $V^{(j)}_{\bullet r}$ is the r-th column of $V^{(j)}$. An order-3, rank-R tensor is visualized in Figure 2. There is internally a tree structure with nodes' value $v_i^j(x)$ that supports dynamic programming to compute $\mathcal{V}_q(x)$ [1]. We propose two different ways (an a-posteriori $\boldsymbol{\delta}$ correction and a robust-tree δ^i correction) in [1] to compute the lower bounds of satisfaction probabilities.

4 Case Study

We consider a two-agent system and a reach-avoid specification. We construct a finite-state abstraction for each agent. With a chosen interface function u_i and relation $\mathcal{R}^i$ we obtain the approximate simulation relation $\hat{\mathbf{M}}^i \preceq^{0.002}_{0.1} \mathbf{M}^i$ using **SySCoRe** tool [3]. We visualized the lower bounds of satisfaction probabilities in Figure 3.

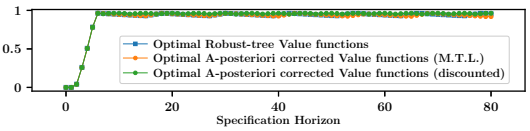

Figure 3: Lower bounds on the satisfaction probabilities of a reach-avoid specification on a 2-agent system as a function of specification horizon.

References

[1] Haesaert, S. et al. 2020. Robust dynamic programming for temporal logic control of stochastic systems. IEEE Transactions on Automatic Control, 66(6), pp.2496-2511.

[2] Wang, R. et al. 2025. Correct-by-Design Control Synthesis of Stochastic Multi-agent Systems: a Robust Tensor-based Solution. arXiv preprint arXiv:2511.06873.

[3] Van Huijgevoort, B. et al. SySCoRe: Synthesis via stochastic coupling relations. HSCC 2023, pp. 1–11.

[4] Hitchcock, F.L., 1927. The expression of a tensor or a polyadic as a sum of products. Journal of Mathematics and Physics, 6(1-4), pp.164-189.

Identification of dynamics of photosynthesis regulation under oscillating light condition

Arviandy G. Aribowo[1], Marcelo Orlando[2], Benjamin Bailleul[2,3], Maarten Schoukens[1], Leyla Özkan[1]
[1]Control Systems Group, Dept. of Electrical Eng., Eindhoven University of Technology, The Netherlands
[2]Institut de Biologie Physico-Chimique, Sorbonne University, CNRS, Paris, France
[3]Station Biologique de Roscoff, Sorbonne University, CNRS, Roscoff, France
Email: A.G.Aribowo@tue.nl

1 Introduction

Photosynthesis is a complex process that strongly influences plant behavior. It involves mechanisms operating on different time scales, ranging from fast dynamics in light harvesting to slow dynamics in the chemical synthesis of organic compounds for plant growth. Moreover, the photosynthetic process exhibits dynamic responses to unfavorable conditions or to its own physiological state. The goals of photosynthesis modeling study is to understand better how photosynthesis responds to light fluctuations and optimize photosynthetic performances accordingly. Although various models have been proposed to describe the influence of oscillating light on photosynthesis, their inherently nonlinear mathematical structures are often too complex for practical applications of regulation of photosynthesis. For example, the nonlinear terms in these models can hinder development and use of standard control design tools for light modulation. Alternatively, system identification methods are a data-driven modeling approach that can capture key system dynamics with reduced complexity, offering a more practical route for investigating and optimizing photosynthetic processes. In this work, we present modeling results demonstrating the application of system identification approaches that capture the input–output dynamics of photosynthetic regulation.[1]

2 Identification of photosynthesis dynamics

The experimental work is conducted with the unicellular green alga *Chlamydomonas reinhardtii* illuminated by oscillating light. Its photosynthetic responses, measured through chlorophyll fluorescence (ChlF), include changes in photochemical quenching (light energy used for photosynthesis) and non-photochemical quenching (light energy dissipated as heat for photoprotection). This illumination is composed of two components of the light signal: a DC (constant) signal and an AC (modulation) signal in terms of amplitude and frequency. To describe their dynamic responses, we employ the frequency-domain-based system identification technique [1] to measure the associated non-parametric and parametric estimations. These estimations of ChlF responses can provide the insights how plants regulate energy flow, showing the associated harmonics in responses to non-linear regulatory feedback, crucial for understanding plant adaptation to fluctuating natural light.

[1]The work is financially funded by the EU-funded DREAM initiative project: HORIZON-EIC-2021-PATHFINDER OPEN project DREAM with the grant agreement no. 101046451.

3 An illustrative example

As an illustration, the application of the Best Linear Approximation (BLA) in [1] is explored to estimate the Frequency Response Functions (FRFs) of the ChlF responses; see further in [2]. In-silico datasets are generated using simulations of the Basic DREAM Model (BDM) [3], with light intensity signals in the form of multisine excitation signals, comprising a set of F frequencies as follows:

$$u(t) = u_{dc} + \sum_{k \in F} A_k \sin(2\pi f_k t + \phi_k), \tag{1}$$

where u_{dc} is the DC light, and A_k, f_k, and ϕ_k are the amplitude, frequency, and phase of the light modulation, respectively. Figure 1 shows the estimation results that distinguish between the FRFs (linear regime) describing the photochemical quenching and the nonlinear distortion associated with the non-photochemical quenching.

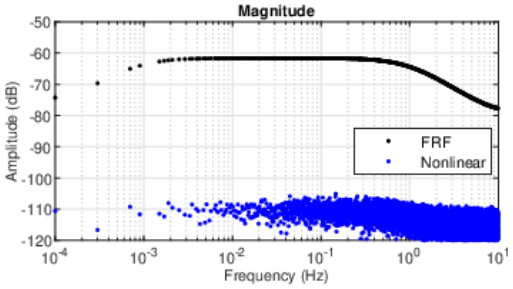

Figure 1: Magnitude spectrum of ChlF responses to multisine light intensity with an amplitude of $38\mu E, m^{-2}, s^{-1}$ and a DC value of $100\mu E, m^{-2}, s^{-1}$.

References

[1] R. Pintelon, J. Schoukens (2012) System identification: a frequency domain approach. *Wiley-IEEE Press, Volume 98, Issue 4, pp 2631–2650.*

[2] C. Portilla et al. (2025) Data-driven modeling of photosynthesis regulation under oscillating light condition: in-silico exploration. *arXiv preprint arXiv:2511.04330.*

[3] D. Fuente et al. (2024) A mathematical model to simulate the dynamics of photosynthetic light reactions under harmonically oscillating light. *Plant Physiology and Biochemistry, Volume 217, 109138.*

Slow neuronal dynamics from fast activating kinetics in conductance-based models

Arthur Fyon*, Julien Brandoit*, Guillaume Drion*, Kevin Jehasse*
*Dept. of Electrical Engineering and Computer Science
University of Liège
Allée de la Découverte 10, 4000 Liège
Belgium
Email: {afyon, jbrandoit, kevin.jehasse, gdrion}@uliege.be

1 Introduction

Neurons exhibit various patterns of electrical activity depending on their functional role. One important pattern is slow pacemaking—regular spontaneous firing at 1–5 Hz. Examples include dopamine neurons in the substantia nigra, which maintain steady dopamine release in the brain, and sinoatrial cells in the heart, which regulate the heartbeat. Generating such slow activity requires a small, steady depolarizing current on the order of tens of picoamperes during the interspike interval. However, the electrophysiological mechanism responsible for this behavior remains highly debated.

For dopamine neurons, hypotheses include sodium leakage channels (unlikely due to their high conductances) or a balanced interplay of larger ion channels (unlikely since blocking individual components does not disrupt pacemaking). For sinoatrial cells, an interaction between two calcium clocks has been proposed (also unlikely, as blocking calcium entry does not eliminate pacemaking).

Recently, a chemical compound, 1-(2,4-xylyl)guanidinium (XG), was found to completely block pacemaking in both cell types [1], suggesting a shared mechanism. XG specifically blocks gating pore currents—tiny currents in the picoampere range that flow within the voltage sensor structure of ion channels rather than through the main pore [2]. While these currents are typically associated with pathological mutations, our findings suggest that distributed non-pathological gating pore currents may be responsible for slow pacemaking. The molecular identity of this ion channel remains unknown, but a complete steady-state current-voltage characteristic has been measured.

2 Results

Using a conductance-based model of dopamine neurons [3], we incorporated this newly characterized current to study its dynamical properties and ability to generate slow pacemaking. The voltage-current relationship of any conductance-based neuron model is expressed as:

$$C\frac{\mathrm{d}V}{\mathrm{d}t} + g_{\text{leak}}(V - E_{\text{leak}}) = -\sum_{\text{ion}\in\mathscr{I}} g_{\text{ion}}(V,t)(V - E_{\text{ion}}) + I_{\text{ext}},$$

where C represents the membrane capacitance, g_{ion} denotes the ion channel conductance (non-negative, gated between 0 for all channels closed and $\bar{g}_{\text{ion}}$ for all channels open), E_{ion} and E_{leak} are the reversal potentials, $\mathscr{I}$ is the index set of intrinsic ionic currents, and I_{ext} is the externally applied current *in vitro* or the combination of synaptic currents. Each ion channel conductance is nonlinear and dynamic, represented by $g_{\text{ion}}(V,t) = \bar{g}_{\text{ion}} m_{\text{ion}}^a(V,t) h_{\text{ion}}^b(V,t)$, where m_{ion} and h_{ion} are gating variables between 0 and 1, modeling the activation and inactivation gates of ion channels, respectively.

Prior models used unphysiological values for sodium and calcium currents to achieve pacemaking. Restoring physiological parameters caused the model to lose pacemaking activity. However, incorporating the new XG-sensitive current recovered pacemaking over a wide parameter range. Importantly, fast activation kinetics (fast opening and closing) are necessary to produce slow pacemaking—a counterintuitive finding confirmed *in vitro* using dynamic clamp experiments.

This work demonstrates how experiments and modeling can work synergistically to reveal novel mechanisms underlying fundamental neurophysiological processes.

References

[1] Fyon, A. et al. (2025). A fast and tiny current as common generator of slow regular pacemaking in brain and heart. *bioRxiv*, 2025.10.30.685563.

[2] Jehasse, K. et al. (2021). The gating pore blocker 1-(2, 4-xylyl) guanidinium selectively inhibits pacemaking of midbrain dopaminergic neurons. *Neuropharmacology*, 197, 108722.

[3] Yu, N., & Canavier, C. C. (2015). A mathematical model of a midbrain dopamine neuron identifies two slow variables likely responsible for bursts evoked by SK channel antagonists and terminated by depolarization block. *J. Math. Neurosc.*, 5, 5.

Towards a dissipativity theory of excitability

Guanchun Tong and Rodolphe Sepulchre
Department of Electrical Engineering (ESAT/STADIUS), KU Leuven
Kasteelpark Arenberg 10 - 3001 Leuven, Belgium
Email: {guanchun.tong, rodolphe.sepulchre}@kuleuven.be

1 Introduction

Excitable systems are primarily characterized by two dynamical features: a *threshold* and a *refractory period*. The threshold corresponds to the system's propensity for all-or-none responses. If the input is below threshold, the system responds passively around the resting equilibrium, while larger inputs crossing the threshold will elicit large-amplitude stereotyped spikes. The refractory period refers to the regime after the spike, in which the system temporarily exhibits a reduced (or absent) ability to generate a new spike. While such a qualitative characterization of excitability has long been acknowledged, a system-theoretic definition remains a challenge in neurobiology and neuromorphic engineering. In this abstract, we propose definitions grounded in the theory of dissipativity [1]: we define the threshold via the concept of *required supply* [3] and the refractory period via the concept of *available storage*. These two concepts form the foundation of dissipativity theory. They admit a formulation in the classical language of optimal control.

2 Definition of threshold via required supply

We focus on a class of systems that admit a lumped circuit realization where a capacitor is in parallel with multiple resistive or memristive elements and voltage sources, all driven by a single port $(v(t), i(t))$. Following KCL, the equation of the circuit is $C\dot{v}(t) = -\sum_{k=1}^{N} i_k(t) + i(t)$, where N is the number of resistive/memristive branches and C is the capacitance. Each current i_k is of the form $g_k(v - E_k)$, where the memductance g_k depends on the past history of the voltage, and the voltage source is E_k.

The *constrained required supply* [2] is the minimum energy that must be supplied to the circuit to reach a certain storage at time $t = 0$:

$$S_{rc}(v) := \inf_{i_{(-\infty,0]}:\, v(0)=v} \int_{-\infty}^{0} i(t)\, v(t)\, dt. \tag{1}$$

Because the external current is determined by v through KCL and branch dynamics, the required supply can be equivalently characterized as:

$$S_{rc}(v) = \frac{1}{2}Cv^2 + \inf_{v_{(-\infty,0]}:\, v(0)=v} \int_{-\infty}^{0} \sum_{k=1}^{N} i_k(t)\, v(t)\, dt. \tag{2}$$

We impose rest boundary conditions $v(-\infty) = 0$ and restrict the search space of $v(t)$ to a finite-dimensional family (e.g., $Ae^{\alpha t}$) to ensure numerical tractability. The *threshold* is defined as a local maximum of $S_{rc}(v)$, which expresses that a higher level of capacitive storage can be reached with less external energy supply.

3 Definition of refractory period via available storage

The natural dual of *required supply* in dissipativity theory is *available storage*. We define the *constrained available storage* $S_{ac}(v)$ as the maximum energy that can be *extracted* from the circuit over the interval $[0, \infty)$. Symmetrically to (2), this splits into the capacitive storage and the maximum energy recoverable from the internal branches:

$$S_{ac}(v) = \frac{1}{2}Cv^2 + \sup_{v_{[0,\infty)}:\, v(0)=v} \left(\int_{0}^{\infty} - \sum_{k=1}^{N} i_k(t)\, v(t)\, dt \right). \tag{3}$$

In the context of the Hodgkin-Huxley model, the refractory period is driven by the potassium channels. After a spike, the persistent opening of these channels creates a sustained high conductance pathway for outward potassium current, pulling the membrane voltage below its resting level. From an energy-based viewpoint, this is the phase where the system is set up to dissipate energy stored in the ionic gradients. Our definition captures this mechanism: we identify the *refractory signature* as a local maximum of $S_{ac}(v)$ (optimized over admissible future voltages, e.g., $Be^{-\beta t}$). A local maximum in the available storage indicates that less energy can be extracted despite a higher level of capacitive storage.

This talk will describe our current research to characterize both the threshold (via required supply) and refractory period (via available storage) of excitable systems using optimization-aided dissipativity analysis.

References

[1] J. C. Willems, "Dissipative dynamical systems part I: General theory", Archive for rational mechanics and analysis, vol. 45, pp. 321-351, 1972.

[2] A. J. van der Schaft, "Cyclo-dissipativity revisited", IEEE Transactions on Automatic Control, vol. 66, no. 6, 2021.

[3] G. Tong, R. Sepulchre, "On the threshold of excitable systems: an energy-based perspective", Proceedings of IEEE Conference on Decision and Control, 2025.

Neuromorphic behaviors

Rodolphe Sepulchre
Center for Dynamical Systems, Signal Processing, and Data Analytics (STADIUS), KU Leuven

1 Introduction

The behavioral approach introduced by Jan Willems [1] regards a dynamical system as a collection of trajectories restricted by a time axis $\mathbb{T}$, a signal space $\mathbb{W}$, and the behavior law $\mathscr{B} \subset \mathbb{W}^{\mathbb{T}}$. The behavior precedes further specifications such as inputs, outputs, latent variables, etc ... The behavioral approach is very useful to clarify system theoretic concepts away from their instantiation in a particular context or application. We use the behavioral approach in an effort to single out what makes a behavior *neuromorphic* and to clarify the distinction between neuromorphic and hybrid behaviors.

2 Mixing the discrete and the continuous

Neuromorphic systems are the building blocks of neuromorphic control, that seeks to imitate neural principles in interconnecting a controller (nervous system) and a plant (body). At the core of the neuromorphic vision lies the fact that interaction in animal nervous systems is *event-based*, whereas it is *trajectory-based* in artificial machines. The *trajectory-based* theory of dynamical systems has created a dichotomy between the continuous and the discrete. Trajectories are *either* continuous variations over time or *discrete* sequences of a finite state machine. In a behavioral setting, this dichotomy is not about the time-axis being $\mathbb{R}$ or $\mathbb{Z}$. It is rather about the signal space $\mathbb{W}$ being either a vector space or a finite set, and about the behavior being either described by differential equations or by logical rules. The dichotomy has led to an increasing separation between continuous behaviors, the focus of regulation, and discrete behaviors, the focus of automation. Cyberphysical systems, that is, systems that interconnect decision-making devices and physical devices result in *hybrid* behaviors. In a hybrid behavior, the signal space is the *union* of a vector space and a finite set, and the behavioral law *concatenates* continuous and discrete laws.

In contrast, we propose to place neuromorphic systems at the *intersection* of the continuous and the discrete. Neuromorphic systems obey the *continuous* differential laws of biophysical circuits inherited from neurophysiology. However, from the excitable nature of neurons, their trajectories consist of *spikes*, that is, *discrete* events. Spikes are continuous voltage variations over time, yet spikes can be counted [2]. This mixed nature is the distinctive property of neuromorphic systems. Neuromorphic systems are *specialized* continuous behaviors in that their trajectories consist of *discrete* events.

3 A behavioral definition of neuromorphic systems

Neuromorphic systems are *physical continuous* behaviors, with time axis $\mathbb{T}$, signal space $\mathbb{W} = \mathbb{R}^n$, and behavior law determined by a specific power-preserving interconnection of electrical port elements. To single out neuromorphic behaviors amongst electrical circuits, we want to further restrict their behavioral description. Our proposal is to associate to the ground behavior, which we assume physical and continuous, a hierarchical sequence of *discrete* behaviors that enable a multi-scale description of the model. This multi-scale description is analog to the wavelet description of a signal: the potential complexity of a ground neuromorphic system is restricted by a hiearchical model which enables coarser but simpler descriptions at coarser scales.

By *discrete* behavior, we simply means a (finite or infinite) discrete-time axis $\mathbb{T} \subset \mathbb{Z}$, a discrete signal space W consisting of a finite alphabet, and a behavioral law $\mathscr{B} \subset \mathbb{W}^{\mathbb{T}}$ that restricts the possible sequences of letters. This is consistent with the aim of Willems to include both differential equations and automata in behavioral theory. The hierarchy of discrete behaviors further restricts the discrete model of a neuromorphic system: events have nested scales, following the hierarchy of a language: words are made of letters, sentences are made of words, etc ... Likewise, bursts are made of spikes, patterns are made of sequences of bursts, etc ....

The talk will illustrate this behavioral description on a neuromorphic snake recently proposed in [3] and discuss ongoing research questions pertaining to a behavioral theory of neuromorphic systems.

References

[1] J. C. Willems, "Paradigms and puzzles in the theory of dynamical systems," in IEEE Transactions on Automatic Control, vol. 36, no. 3, pp. 259-294, March 1991

[2] R. Sepulchre, "Spiking Control Systems," in Proceedings of the IEEE, vol. 110, no. 5, pp. 577-589, May 2022,

[3] Y. Huo, F. Forni, R. Sepulchre. "A Neuromorphic Architecture for Scalable Event-Based Control", arxiv 2511.11924, 2025.

Acknowledgment This work was supported by European Research Council under the Advanced ERC Grant Agreement SpikyControl 101054323.

Mitigating mismatch in neuromorphic circuits using neuromodulation loops

Christian Fernandez Lorden[1], Alessio Franci[1,2]
[1]Neuroengineering Laboratory, University of Liège, Belgium
[2] WEL Research Institute, Wavre, Belgium
email : Christian.FernandezLorden@uliege.be, afranci@uliege.be

1 Introduction

Carver Mead and collaborators envisioned fast, low-power, and efficient computing inspired by biological systems. This vision has been pursued by research groups and companies, leading to the development of multiple neuromorphic chips [1, 2]. In recent years, neuromorphic chips have demonstrated their envisioned high energy-efficiency and have even shown promising results on real-world tasks [3].

A fundamental challenge hindering the deployment of analog neuromorphic chips is device mismatch. Mismatch refers to the inherent variability in circuit parameters introduced during production. Designs that do not mitigate mismatch are impractical to deploy at scale; either each chip must be precisely calibrated during production or design parameters need to be adjusted individually for each chip.

Biological systems face similar variability yet operate reliably despite it, owing to homeostatic and neuromodulatory mechanisms [4]. The role of neuromodulation in real neural circuits is crucial, as blocking it leads to the collapse of the circuit's function [5]. Inspired by biology, we propose using neuromodulation-based regulation loops to stabilize the behavior of neuromorphic circuits in the presence of mismatch. We demonstrate this approach using simulations performed with our Neuromorphic Control Toolbox [6, 7], which provides all the necessary blocks (neurons, synapses, neuromodulation) as well as tools for mismatch emulation.

2 Counteracting Mismatch in a Half-Center Oscillator

To illustrate a regulatory loop using neuromodulation, we study the example of a half-center oscillator (HCO) affected by mismatch. In simulation, mismatch is applied to all neuronal parameters, as well as to all parameters of the neuromodulation loop, since the loop should also be implemented on chip. The objective of the neuromodulation loop is to balance the HCO neurons, ensuring that each neuron's burst duration is equal.

Each neuron participates in a neuromodulation loop that adjusts the synaptic strengths based on the bursting activity of both neurons. When the HCO is balanced, the activities of the two neurons are equal but in opposite phases, and their modulatory effects cancel each other out. When there is an imbalance, neuronal activities no longer cancel, leading to a gradual drift in synaptic strengths until balance is restored.

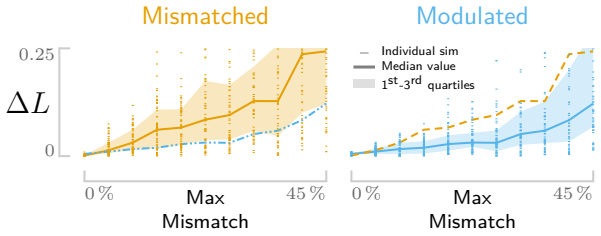

Figure 1: Comparison of the relative difference in burst duration (ΔL) between the two neurons of an HCO circuit as a function of the mismatch strength (maximum per-parameter deviation from nominal, in %) with and without a neuromodulation loop.

Figure 1 shows that introducing mismatch perturbs the HCO and results in unequal burst durations. However, it also demonstrates that activating the neuromodulation loop reduces the deviation from the balanced state. These preliminary results indicate the potential effectiveness of using neuromodulatory feedback for stabilizing neuromorphic circuits.

References

[1] O. Richter et al., "DYNAP-SE2: a scalable multi-core dynamic neuromorphic asynchronous spiking neural network processor," Neuromorphic Computing and Engineering, 2024

[2] L. Mendolia et al., "A Neuromodulable Current-Mode Silicon Neuron for Robust and Adaptive Neuromorphic Systems," arXiv, 2025

[3] J. Zhao et al., "Closed-Loop Spiking Control on a Neuromorphic Processor Implemented on the iCub," IEEE Journal on Emerging and Selected Topics in Circuits and Systems, 2020

[4] J.-M. Goaillard and E. Marder, "Ion Channel Degeneracy, Variability, and Covariation in Neuron and Circuit Resilience," Annual Review of Neuroscience, 2021

[5] J. Golowasch, "Neuromodulation of central pattern generators and its role in the functional recovery of central pattern generator activity," Journal of Neurophysiology, 2019

[6] C. Fernandez Lorden et al., Neuromorphic Control Systems. [Online]. Available: https://nl.mathworks.com/matlabcentral/fileexchange/180432-neuromorphic-toolbox-for-control

[7] C. Fernandez Lorden, "Neuromorphic Control Systems: A MATLAB® toolbox for neuromorphic control analysis and design," arXiv, 2026

Input-driven timing reliability in mismatched neuromorphic neurons

Loris Mendolia
Montefiore Institute, University of Liège
Allée de la Découverte 10, 4000 Liège, BE
Email: lmendolia@uliege.be

Alessio Franci
Montefiore Institute, University of Liège
Allée de la Découverte 10, 4000 Liège, BE
Email: afranci@uliege.be

Abstract

Achieving reliable event timing for reliable sensory processing and motor coordination is a central requirement for neuromorphic systems intended to operate with nonlinear, low-power, and mismatch-prone analog hardware. When driven by constant inputs, excitable neuromorphic neurons typically enter autonomous oscillatory regimes, where device mismatch induces variability in firing rates, phases, and waveform shape, resulting in unreliable event timing across trials and across neurons. This unreliability arises because, under constant stimulation, neurons operate as autonomous oscillators whose limit-cycle dynamics are sensitive to parametric variations. In this work, inspired by previous biological and computational findings [1, 2], we demonstrate experimentally that time-varying inputs can enhance event timing reliability, not only across single-neuron trials, but also in a population of heterogeneous silicon neurons, overcoming hardware mismatch beyond careful circuit design and parameter tuning.

Using a 16-neuron array of mixed-feedback analog silicon neurons capable of both spiking and bursting excitability, and exhibiting significant transistor mismatch, we compare neural responses to constant (DC) stimulation and to repeated band-limited stochastic inputs of identical mean amplitude. Timing reliability is derived from SPIKE-distance metrics [3], across repeated trials of individual neurons and throughout the neuron population.

While DC inputs produce autonomous oscillations with high trial-to-trial variability and no inter-neuron coherence, fluctuating inputs drive neurons to behave as open, input-triggered excitable systems in which events are triggered by shared input fluctuations. In this regime, event timing becomes highly reproducible across trials, and better aligned across mismatched neurons, despite remaining heterogeneity in firing rates and waveform details. Thus, time-varying inputs reduce the mismatch-induced variability in event timing at the population level without requiring calibration, coupling, or network interactions.

We further discuss the existence of distinct optimal input timescales for spiking and bursting neurons, by observing that timing reliability depends on the input bandwidth. Spiking and bursting neurons both display an optimum input frequency for maximum reliability, but bursting neurons require significantly lower frequencies to obtain better timing reliability. This behavior is consistent with a trigger-based mechanism, where input variations must correspond to the intrinsic timescales of the excitable dynamics, while faster fluctuations will interfere with event formation.

These results demonstrate that appropriate input design can transform heterogeneous neuromorphic neurons from unreliable autonomous oscillators into robust event-triggered systems. More broadly, they highlight how excitable dynamics can be exploited to stabilize event-based computation under parametric uncertainty, offering a robust alternative to rate-based operation and explicit mismatch compensation in analog neuromorphic hardware.

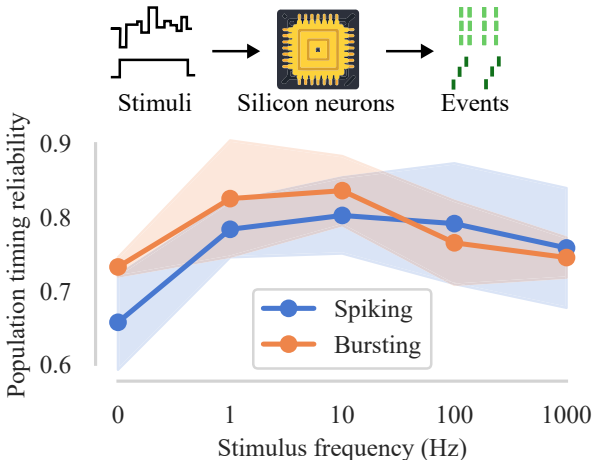

Figure 1: Event timing reliability increases across mismatched neurons under time-varying stimulation and shows different ideal input bandwidths for spiking and bursting. For each stimulus, population timing reliability is defined as $R = 1 - \overline{D}_{i \neq j}$ where $D_{i,j}$ is the SPIKE-distance averaged across pairs of trials of neurons i and j. Shading shows the standard deviation across neuron pairs.

References

[1] Z. F. Mainen and T. J. Sejnowski, "Reliability of Spike Timing in Neocortical Neurons," Science, vol. 268, no. 5216, pp. 1503–1506, Jun. 1995.

[2] T. M. Kirby et al., "Reliability of Event Timing in Silicon Neurons," arXiv:2112.14134, Dec. 2021.

[3] T. Kreuz et al., "Time-resolved and time-scale adaptive measures of spike train synchrony," Journal of Neuroscience Methods, vol. 195, no. 1, pp. 92–106, Jan. 2011.

Joint Optimization of Hybrid Battery Sizing and Parameterized Energy Management System

Maxime Monsieur, Sam Weckx and Ankit Surti
MotionS, Flanders Make VzW
3920 Lommel, Belgium
maxime.monsieur@flandersmake.be

1 Motivation

Hybrid battery energy storage systems (HBESS), which combine high-energy (HE) and high-power (HP) cells, can outperform single-technology battery systems, particularly in applications with moderate power-to-energy ratios [1]. However, their design is inherently more complex due to the strong coupling between battery sizing and the energy management strategy (EMS). While fully optimal power-split strategies can be derived during offline sizing, these strategies are generally difficult to implement in practice unless their behavior can be captured by an implementable control structure or reproduced by an online optimization scheme with comparable performance. To address this, we focus on parameterized EMS architectures with fixed structures, optimizing their parameters jointly with battery sizing. We investigate how the choice of EMS architecture limits the achievable power split and, in turn, affects the optimal HBESS sizing.

2 Approach

HBESS sizing, including the parametrized EMS that splits the power demand P between the HE and HP cells (respectively P_{HE} and P_{HP}), is formulated as a mixed-integer nonlinear problem using CASADI, as described in [1]. The optimization is performed for a user-defined power demand profile with predefined cell chemistries, and the discrete sizing problem is solved iteratively via a continuous relaxation, resulting in a (locally) optimal solution. The objective is to minimize the initial battery pack investment cost, neglecting operational costs.

The EMS must satisfy the power balance $P[k] = P_{HE}[k] + P_{HP}[k]$ for all time steps k. Three EMS strategies are considered: (i) *optimal*, where the solver has full freedom to choose P_{HE} and P_{HP} at each step, representing an idealized, practically unattainable benchmark; (ii) *threshold*, where HE cells supply power up to a limit $P_{\text{threshold}}$, with HP providing the remainder, and $P_{\text{threshold}}$ is treated as an optimization variable; (iii) *filter*, where HE power is a low-pass filtered version of the demand, $P_{HE}[k] = a \cdot P[k] + (1-a) \cdot P_{HE}[k-1]$, with the filter coefficient a optimized alongside battery sizing.

3 Results

Figure 1 compares the HBESS sizing results against the lowest-cost monotype solution for 2472 demand profiles, with peak power ranging from 200 to 6000 kW and energy content from 200 to 2000 kWh, thereby spanning a broad range of power-to-energy ratios. The cost rises on average 5.9% and 14.5% for threshold and filter EMS, respectively, relative to the optimal split, reducing HBESS benefits. Future work will explore parameterized EMS designs that better approximate the optimal power split while implicitly respecting constraints and secondary objectives, such as charge sustainability.

References

[1] S. Weckx, A. Surti, Z. Tao, "Evaluating Hybridization Potential Using Load Profile Metrics: A Rule-of-Thumb Approach," Batteries, 2025.

Acknowledgment This research was supported by Flanders Make, the strategic research center for the manufacturing industry, in the framework of the OptiBATT_IRVA project.

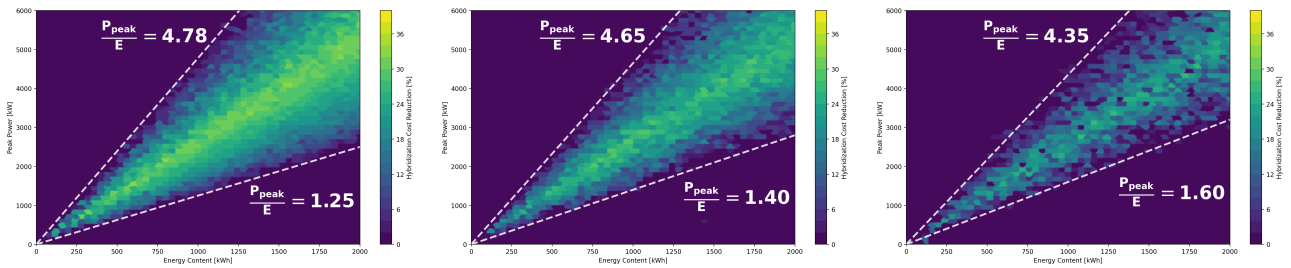

(i) optimal (ii) threshold (iii) filter

Figure 1: Cost reduction of HBESS sizing for different EMS architectures compared with the cheapest monotype solution.

Equivalent circuit model parameter identification for lithium-ion batteries and adaptation to aging

Ali Rahdarian
Depart. SAAS
Université libre de Bruxelles
CP 165/55, 50 Av.F.D. Roosevelt
B-1050 Brussels
Belgium
Email: ali.rahdarian@ulb.be

Michel Kinnaert
Depart. SAAS
Université libre de Bruxelles
CP 165/55, 50 Av.F.D. Roosevelt
B-1050 Brussels
Belgium
Email: michel.kinnaert@ulb.be

1 Introduction

Compared with other commonly used batteries, lithium-ion batteries continue to draw more attention as a promising energy storage technology. In most applications, a battery system consists of the battery and the battery management system (BMS), which is made up of hardware and software. A BMS controls the charging and discharging of the battery while guaranteeing reliable and safe operation. A BMS is also responsible for other tasks, such as cell balancing and thermal management of the battery pack. Designing BMS software for Li-ion batteries relies on a model that can describe the battery dynamics. For this purpose, equivalent circuit models (ECMs) or physics-based models can be used. ECMs are more popular due to their simplicity and lower computational cost compared to physics-based models.

Using the model, the State of charge (SOC) of the battery can be monitored, which is one of the most important tasks of a BMS. SOC represents the remaining energy of the battery and is essential for preventing overcharge and overdischarge. In recent years, research has been focused on model-based SOC estimation approaches based on ECMs or physics-based models. The main challenge for maintaining accurate SOC estimation over the battery's lifetime is aging [1]. As a result of aging, models that have been identified based on the fresh cell data lose their accuracy. ECM parameters, including open circuit voltage (OCV), capacity, and overpotential parameters, change. This decreases the accuracy of model-based SOC estimation methods such as the extended Kalman filter (EKF). The impact of aging on the accuracy of SOC estimation has been studied extensively in the literature. In addition, it has been shown that to maintain an accurate SOC estimate over time, model parameters must be updated upon aging. For instance, in [2], the authors analysed the impact of aging on the accuracy of the SOC estimation using EKF. It was seen that when the battery ages, if parameters are not updated, the estimator's performance deteriorates significantly. However, there are significant challenges in updating ECM parameters over time. In practice, diagnostic tests such as hybrid pulse power characterization (HPPC) and OCV characterization tests can not be performed during normal battery operation, as they are time-consuming. Therefore, parameter updating must be based on data obtained from BMS sensors, which exhibit lower precision compared to testers. In response to these challenges, multiple studies have been focused on updating ECM parameters, including capacity. In [3], a method based on weighted total least squares (WTLS) was developed for estimating capacity from operational data. By taking into account the impact of noise on current and uncertainity in SOC estimate, the method provides an unbiased estimate of the capacity. Owing to its low computational complexity, the method enables real-time capacity estimation during normal battery operation. However, updating capacity alone is insufficient to fully compensate for aging effects, as the OCV characteristics and overpotential parameters also evolve with aging. In most existing studies, the OCV characteristics are assumed to be independent of aging; however, this assumption does not hold in practice. The aim of our work is to develop a straightforward method to identify ECM parameters from fresh cell data and update all the ECM parameters, including the OCV characteristics, using the operational data. For this purpose, the system identification for the fresh cell is formulated as a linear regression problem and solved using the data obtained from the tester. As the battery ages, parameter update is carried out using total least squares to account for the measurement inaccuracies and noise in the BMS sensor data.

References

[1] J.P Rivera-Barrera, N. Muñoz-Galeano, H. O Sarmiento-Maldonado, "SOC Estimation for Lithium-ion Batteries: Review and Future Challenges," Electronics, vol. 6, no. 4, 2017.

[2] X. Hu, S. Li, H. Peng, and F. Sun, "Robustness analysis of State-of-Charge estimation methods for two types of Li-ion batteries," Journal of power sources, vol. 217, pp. 209–219, 2012.

[3] G. L. Plett, "Recursive approximate weighted total least squares estimation of battery cell total capacity," vol. 196, pp. 2319-2331, 2011.

Load-based Control Performance Testing Method for Air Conditioning Systems

Cuiling Wang[12], Baolong Wang[2], Hussain Kazmi[1]
1 Department of Electrical Engineering, KU Leuven, Leuven, Belgium
2 Department of Building Science, Tsinghua University, Beijing, China
Email:cuiling.wang@kuleuven.be

1 Introduction

The dynamic performance of advanced control methods for air-conditioning systems is commonly evaluated using numerical simulations or field experiments. However, field experiments are difficult to repeat and time-consuming due to weather and occupancy variations, while numerical simulations often fail to accurately represent the complex behavior and internal control logic of real systems. This paper presents a laboratory-based dynamic performance testing platform for air-conditioning control methods, which enables repeatable and low-cost performance evaluation of control strategies on real systems under diverse operating scenarios, such as different building and climate types. A quantitative assessment of the demand response performance of an advanced controller (model predictive control method, MPC) is conducted using the platform.

2 Method overview

The testing platform is based on the dynamic load-based testing method, which is commonly used for experimental evaluation of the thermal performance of air-conditioning systems (AC). Building upon this principle, the platform integrates Internet-of-Things technologies and dynamic simulation to enable laboratory-based performance testing of dynamic control performances of AC systems under realistic and repeatable operating conditions. The platform consists of three core modules— the building, the AC system, and the controller—coordinated through well-defined signal exchange mechanisms. As shown in Fig.1 , the building is emulated to simulate the dynamic indoor air temperature, while the AC system is a real physical system. A cooling capacity measurement module is employed to monitor the real-time cooling output of the AC system, which is converted into digital signals and used as inputs to the building emulator. When the ambient temperature surrounding the indoor unit matches the simulated indoor air temperature and the cooling load is known, the system performance can be regarded as equivalent to that observed under real building operating conditions. To achieve this equivalence, a condition generator is used to ensure that the air temperature within the chamber housing the indoor unit continuously follows the simulated indoor air temperature, while the air temperature in the chamber containing the outdoor unit corresponds to the outdoor weather conditions. Advanced controller are implemented on a supervisory computer to determine optimal AC system operation. During real-time operation, the controller communicates bidirectionally with the AC system to acquire system states and issue control commands. For more complex model-based controller, real-time data communication with the building emulator is required.

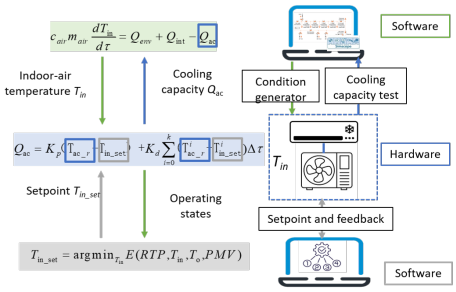

Figure 1: Schematic of the testing platform operation.

3 Results overview

Under dynamic operating conditions, the developed platform demonstrates good performance in terms of regulation speed, steady-state error, and response delay, as shown in Fig. 2. In addition, as shown in Fig. 3, the platform effectively captures the dynamic power variations of the AC system under MPC control.

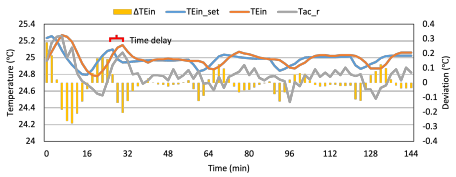

Figure 2: Platform performance under dynamic conditions.

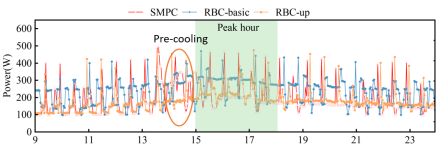

Figure 3: Dynamic control performances of MPC.

Stochastic Model Predictive Control for Energy Management System of Grid with Hybrid Energy Storage System

Taranjitsingh Singh, Boulaid Boulkroune, Maxime Monsieur and Chris Ganseman
Flanders Make, Belgium
firstname.lastname@flandersmake.be

1 Motivation

Hybrid Energy Storage Systems (HESS) integrated with renewables, grid and diverse technologies offer improved efficiency and resilience but face operational challenges due to the uncertainty in demand, renewable energy generation and market price conditions. Conventional strategies like Standard Rule-Based Control (RBC) often underperform when forecasts deviate. deterministic Model Predictive Control (MPC) provides a structured approach for grid Energy Management System (EMS), yet this strategy ignores uncertainty, thereby risking inefficiency and constraint violations.

This work proposes a scenario-based Stochastic MPC (SMPC) framework for a hybrid system with grid connection, renewable sources, and two energy storage units, namely, High-Energy (HE) and High-Power (HP) batteries. The objective is to maximize the renewables penetration and battery usage thereby reducing the reliance on grid power while respecting system and battery constraints.

2 Approach

The Optimization problem for EMS is formulate over a finite horizon MPC that dispatches battery charge/discharge powers along with grid power while enforcing operational constraints.

Demand uncertainty is represented by a finite set of $\mathscr{S}$ demand scenarios over a horizon N. The Stochastic MPC (SMPC) problem (1) minimizes the expected grid usage with an input penalty on battery dispatch powers. To ensure implementability, the non-anticipativity constraint is considered, where the first control move is identical across all scenarios, while future inputs and state trajectories may depend on the scenario [1].

$$\min_{\mathbf{u}} \ \frac{1}{S}\sum_{s=1}^{S}\sum_{t=0}^{N-1}\left(\left\|p^{(s)}_{\text{grid},t}\right\|_2^2+\lambda_\Delta\left\|\Delta p^{(s)}_{\text{HE\&HP},t}\right\|_2^2\right)$$
$$\text{s.t. power balance, storage dynamics, operational bounds, terminal SoC constraint, and non-anticipativity.} \tag{1}$$

3 Results

For comparison, we consider a standard RBC [2] and deterministic MPC that uses single nominal demand forecast.

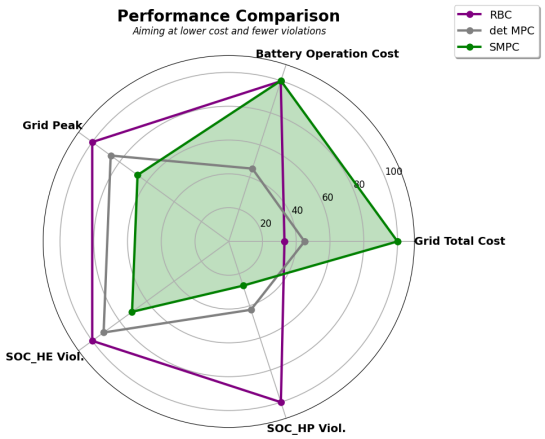

Figure 1: Performance comparison of RBC, MPC, and SMPC.

Figure 1 compares RBC, deterministic MPC, and SMPC over cost, grid peak, and SoC constraint violations. Deterministic MPC achieves low grid peak but results in increased SoC violations under uncertainty. In contrast, SMPC reduces SoC violations compared to deterministic MPC by enclosing against uncertain demand scenarios (by $\approx 22\%$ violations in HE and 35% in HP and by $\approx 23\%$ peak import) while maintaining competitive total cost with respect to RBC.

References

[1] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, "The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations," *Automatica*, vol. 50, no. 12, pp. 3009–3018, 2014.

[2] M. U. Safder, M. A. Hossain, M. J. Sanjari, and J. Lu, "Rule-based energy management system for autonomous voltage stabilization in standalone dc microgrid," *Energy Science & Engineering*, vol. 12, no. 10, pp. 4278–4296, 2024.

Acknowledgment This work was financially supported by the European Union's Horizon Europe Research and Innovation Program through HiHELIOS Project under Grant Agreement No 101137626.

Control of Shared Battery Energy Storage Systems in Renewable Energy Communities: a Convex Optimization Approach

Alessandro Cicalese
ENTEG
University of Groningen
Nijenborgh 4
9747 AG Groningen
The Netherlands
Email: a.cicalese@rug.nl

Francesco Lo Iudice
DIETI
University of Naples Federico II
via Claudio 2
80125 Napoli NA
Italy
Email: francesco.loiudice@unina.it

1 Introduction

Renewable Energy Communities (RECs) are emerging as a key solution for enabling decentralized and sustainable energy systems through local production and sharing of renewable energy. The effective operation of such communities strongly relies on Shared Battery Energy Storage Systems (SBESS), which allow the mitigation of renewable intermittency and the improvement of economic efficiency by increasing self-consumption.

This abstract presents a convex optimization framework for the optimal control of SBESS in residential RECs. The proposed approach is designed to ensure computational efficiency, physical implementability, and scalability, while complying with realistic operational and market constraints.

2 Problem formulation

The energy management problem is formulated as a finite-horizon optimal control problem aimed at minimizing the total energy cost of the community under time-of-use tariffs, renewable incentives, and feed-in mechanisms. The optimization model accounts for battery state-of-charge dynamics, power limits, and energy balance constraints.

Under realistic assumptions on energy prices, a convex reformulation of the original problem is derived. In particular, it is shown that the renewable charging decision can be fixed without loss of optimality, ensuring maximum renewable self-consumption while preserving the convex nature of the problem. This allows the avoidance of binary variables and significantly reduces computational complexity.

3 Convex reformulation and temporal decomposition

It is analytically proven that the optimal management of the SBESS over multiple consecutive days is equivalent to solving a sequence of independent single-day optimization problems. This property enables a substantial reduction in computational burden, making the approach suitable for long-term simulations and real-time implementations.

Moreover, the proposed formulation guarantees the existence of physically implementable solutions by avoiding simultaneous charging and discharging of the storage system, which is a physically infeasible operating condition.

The convex reformulation of the proposed optimization problem also enables its decomposition within a distributed computational framework. In particular, the absence of non-convex coupling constraints allows the original centralized problem to be partitioned into smaller subproblems that can be solved in parallel, while preserving global optimality. This feature is especially relevant for large-scale Renewable Energy Communities, where computational scalability and privacy requirements motivate the adoption of distributed optimization architectures.

4 Numerical validation

Numerical validation is conducted using the AUSGRID dataset, which provides real residential consumption and photovoltaic generation data from Australian households. Simulations consider a community of multiple prosumers operating under realistic time-of-use tariffs and battery parameters. The results demonstrate improved computational efficiency, enhanced renewable energy utilization, and a reduction in total community energy costs.

References

[1] E. L. Ratnam, S. R. Weller, C. M. Kellett, and A. T. Murray, "Residential load and rooftop PV generation: An Australian distribution network dataset," *International Journal of Sustainable Energy*, vol. 36, no. 8, pp. 787–806, 2017.

[2] L. Glielmo, E. Musicò, C. Ancona, and F. Lo Iudice, "An optimal control approach for enhancing efficiency in renewable energy communities," *IEEE Control Systems Letters*, vol. 8, pp. 3039–3044, 2024.

[3] A. S. Alahmed, M. J. Khojasteh, and M. C. Deza, "A decentralized market mechanism for energy communities under operating envelopes," *IEEE Transactions on Control of Network Systems*, vol. 12, no. 1, pp. 313–324, Mar. 2025.

Krasovskii Passivity-Based Control for Grid-Forming Inverters under Matching Control

Xinjiang Cai
Engineering and Technology institute Groningen
University of Groningen
9747 AG, Groningen, the Netherlands
Email: x.cai@rug.nl

Michele Cucuzzella
Engineering and Technology institute Groningen
University of Groningen
9747 AG, Groningen, the Netherlands
Email: m.cucuzzella@rug.nl

1 Introduction

Grid-forming (GFM) inverter is a key enabler for converter-dominated grids, where inverters must autonomously establish voltage and frequency while ensuring stability under network uncertainty and weak damping. Matching control [1] has emerged as a principled, energy-based GFM paradigm, which leverages port-Hamiltonian structure, together with energy shaping and damping injection, to replicate synchronous-machine-like behavior. However, a practical limitation is that available sufficient stability certificates can become conservative when parasitic dissipation is small, which is common in modern hardware. To improve robustness with minimal architectural disruption, we investigate krasovskii passivity as a control-design asset: a passivity notion defined on an extended system with input derivatives, enabling dynamic controllers that systematically manufacture closed-loop dissipation.

2 Methodology

Consider a single grid-forming inverter with DC-link capacitor and AC-side filter, which is shown in Fig. 1. A Krasovskii passivity-based control (K-PBC) method [2, 3] is tailored for the matching-controlled inverter model by representing it as an input-affine system.

$$\dot{x} = f(x,u) = f_0(x) + g(x)u,$$

exposing a port (u,y) with $u = [i_{dc}; i_l]$, collecting the DC-side current injection and the AC-side load/network current, and $y = [v_{dc}; v]$, collecting the DC and AC terminal voltages. Following the Krasovskii framework, an extended system is constructed by augmenting the plant with

$$\dot{x} = f(x,u), \quad \dot{u} = u_K,$$

and select a Krasovskii storage that is quadratic in the state derivative

$$S_K(x,u) = \frac{1}{2}\dot{x}^\top M \dot{x}, \quad M = \mathrm{diag}(C_{dc}, L\mathbb{I}_2, C\mathbb{I}_2).$$

Differentiating S_K along the extended system leading to the Krasovskii passivity inequality

$$\dot{S}_K \le u_K^\top y_K, \quad y_K := g(x)^\top M \dot{x} = \frac{\partial S_K}{\partial u},$$

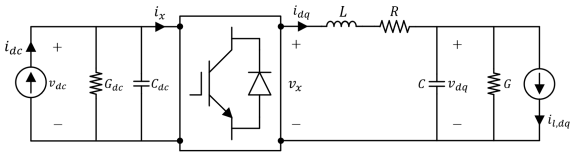

Figure 1: Circuit diagram of the investigated inverter

under a standard sufficient condition $\mathrm{sym}(MF) \preceq 0$, which can be checked via a Schur-complement bound and turned into an explicit, operating-region-dependent inequality on the dissipation parameters.

Once Krasovskii passivity is certified, a K-PBC controller could be designed. To enforce tracking, the passive damping interconnection could be augmented with two integrators

$$\dot{\xi} = v_{dc} - v_{dc,ref}, \quad \dot{\nu} = r - r_{ref},$$

with r being the output voltage of the inverter. Under the established framework, the closed-loop system will drive $v_{dc} \to v_{dc,ref}$ and $r \to r_{ref}$ simultaneously.

3 Outcomes

(1) **Closed-loop stability and robustness** under weak physical damping by injecting dynamic dissipation through K-PBC.
(2) **Frequency and voltage regulation compability** by integrating K-PBC with matching control, while preserving the passivity-based stability under realistic operating conditions.

References

[1] C. Arghir, T. Jouini, and F. Dörfler, "Grid-forming control for power converters based on matching of synchronous machines," *Automatica*, vol. 95, pp. 273–282, Sep. 2018.

[2] Y. Kawano, K. C. Kosaraju, and J. M. A. Scherpen, "Krasovskii and shifted passivity-based control," *IEEE Trans. Autom. Control*, vol. 66, no. 10, pp. 4926–4932, Oct. 2021.

[3] K. C. Kosaraju, M. Cucuzzella, J. M. A. Scherpen, and R. Pasumarthy, "Differentiation and passivity for control of Brayton-Moser systems," *IEEE Trans. Autom. Control*, vol. 66, no. 3, pp. 1087–1101, Mar. 2021.

Guaranteed Safe Navigation in Polytopic Free-Space: a Control Barrier Function Approach

Dries Dirckx, Alejandro González-García, Jan Swevers, Wilm Decré
MECO Research Team, KU Leuven
Flanders Make @ KU Leuven
Celestijnenlaan 300, B-3001 Heverlee, Belgium
dries.dirckx@kuleuven.be

1 Introduction

Autonomous navigation with provable safety guarantees is a key cornerstone of succesful robot integration in a complex world. Currently, discrete Control Barrier Functions (DCBF) are a popular technique to ensure this safety. A recent method [1] employs DCBFs for polytope-to-polytope obstacle avoidance in a Model Predictive Control (MPC) scheme. However, this method requires modeling the occupied space as a set of polytopes, which leads to non-convex constraints and high computational complexity for complex environments. In this work, a novel free-space DCBF formulation is proposed, built instead on a decomposition of the free space as a set of convex polytopes. This results in a convex DCBF formulation that more favourably scales with the complexity of the environment.

2 Free-Space Control Barrier Function

Consider a discrete-time dynamical system $\boldsymbol{x}_{k+1} = \boldsymbol{f}(\boldsymbol{x}_k, \boldsymbol{u}_k)$, with $\boldsymbol{x}_k \in \mathbb{X} \subset \mathbb{R}^n$ the state at time step k and $\boldsymbol{u}_k \in \mathbb{U} \subset \mathbb{R}^m$ the control input. The safe set $\mathbb{C}$ is defined as the 0-superlevel set $\mathbb{C} = \{\boldsymbol{x} \in \mathbb{X} : h(\boldsymbol{x}) \geq 0\}$ of a continuously differentiable function h. A function h is a discrete CBF for such a system if it satisfies the following condition:

$$\forall \boldsymbol{x}_k \in \mathbb{C}, \exists \boldsymbol{u}_k \in \mathbb{U} : h(\boldsymbol{f}(\boldsymbol{x}_k, \boldsymbol{u}_k)) \geq \gamma_k h(\boldsymbol{x}_k), \qquad (1)$$

with $\gamma_k \in [0,1]$ [1]. The inner distance between two polytopes $\mathbb{F}$ and $\mathbb{P}$, representing the free space and robot respectively, is visualised in Figure 1 and described as:

$$d_m(\mathbb{F}, \mathbb{P}) = \min_{i,j} \{\boldsymbol{w}_i^T \boldsymbol{p}_j + b_i \mid \boldsymbol{p}_j \in \mathbb{P},\ \boldsymbol{w}_i \in W_{\mathbb{F}}\}, \qquad (2)$$

where $W_{\mathbb{F}}$ is the set of inward pointing normals of the faces of $\mathbb{F}$ and b_i the corresponding offsets. The minimum function is however not continuously differentiable, making it unsuitable as a CBF. As such, the smooth Log-Sum-Exp (LSE) function approximation is adopted, with α defining the approximation accuracy, resulting in the following convex DCBF formulation for navigation in a polytopic free space:

$$d_C(\mathbb{F}, \mathbb{P}) = d_m - \frac{1}{\alpha} \log \left[\sum_{i,j} \exp\left(-\alpha(\boldsymbol{w}_i^T \boldsymbol{p}_j + b_i - d_m)\right) \right].$$

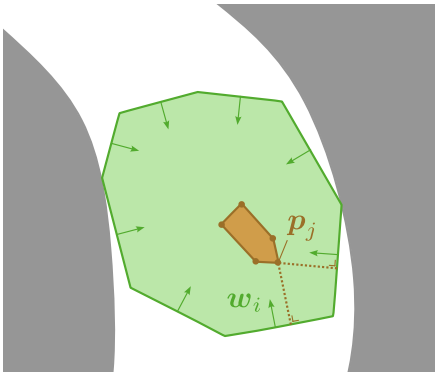

Figure 1: Convex free-space decomposition and the distance between inner ($\mathbb{P}$) and outer ($\mathbb{F}$) polytope.

	N_v	4	8	12	16
dCBF	t_{wall} [ms]	4.01	11.42	30.15	59.96
LSE-CBF	t_{wall} [ms]	3.68	5.67	8.84	8.85

Table 1: Scalability analysis of both formulations.

Substituting d_C as h in Equation 1 ensures that the robot $\mathbb{P}$ remains inside the free space polytope $\mathbb{F}$. Note that d_C remains a scalar function, independent of the number of faces of $\mathbb{F}$ and vertices of $\mathbb{P}$, creating only a single MPC constraint.

3 Experiments

Table 1 shows initial results from a simple point-to-point problem with an increasing amount of vertices N_v. For DCBF, the vertices reflect the amount of occupied polytopes, with each polytope modelled as a rectangle. This shows that the proposed CBF scales better with more complex environment models due to the scalar nature of d_C.

References

[1] Akshay Thirugnanam, Jun Zeng, Koushil Sreenath, "Safety-Critical Control and Planning for Obstacle Avoidance between Polytopes with Control Barrier Functions", IEEE International Conference on Robotics and Automation (ICRA), 2022.

Dual-arm cooperative control under state and input constraints

Zemerart Asani[1,2], Mohayad Omer[1], Maxime Martin[2], Greet Van de Perre[1],
Bram Vanderborght[1], Emanuele Garone[2]
Brubotics, Vrije Universiteit Brussel, and IMEC, Belgium[1]
SAAS, Université Libre de Bruxelles, Belgium[2]
Email: zemerart.asani@vub.be

1 Introduction

Dual-arm cooperative manipulation enables the transportation and placement of heavy rigid objects beyond the payload capability of a single manipulator. However, safe coordination is challenging due to the rigid grasps and the presence of strict physical limits (joint range, velocity, and torque bounds). Classical cooperative control strategies typically do not explicitly enforce such constraints, while optimization-based constrained control can be computationally demanding. In this work, we propose a computationally efficient constrained control scheme based on the Explicit Reference Governor (ERG).

2 Modeling and control

We consider two 7-DoF robotic arms rigidly grasping a common object. The constrained dynamics are written as [1]

$$\begin{cases} M(q)\ddot{q} + C(q,\dot{q})\dot{q} + g(q) = S\tau + E^{\top}(q)\lambda, \\ h(q) = 0. \end{cases} \tag{1}$$

The ERG ensures constraint satisfaction by continuously manipulating the derivative of the applied reference

$$\dot{q}_v = \rho(q_v, r)\Delta(q, \dot{q}, q_v), \tag{2}$$

where $\rho(q_v, r)$ is the Navigation Field (NF) and $\Delta(q, \dot{q}, q_v)$ is the Dynamic Safety Margin (DSM)[2]. The navigation field is composed of an attraction ρ^{att} and repulsion term ρ^{rep}. The DSM is computed by simulating the system forward for a given time horizon T [2].

The dual-arm robotic system is subject to the following constraints:

- **Joint range constraints:**
$$q_{\min,i} \leq q_i \leq q_{\max,i}, \qquad \forall i = 1,\ldots,2n. \tag{3}$$
- **Joint velocity constraints:**
$$\dot{q}_{\min,i} \leq \dot{q}_i \leq \dot{q}_{\max,i}, \qquad \forall i = 1,\ldots,2n. \tag{4}$$
- **Actuator saturation:**
$$\tau_{\min,i} \leq \tau_i \leq \tau_{\max,i}, \qquad \forall i = 1,\ldots,2n. \tag{5}$$

3 Numerical results

Simulations on a dual-arm system demonstrate that the proposed ERG-based cooperative architecture achieves safe cooperative manipulation while satisfying the constraints.

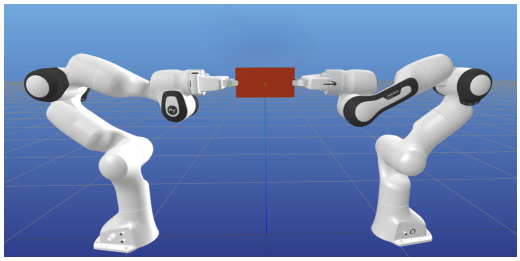

Figure 1: Dual arm robotic system manipulating a load.

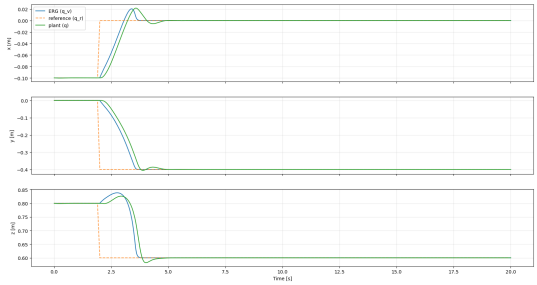

Figure 2: Payload position and ERG commanded reference.

Acknowledgments: This work was supported by the Research Foundation - Flanders (FWO) through the PhD fellowship grant numbers 1SHBX24N, 1S47925N, and by Innoviris and FARI (BrickieBots 2.0), and by the Artificiele Intelligentie Vlaanderen from the Flemish Government.

References

[1] Zemerart Asani, Michele Ambrosino, Bram Vanderborght, and Emanuele Garone. Implementation of a heterogeneous multi-robot system for a construction task. *IFAC-PapersOnLine*, 56(2):3373–3378, 2023.

[2] Zemerart Asani, Michele Ambrosino, Bram Vanderborght, and Emanuele Garone. Towards safer and more efficient construction-the development of a multi-robotic system for heavy block assembly. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13–16. The International Association for Automation and Robotics in Construction ..., 2024.

Slip Avoidance in Grasping Using Control Barrier Functions

Mohaddese Hazini
University of Twente
m.hazini@utwente.nl

Janset Dasdemir
University of Twente
j.dasdemir@utwente.nl

Wouter Hakvoort
University of Twente
w.b.j.hakvoort@utwente.nl

1 Introduction

A successful interaction between a manipulator and its environment is achieved when both performance and safety are satisfied. While notions of stability often guide performance, safety can be addressed through the concept of invariance. If a safe set is defined as the super-level set of a continuously differentiable function h, $\mathscr{C} = \{\zeta \in D \subset \mathbb{R}^n : h \geq 0\}$, and if it can be ensured that the controller renders this set forward invariant, then the system can be guaranteed to remain within a safe region. This is the motivation behind Control Barrier Functions (CBFs), which were developed to formalize safety constraints [1]. This study focuses on the implementation of CBFs to avoid slip and thereby guarantee a safe grasp of an object with a flexure-based gripper.

2 Problem statement

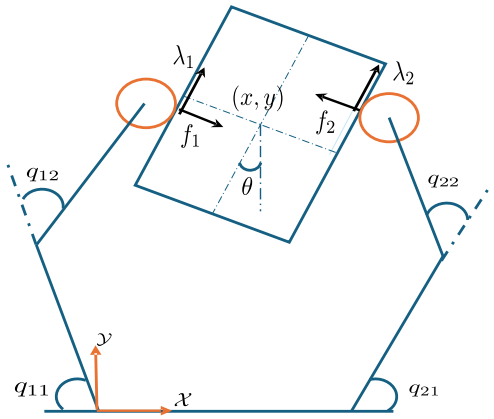

Figure 1: Schematic view of an underactuated gripper pinch-grasping an object with general system coordinates $\zeta = [q_{11}, q_{12}, q_{21}, q_{22}, x, y, z]^T$.

An underactuated gripper is considered to pinch-grasp an object, as depicted in Fig. 1. There is one contact point per finger that executes contact forces $\begin{bmatrix} f_i & \lambda_i \end{bmatrix}^T$ on the object in the normal and lateral direction, respectively. In previous studies, the nominal controller $U_{nominal}(\zeta)$ was designed using the Lyapunov method under the assumption that slip does not occur. However, in practical scenarios, the friction may not always be sufficient to satisfy this condition.

To formulate the safety, the continuously differentiable function h is defined, which is negative when slip occurs and positive otherwise. It is formulated for two contact points as follows:

$$h_i = (f_i \cdot \mu)^2 - \lambda_i^2, \quad i = 1, 2, \tag{1}$$

where μ is the friction coefficient. As long as both h_1 and h_2 remain positive, the controller operates within the safety region. As shown in Figure 2, when the friction coefficient μ decreases from 0.8 to 0.2, slip occurs at at both contact points.

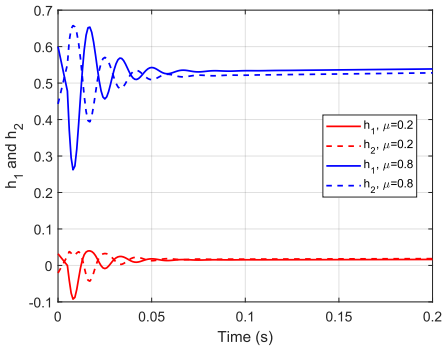

Figure 2: h_1 and h_2 vs Time for Different values of μ

3 Discussion

The goal of this study is to incorporate the CBF in the controller design to ensure forward invariance of the safe set $\mathscr{C}$. The control input can be computed by solving the following Quadratic Program (QP), known as the CBF-QP formulation which minimally modifies the nominal controller through quadratic optimization [1]:

$$u = \arg\min \frac{1}{2}\|u - U_{nominal}\|^2 \quad \text{, subject to} \quad \dot{h} \geq -\alpha(h),$$

where $\alpha(\cdot)$ is a class $\mathscr{K}_\infty$ function.

Similar studies have implemented CBF to constrain the force overshoot in robot–environment interactions [2]. In most existing approaches, $h := h(\zeta)$ depends solely on the system's state, and its time derivative satisfies

$$\dot{h}(\zeta) = L_f h(\zeta) + L_g h(\zeta) u,$$

which is affine in the control input u. As a result, the CBF constraint remains linear in u, rendering the associated quadratic program convex and solvable. In contrast, in the current study $h := h(u, \zeta)$ depends explicitly on the control input u, hence,

$$\dot{h}(\zeta, u) = \frac{\partial h}{\partial \zeta}\dot{\zeta} + \frac{\partial h}{\partial u}\dot{u},$$

which introduces an explicit dependence on $\dot{u}$. Consequently, the CBF constraint is no longer affine in u, and standard CBF-QP formulations are no longer directly applicable without additional assumptions or reformulations. Ongoing work is on a new strategy to ensure safety for the considered case.

References

[1] Ames et al.,Eur.Control Conf. (ECC), pp.3420–3431, 2019.
[2] Liang et al.,Int. Conf.Autom. Comput.(ICAC),pp.1–7,2023.

Reinforcement Learning for Navigation Skill Orchestration

Daan Meulendijks[1], Jordy Senden[1], Mehrdad Mohammadi[2], Ömür Arslan[1], René van de Molengraft[1]
Depts. of Mechanical Engineering (1) and Industrial Engineering and Innovation Sciences (2)
Eindhoven University of Technology
{d.p.meulendijks, j.p.f.senden, m.mohammadi1, o.arslan, m.j.g.v.d.molengraft}@tue.nl

1 Introduction

For autonomous navigation, mobile robots rely on a proper composition of various interdependent navigation skills, such as path planning, path following, replanning, and (re)mapping, often coordinated using fixed rule-based systems (e.g., behavior trees) [1]. While reliable and transparent, such rule-based approaches usually struggle to adapt to changing environmental, computational, or sensing conditions, motivating data-driven, learning-based decision-making for skill orchestration and execution. In this work, we propose a reinforcement-learning (RL)-based skill orchestrator for the automated supervisory composition of navigation skills in the Nav2 navigation stack for the ROS robot software ecosystem [2]. We apply reinforcement learning to achieve a balance between improving skill- and task-level performance while reducing unnecessary computation, and to enhance robustness and generalizability to environmental variability. Overall, our navigation skill orchestrator enables safe and interpretable skill integration based on existing navigation skills, rather than designing a single end-to-end black-box neural navigation policy.

2 Skill Orchestrator Framework

We build our navigation skill orchestrator within a standard RL framework, specifically Deep Q-Network (DQN), using a systematic model-based design of the state representation, action space, and reward mechanism, as illustrated in Figure 1.

As a state representation, our skill orchestrator uses a composition of navigation performance metrics (e.g., path cost and navigation progress), environmental context (e.g., unexpected obstacles), and skill-level status information (e.g., last skill usage time and skill performance). At each decision step, occurring at a fixed rate, the orchestrator selects an optimal navigation skill from a discrete action set of navigation skills based on the learned action-value (Q) function. This allows adaptive handling of unexpected situations, such as dynamic obstacles that invalidate the current plan as shown in Figure 2.

We use an additive reward combining task success and penalties for computation, unsafe behavior, and excessive skill switching:

$$r_t = r_{\text{goal}} + \lambda_1 r_{\text{comp}} + \lambda_2 r_{\text{safe}} + \lambda_3 r_{\text{switch}}, \qquad (1)$$

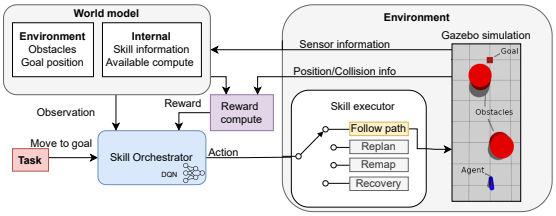

Figure 1: The navigation skill orchestrator uses environmental context, skill status, and task performance to learn to adaptively select an optimal navigation skill relevant to the robot's instantaneous situation, enabling it to handle unexpected scenarios in a safe and interpretable manner.

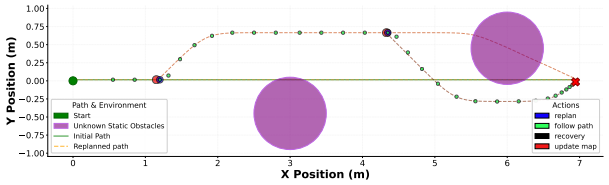

Figure 2: An example execution of the navigation skill orchestrator, which adaptively selects different skills to appropriately handle the robot's instantaneous situation and needs during navigation.

where r_{goal}, r_{comp}, r_{safe} and r_{switch} denote the goal progress, computation, safety, and skill-switching rewards, with their corresponding weights.

In summary, we present an RL-based framework for explainable learning of adaptive skill orchestration for mobile robot navigation in unstructured and dynamic environments. We believe that learning skill composition allows for maintaining a trace of decision making and increases transparency compared to end-to-end black-box learning models. Our ongoing work focuses on expanding skill sets and improving resource-aware generalization across tasks and environments.

References

[1] M. Iovino *et al.*, "A survey of behavior trees in robotics and ai," *Robotics and Autonomous Systems*, vol. 154, p. 104096, 2022.

[2] S. Macenski *et al.*, "The marathon 2: A navigation system," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2718–2725, 2020.

Optimization-Driven Manipulator Design for Confined Workspace

Chaoyue Fei, Hafsa Nouhi, Gaoyuan Liu, Bram Vanderborght, Greet Van de Perre
Brubotics, Vrije Universiteit Brussel, and IMEC, Belgium
Email: chaoyue.fei@vub.be

1 Introduction

Consider a warehouse scenario where the robotic manipulator's design is critical for ensuring efficient and effective operations. The robot operates continuously, maintains the flow of goods, and often functions in a restricted environment. Standard "off-the-shelf" commercial robots are typically designed for general-purpose applications and may not be suitable for specific task and constraint combinations [1]. Mapping the robot design to the intended tasks can improve efficiency, performance, and safety.

2 Optimization-Driven Design

As shown in Fig. 1, obstacles are 0.5m from the robot, and goal poses are sampled in the yellow region. Using a standard UR5 as the initial manipulator, we evaluate its suitability with an RRT planner. We generate 1000 random goal poses and run the RRT (max 1000 iterations) from a random collision-free configuration to each goal. Table 1 reports a success rate of only 6.5%, indicating that the UR5 is unsuitable for this confined workspace.

In the following formulation we adjust the UR5's non-zero Denavit–Hartenberg parameters while fixing the link twists $\alpha^0 = [90^\circ, 90^\circ, -90^\circ]^T$. The objective is to reduce the overall reach error δ and simultaneously improve the robot's manipulability, quantified by the global normalized index G_n [2]. The optimization problem is formulated as:

$$\text{(P1)} \quad \begin{aligned} &\min_{\mathbf{a},\mathbf{d}} f(\mathbf{a},\mathbf{d};\alpha^0) = \delta - G_n \\ &\text{s.t. } a_k^- \le a_k \le a_k^+, \ k \in \{1,2,5\} \\ &\quad d_l^- \le d_l \le d_l^+, \ l \in \{0,3,4\} \end{aligned}$$

where $\mathbf{a} = [a_1, a_2, a_5]^T$ and $\mathbf{d} = [d_0, d_3, d_4]^T$ are the design variables; $\delta = \sum_{\hat{i}} \left| (r_{\max,\hat{i}} - r_{\min,\hat{i}}) - R_{\text{target},\hat{i}} \right|$ measures the deviation between the robot's reachable workspace and a desired target range; G_n captures the average manipulability over the workspace, normalized by the maximum value.

3 Results

The optimization is performed using Differential Evolution (DE). To ensure the robot's reach is neither excessively large nor overly constrained, the target reach range $R_{\text{target},i}$ is set to $[1.2\text{m}, 1.2\text{m}, 1.2\text{m}]$ along the x, y, and z axes, respectively.

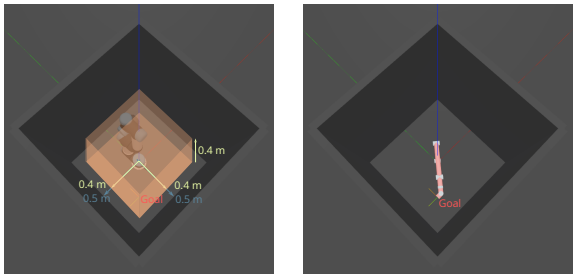

(a) UR5 (b) Optimized Design

Figure 1: The obstacles are positioned 0.5 m from the robot's origin. The goal poses are sampled in yellow area.

Robot	G_n	δ(m)	Succ. Rate	Avg. Iter.
UR5	0.20	2.00	6.5%	828
P1	0.22	0.04	33.2%	782

Table 1: Robot performance: reach error, global normalized manipulability, and RRT performance

As shown in Table 1, the optimized design P1 achieves a reach error within 0.05m while improving the global normalized manipulability from 0.20 to 0.22. The enhanced kinematic performance translates to improved motion planning efficiency: as demonstrated in Table 1, the RRT success rate increases from 6.5% to 33.2% when finding collision-free paths. This optimization-driven design approach yields a robot configuration that not only achieves the desired workspace with high accuracy but also exhibits enhanced dexterity over the workspace.

Acknowlegment: This work was funded by FWO research grant G0A9623N and 12APZ24N.

References

[1] Lefebvre, T., Wauters, J., Ostyn, F., Crevecoeur, G. (2023). Towards Task Tailored Articulated Robot Designs. In 2023 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM). https://doi.org/10.1109/AIM46323.2023.10196183

[2] Patel, S., Sobh, T. (2015). "Manipulator Performance Measures - A Comprehensive Literature Survey." In Journal of Intelligent Robotic Systems (Vol. 77, Issues 3–4, pp. 547–570). Springer Science and Business Media LLC. https://doi.org/10.1007/s10846-014-0024-y

Dynamic Spline Model for Soft Robotic Grippers

Yulyan Wahyu Hadi, Bayu Jayawardhana
Engineering and Technology Institute (ENTEG), University of Groningen, The Netherlands
y.w.hadi@rug.nl, b.jayawardhana@rug.nl

1 Introduction

Soft robotic grippers enable adaptable interaction with objects of varying shapes and stiffness. However, their continuous deformation and nonlinear behavior pose significant challenges for real-time modeling and control. Previous studies have explored the Rigid Link Discretization approach [1] that simulate discrete rigid body. Numerical methods like FEM are computationally expensive for high accuracy, while simplified approaches such as constant-curvature models lack fidelity under external loads. Although continuum mechanics methods like Cosserat rod theory offer high accuracy, their computational cost limits real-time use. To address these challenges, this work proposes a continuous spline-based dynamic model for soft grippers, which captures curvature and deformation using a compact representation. The model balances accuracy and computational efficiency, making it suitable for real-time estimation and control.

2 Dynamic Spline

Splines are represented by parametric curves composed of piecewise polynomial segments [2]. In this study, Bézier curves with Bernstein polynomials are employed as the blending basis. a Bézier curve of degree n can be expressed as

$$\Big(\mathbf{p}(u)\Big)(t) = \sum_{i=0}^{n} B_i^n(u)\,\mathbf{P}_i(t),$$

that describe any position of the beam along the spline curve $\mathbf{p}(u)$, $u \in [0,1]$, where $\mathbf{P}_i(t)$ are the positions of the i-th control point. The Euler-Lagrange modeling framework is used to describe the equations of motion of the soft gripper dynamics

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{\mathbf{q}}}\right) - \frac{\partial T}{\partial \mathbf{q}} + \frac{\partial U}{\partial \mathbf{q}} - \left\{\frac{\partial \psi}{\partial \mathbf{q}}\right\}^{\top} \lambda = \mathbf{Q}, \quad i = 1,\ldots,n$$
$$\{\psi\} = \mathbf{0},$$

where $\mathbf{q}$ and $\dot{\mathbf{q}}$ are the generalized coordinate and its time derivative, T and U are the kinetic and potential energy of the system, ψ and λ are the vector of constraints, and its Lagrange multipliers, and $\mathbf{Q}$ is the generalized force. Since the spline is a continuous curve and assuming uniform mass density, the kinetic energy T can be written as

$$T(\mathbf{q},\dot{\mathbf{q}}) = \frac{1}{2}\int_0^1 \left\{\big(\dot{\mathbf{p}}(u)\big)(t)\right\}^{\top} \mathbf{M} \left\{\big(\dot{\mathbf{p}}(u)\big)(t)\right\} \left\|\frac{d\mathbf{p}(u)}{du}\right\| du,$$

where $\big(\dot{\mathbf{p}}(u)\big)(t) = \frac{d\left(\mathbf{p}(u)\right)(t)}{dt}$ is the time derivative of the spline parametric expression, $\mathbf{M}$ is the uniform mass matrix of the spline, and $\left\|\frac{d\mathbf{p}(u)}{du}\right\| du$ is the arc length. The potential energy from bending and stretching are expressed as

$$U_{\text{bending}}(t,\mathbf{q}) = \frac{1}{2}\int_0^1 EI\left(\varepsilon_b(u) - \varepsilon_b^0(u)\right)^2 \left\|\frac{d\mathbf{p}(u)}{du}\right\| du,$$

$$U_{\text{stretching}}(t,\mathbf{q}) = \frac{1}{2}\int_{u_{\min}}^{u_{\max}} EA\left(\varepsilon_s(u) - \varepsilon_s^0(u)\right)^2 \left\|\frac{d\mathbf{p}(u)}{du}\right\| du,$$

where ε_b^0, ε_s^0, ε_b, and ε_s are the bending and stretching strains of the free-form spline and its deformed pose, E, I, and A are the Young modulus, moment of inertia of the cross section, and the cross-sectional area of the spline. The generalized force and torque component is obtained as

$$Q_{i,F} = \frac{\partial L_F}{\partial \mathbf{P}_i}, \quad Q_{i,T} = \frac{\partial L_T}{\partial \mathbf{P}_i}.$$

3 Results

The beam exhibits increased flexibility and displacement as the number of control points increases, as in Figure 1. It validates the spline-based approach for modeling static deformation of the soft gripper.

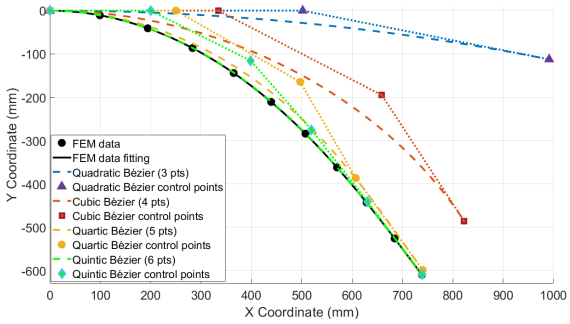

Figure 1: Beam displacement with -0.5N external force applied at the tip in y direction for different spline resolutions

References

[1] Y. W. Hadi, et al, "Developing simulation models for soft robotic grippers in Webots," in *Proc. IEEE 8th Int. Conf. on Soft Robotics (RoboSoft)*, 2025.

[2] P. Valentini, et al, "Modeling elastic beams using dynamic splines," *Multibody System Dynamics*, 2011.

Hardware-in-the-Loop Setup for Testing and Validation of Autonomous Driving Applications

Shuhao Zhang, Jan Swevers, Wilm Decré
MECO Research Team, Department of Mechanical Engineering,
KU Leuven Flanders Make@KU Leuven, Leuven, Belgium
shuhao.zhang@kuleuven.be

1 Introduction

To test and validate autonomous driving applications, a hardware-in-the-loop (HiL) setup offers the advantages of combining rapid prototyping from simulation-based design with a safe and cost-effective alternative to real-vehicle testing. This abstract presents a modular HiL experimental framework with flexible and easily integrated design, as illustrated in Fig. 1, for testing and validating a wide range of autonomous driving applications.

2 HiL framework for autonomous driving

The proposed HiL framework shown in Fig.1 is composed of three main functional blocks:
1. Projection Box: the framework employs a camera-in-the-loop setup by mounting a camera and a PC monitor within a custom aluminum alloy frame box as illustrated in Fig.2. The box is then covered with a masking object, such as a black blanket, to minimize interference from external light. The monitor, connected to the host computer, displays the real-time view from a virtual vision sensor mounted on the simulated vehicle. Instead of directly using the synthetic visual data from the simulator, we employ an Intel RealSense Depth Camera D435 to capture the image displayed on the monitor, thereby introducing a more realistic condition similar to that encountered during the practical deployment.
2. Host Computer: The framework simulates driving scenarios and vehicle dynamics on the host computer. Driving scenarios are designed using Simcenter Prescan software, which enables the creation of digital twins of various environments. The vehicle simulator is implemented as a ROS2 node with customized vehicle models. Through the ROS2-Prescan bridge component, the camera view from Prescan is acquired in real-time and synchronized with the corresponding vehicle pose. The host computer enhances the flexibility of simulating both the environment and the vehicle, with its computational resources independent of the other components of the framework.
3. Embedded Computing Board: The framework leverages an NVIDIA Jetson AGX Orin to implement key components, including a vision-based detector, a motion planner, and a tracking controller, within a ROS 2 network. The tracking controller closes the control loop between the vehicle simulator running on the host computer and the planner on the Jetson by subscribing to both the current vehicle states and the planner-generated reference states, and by publishing control inputs to the simulator at a fixed frequency, thereby enabling real-time trajectory tracking. In contrast, the planner operates in an event-driven manner rather than at a fixed update rate, allowing replanning to be triggered by new perception results when necessary—for example, when a previously selected parking spot becomes occupied in an AVP scenario. Image data are not transmitted over the ROS 2 network; instead, they are directly acquired from the on-board camera via the RealSense API on the Jetson platform.

Acknowledgment This work has received funding from the Flemish Agency for Innovation and Entrepreneurship (VLAIO) under research project No. HBC.2021.0939 (BECAREFUL), from the European Union's Horizon 2020 research, and innovation program under grant agreement No. 953348 (ELO-X) and from the Flanders Make SBO project ARENA: Agile and Reliable Navigation.

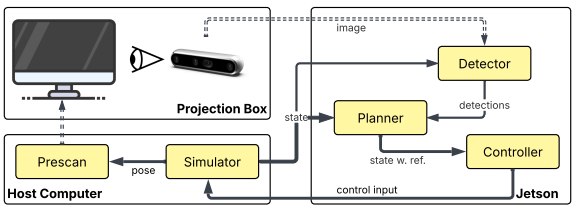

Figure 1: Architecture of the HiL framework.

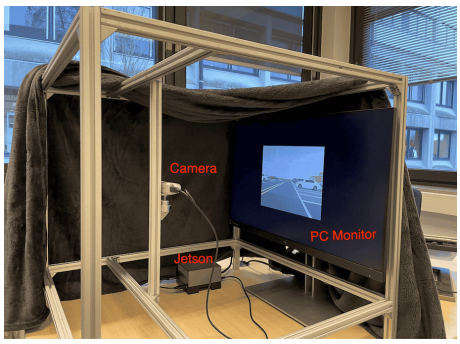

Figure 2: The projection box.

Effects of Model Mismatch on Equivalent CACC Controller Realizations

Mischa Huisman*
M.R.Huisman@tue.nl

Carlos Murguia Rendon*
C.G.Murguia@tue.nl

Erjen Lefeber*
A.A.J.Lefeber@tue.nl

Nathan van de Wouw*
N.v.d.Wouw@tue.nl

1 Introduction

Cooperative Adaptive Cruise Control (CACC) depends on accurate longitudinal vehicle models. Most CACC designs rely on a simplified double-integrator representation obtained via input–output linearization [1],

$$\dot{q}_i = v_i, \qquad \dot{v}_i = a_i, \tag{1a}$$

$$\dot{a}_i = -\tfrac{1}{\tau}a_i + \tfrac{1}{\tau}u_i, \tag{1b}$$

where q_i, v_i, and a_i denote the position, velocity, and acceleration of the ith vehicle, respectively, u_i is the desired acceleration, and τ is a driveline constant. While this model forms the basis of CACC controllers [2, 3], it implicitly assumes perfect knowledge of the underlying vehicle dynamics. In practice, model parameters are uncertain, which may alter closed-loop behavior. This work shows that controller realizations that are equivalent under nominal conditions can exhibit distinct responses in the presence of uncertainty.

2 Problem Formulation

Input–output linearization in [1] accounts for forces such as drag, viscous friction, rolling resistance, and road slope, each depending on uncertain parameters. To capture the resulting model mismatch, we augment the acceleration dynamics with a state-dependent uncertainty term,

$$\dot{a}_i = -\tfrac{1}{\tau}a_i + \tfrac{1}{\tau}u_i + f_\Delta(v_i, a_i, u_i), \tag{2a}$$

$$f_\Delta(v_i, a_i, u_i) = \Delta_1 v_i + \Delta_2 a_i + \Delta_3 u_i + \Delta_4 v_i^2 + \Delta_5 v_i a_i + \Delta_6, \tag{2b}$$

where f_Δ captures the uncertainty induced by imperfect input–output linearization. As shown in the simulation study, this uncertainty causes nominally equivalent CACC controllers to produce different vehicle responses when $\Delta = [\Delta_1\ \Delta_2\ \Delta_3\ \Delta_4\ \Delta_5\ \Delta_6] \neq 0$.

3 Simulation Study

A two-vehicle platoon governed by (1)–(2) is simulated under nominal conditions ($\Delta = 0$) and in the presence of uncertainty $\Delta \neq 0$. Two CACC controller realizations from [2, 3] are considered with $\tau = 0.1\,\text{s}$, $h = 0.5\,\text{s}$, $r = 3\,\text{m}$, and gains $(k_p, k_d) = (0.2, 0.7)$. The leader tracks a constant reference velocity of $20\,\text{m/s}$. Figure 1 shows the inter-vehicle distance and the follower velocity for $\Delta = \begin{bmatrix}0.1 & 0.05 & 0.01 & 0.005 & 0.01 & 0.1\end{bmatrix}$. While both controller realizations yield identical nominal closed-loop behavior, their responses differ under uncertainty. This motivates the use of *Plant Equivalent Controller realizations* [4] as a means to mitigate performance degradation due to model uncertainty and to maintain closed-loop behavior close to the nominal design.

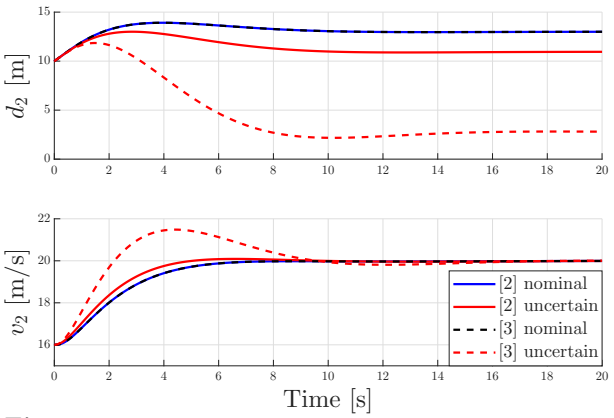

Figure 1: Comparison of follower velocity v_2 and inter-vehicle distance $d_2 = q_1 - q_2$ for two CACC controller realizations [2, 3] under nominal conditions and in the presence of model uncertainty.

References

[1] S. Sheikholeslam and C. A. Desoer, "Longitudinal Control of a Platoon of Vehicles," in *Amer. Control Conf.*, 1990.

[2] J. Ploeg et al., "Design and experimental evaluation of cooperative adaptive cruise control," in *Conf. on Intel. Trans. Sys.*, 2011.

[3] E. Lefeber et al., "Cooperative Adaptive Cruise Control of Heterogeneous Vehicle Platoons," *IFAC-PapersOnLine*, vol. 53, 2020.

[4] M. Huisman et al., "Plant equivalent controller realizations for attack-resilient cyber-physical systems," *arXiv preprint arXiv:2512.13229*, 2025.

*Dynamics & Control, Mechanical Engineering, Eindhoven University of Technology, PO Box 513, 5600 MB Eindhoven, The Netherlands

Control-Oriented Identification Modeling for Ship Maneuvering: A Physics-Informed and Scenario-Based Approach

Yan-Yun Zhang[1,2*], Katrien Eloot[3], Jan Swevers[1,2]
[1]Department of Mechanical Engineering, KU Leuven, Leuven, Belgium
[2]Flanders Make@KU Leuven, Leuven, Belgium
[3]Flanders Hydraulics, Antwerp, Belgium
Email: yanyun.zhang@kuleuven.be, katrien.eloot@mow.vlaanderen.be, jan.swevers@kuleuven.be

1 Introduction

Accurate ship motion prediction is important for both human and machine with the foresight to make informed navigational decisions. These predictions, however, are challenging due to the complex hydrodynamic interactions and is fundamentally dependent on the mathematical models used to represent the vessel's motion response. Traditional methods like the MMG standard method [1] prioritize high-fidelity physical replication across a global range of operating conditions and are frequently computationally intensive and rely on extensive, often unrealistic, excitation datasets for parameter identification. Conversely, *control-oriented* models prioritize short-horizon predictive accuracy, numerical stability, and real-time identifiability. To bridge this gap, we propose a physics-informed and scenario-based identification modeling framework adopting a *complexity on demand* strategy, by tailoring model structures to specific maneuvers. The approach ensures real-time identifiability while maintaining physical consistency and is validated using high-fidelity simulation data from Flanders Hydraulics.

2 Methodology

The framework employs a two-staged approach: (I) baseline model establishment, and (II) online parameter updating.

2.1 Baseline Model Establishment

Baseline model establishment involves three core components: **scenario-based model structuring**, which tailors Newtonian equations for course-keeping and course-changing maneuvers based on dominant physical phenomena; **data preparation**, where identification datasets are optimized for scenario-specific persistent excitation; and **parameter identification**, which evaluates two optimization criteria: (i) one-step ahead prediction error (derivative matching) and (ii) multi-step ahead simulation error (time-series curve fitting).

*Yan-Yun Zhang is a researcher funded by the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 955.768 (MSCA-ETN AUTO-Barge).

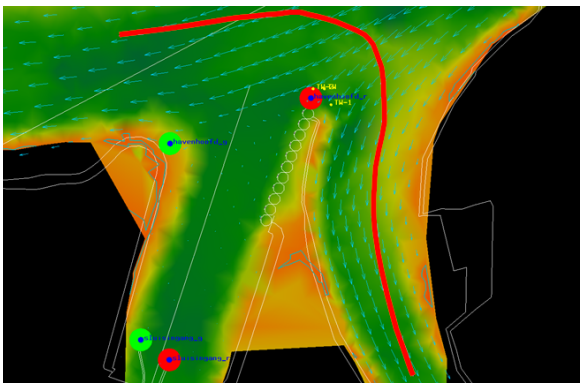

Figure 1: Overview of the Simulated Scenario: Realistic Voyage under Dynamic Environment

2.2 Online Parameter Updating

To address dynamic disturbances, the framework employs Recursive Least Squares (RLS) for online parameter updates, incorporating parameter bounding to ensure stability and prevent non-physical convergence during low-excitation maneuvers.

3 Results and Conclusion

The proposed framework is validated through a simulated voyage under dynamic disturbances, as illustrated in Figure 1. Modeling results demonstrate that multi-step simulation error optimization achieves superior global accuracy and anti-noise robustness. Furthermore, the baseline models are proven effective for ensuring physical consistency during online updates. In conclusion, the proposed modeling framework successfully balances physical constraints with real-time identification efficiency, providing a robust foundation for autonomous control application for marine vehicles.

References

[1] Yasukawa, H., Yoshimura, Y., 2015. Introduction of MMG Standard Method for Ship Maneuvering Predictions. Journal of Marine Science and Technology, (20):37–52.

Compensating the effects of hang for agricultural vehicles with AWS

Patrick Conjarts, Tom van der Sande
Mech. Engineering (D&C), TU Eindhoven
Groene Loper 3, 5612AE Eindhoven - NL
Email: p.m.d.conjarts@student.tue.nl
t.p.j.v.d.sande@tue.nl

Franz Fuchshumer
Research & Advanced Eng., AGCO GmbH (Fendt)
Johann-Georg-Fendt Str. 4, 87616 Marktoberdorf - DE
Email: franz.fuchshumer@agcocorp.com

Even though modern agricultural vehicles can have a working width of over 50 meters, accuracy requirements of these machines are on the centimetre level, while field conditions are often unpredictable. The lateral dynamics of an agricultural vehicle operating in hilly regions may be affected by the slope of the field, influencing the efficiency and accuracy of the operation at hand. Here a control strategy is proposed for agricultural vehicles equipped with all-wheel-steering (AWS), to prevent the machine from drifting in a downhill direction when driving perpendicular to the slope. In order to validate the control strategy in practice, a prototype vehicle was made available by *AGCO GmbH (Fendt)*.

A schematic representation of the vehicle is shown in Figure 1. Contrary to most literature (e.g. [1]), this research considers a variation of implements to be attached to the tractor, which influences the position of both the centre of mass (CM) and the desired control point (CP) along the longitudinal axis of the vehicle, not necessarily constrained to the wheelbase l_{wb}. Certain implements additionally enforce minimizing the body-slip angle of the vehicle, to limit lateral forces on the implement. This creates the requirement that the cross track error, γ, and heading angle error, Θ, can be controlled independently.

By considering a single-track model to represent an AWS-vehicle, it is possible to represent the effects of hang using

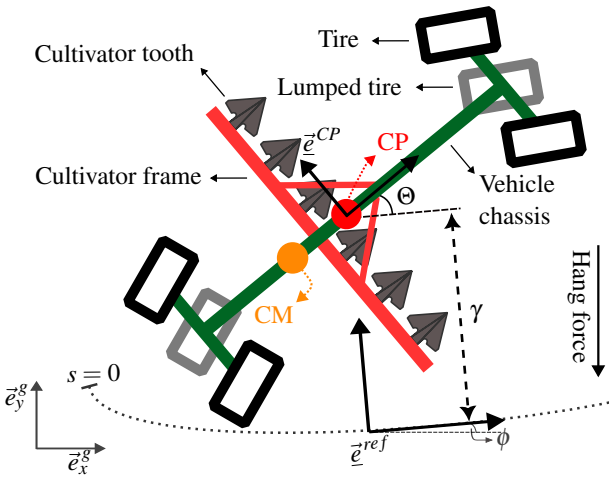

Figure 1: Top view and definitions of the schematic representation of the vehicle w.r.t. a curved reference path

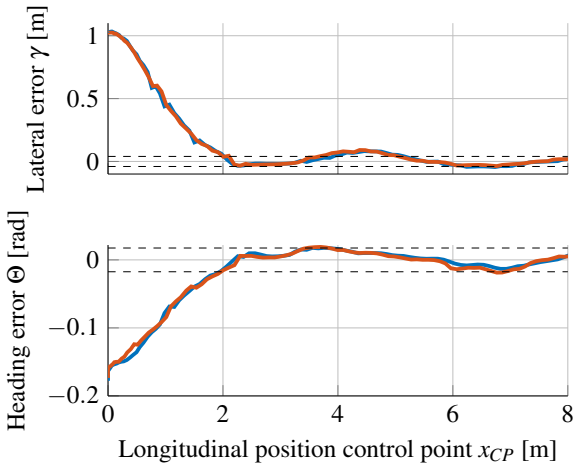

Figure 2: Experimental comparison between kinematic controller (—) and controller with observer-based side-slip angle estimation (—), where $\gamma_{ref} = 0\,\text{m}$, $\Theta_{ref} = 0\,\text{rad}$

a side-slip angle on the front and rear axle [2]. Based on a coordinate transformation and input-output linearization, a position-based control law was found which separates control of the position and orientation of the control point. By opting for a position-based control law, the body-slip angle can be minimized even at low operating velocities. Since knowledge on tire characteristics is not available, an observer can be used to estimate the side-slip angles [2].

Experiments have been conducted under a 10 degree hang, of which the results are shown in Figure 2. In the vehicle configuration used for this experiment, CM is positioned at approximately $l_{wb}/2$ from the front axle, while CP is positioned at approximately $l_{wb}/4$ from the front axle. Given an initial lateral error of 1 m and a constant velocity of 1 m/s, the vehicle was controlled to a straight reference at a desired accuracy of 4 cm within a longitudinal distance of 5.5 m.

References

[1] J. Zhu et al., "Design and Simulation of Steering Control Strategy for Four-Wheel Steering Hillside Tractor," Agriculture, 2024

[2] P. Conjarts et al., "Compensating the effects of hang for agricultural vehicles equipped with all-wheel steering," 2025.

Ethical and safety-critical decision making in autonomous driving with V2V and V2X communication

Taha Bouzid, Mircea Lazar, and Sofie Haesaert

1 Introduction

Mature Vehicle-to-Vehicle (V2V) and Vehicle-to-Everything (V2X) communication between vehicles and the infrastructure presents novel opportunities and challenges for ethical and safety-critical decision-making in autonomous driving. Information received from these heterogeneous V2V and V2X sources should generally improve safety. However, when these sources are conflicting and their reliability is uncertain, their usage can potentially lead to unsafe or unethical decisions. Such situations are particularly critical in dynamic traffic environments, where autonomous vehicles must react in real time under strict safety constraints. To address these challenges, it is essential to establish formal safety guarantees to achieve safety-critical and ethical decision-making based on these heterogeneous information sources while ensuring smooth, human-like control.

2 Autonomous driving with V2V and V2X

Temporal logic is widely used to establish formal safety guarantees for deterministic planning and control problems in autonomous driving [1], as it can provide an expressive framework to encode both spatial and temporal constraints, making itself well-suited to encode traffic rules and norms [2]. In parallel, temporal-logic-based frameworks have been extended to uncertain models for autonomous driving, providing probabilistic safety guarantees under uncertain intentions [3]. However, dealing with the uncertainty that emerges from potentially unreliable V2V/V2X information remains an open challenge.

Case Study. We consider a scenario with a pedestrian, illustrated in Figure 1, in which an ego vehicle receives perception information via V2V and V2X from surrounding vehicles and infrastructure. Predictions of the pedestrian can depend on the incorporation of the external information sources and exhibit varying levels of reliability. The ego vehicle must make real-time decisions based on these uncertain inputs while adhering to formal safety constraints and traffic rules, and while ensuring smooth, human-like control. However, the uncertainty about the correctness of information received from external sources can create ethical dilemmas.

The research has received funding from the European Union under the Horizon Europe Grant Agreement AIGGREGATE, no. 101202457. All authors are with Control Systems Group, Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands. Email: {t.bouzid, m.lazar, s.haesaert}@tue.nl

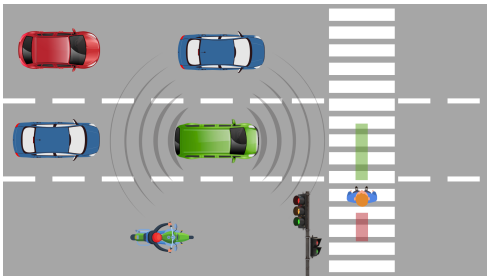

Figure 1: Urban scenario with green ego vehicle and a pedestrian with different predictions (green and red shaded areas).

Problem Setting. We consider an ego autonomous vehicle operating in a dynamic traffic environment with other objects, including vehicles, pedestrians, and cyclists. The ego vehicle receives information about each surrounding object j from a heterogeneous set of V2V/V2X sources Φ_j (e.g., other vehicles and infrastructure units). The reliability of these sources of information is uncertain, and as such leads to uncertainty in the object states and their predictions. More precisely, this uncertainty implies that predictions are not just defined as probability distributions as in [3], but as sets of probability distributions $\mathbf{P}_j$. As such, the true but unknown probability distribution associated with the object j satisfies $p_j \in \mathbf{P}_j$. This leads to the following problem statements.

Problem 1: Leverage formal specification languages to define an ethical framework for safety-critical decision making based on uncertain information from potentially unreliable sources Φ_j.

Problem 2: Develop decision making algorithms that maximize passenger comfort while guaranteeing ethical and safe decisions.

References

[1] P. Halder, H. Homburger, L. Kiltz, J. Reuter, and M. Althoff, "Trajectory planning with signal temporal logic costs using deterministic path integral optimization," 2025. [Online]. Available: https://arxiv.org/abs/2503.01476

[2] M. Althoff, S. Maierhofer, G. Wursching, Y. Lin, F. Lercher, and R. Stolz, "No more traffic tickets: A tutorial to ensure traffic-rule compliance of automated vehicles," *Proceedings of the IEEE*, 2025.

[3] Z. Zhang, Z. Sun, and S. Haesaert, "Intention-aware control based on belief-space specifications and stochastic expansion," 2024. [Online]. Available: https://arxiv.org/abs/2404.09037

Criticality measure for scenario-based testing of automated vehicles

Kevin Tom Kurian[1], Erjen Lefeber[1], Jeroen Ploeg[2], Igo Besselink[1], and Mohsen Alirezaei[1,2]

[1] Eindhoven University of Technology, Eindhoven, The Netherlands

[2] Siemens Digital Industries Software, Helmond, The Netherlands

Email: {k.tom.kurian, a.a.j.lefeber, i.j.m.besselink, m.alirezaei}@tue.nl

1 Background

Scenario-Based Testing (SBT) is a widely used method for safety assurance of Automated Driving System equipped Vehicles (ADS-Vs) [1], a topic of significant research in the intelligent vehicles community. To comply with ISO 21448, a key standard for ADS-V safety assurance, SBT can be formulated as an optimization problem to identify unknown critical scenarios within the ADS-V's Operational Design Domain (ODD). The ODD may contain infinitely many scenarios, and only the highly critical scenarios in the ODD are informative and relevant for safety testing. Thus, a cost function rating scenario criticality is a prerequisite for selecting a representative, informative set of scenarios from the ODD using optimization.

Although many measures for scene/scenario criticality exist [2], they do not satisfy all the requirements of a cost function in SBT optimization, such as indicating proximity to collisions and validity for multiple scenario types like intersections, highways, roundabouts, etc. However, Kurian et al. [3] showcase scenario criticality measure for 1-D motion of the ego vehicle in a highway, and an intersection scenario, that evaluates the proximity to a collision. This work extends the measure to scenarios with 2-D planar motion for the ego vehicle.

2 Formulation (Intersection scenario)

An intersection scenario from the perspective of the ego vehicle, during an unprotected left turn, is shown in Figure 1. The top view of the complete scenario, including the motion path of the actors is shown in Figure 2 (a).

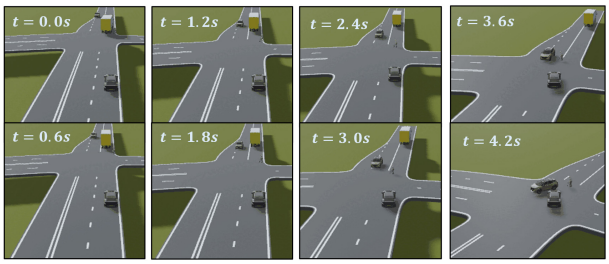

Figure 1: Intersection scenario: scenes over time.

The Criticality Measure (CM) is computed retrospectively for a scenario using the reachable set of the Host Vehicle (HV) of the automated vehicle. The computation of the CM at $t = 1s$, for a horizon of $2s$ is shown in Figure 2 (b). At each time instant, over an analysis horizon, the fraction (ratio) of the collision trajectories in the reachable set of the HV is used to compute the expected severity. The alternate trajectories for the HV are cubic polynomial splines connecting the ego vehicle's pose at $t = 1s$ with final poses sampled from heuristic probability distributions. The other actors are assumed have a pre-specified fixed behavior for the horizon, i.e., non-interactive actors, as per Figure 2 (b). Each trajectory is checked for a collision and aggregated to compute the criticality measure at each time instant as in Figure 2 (c).

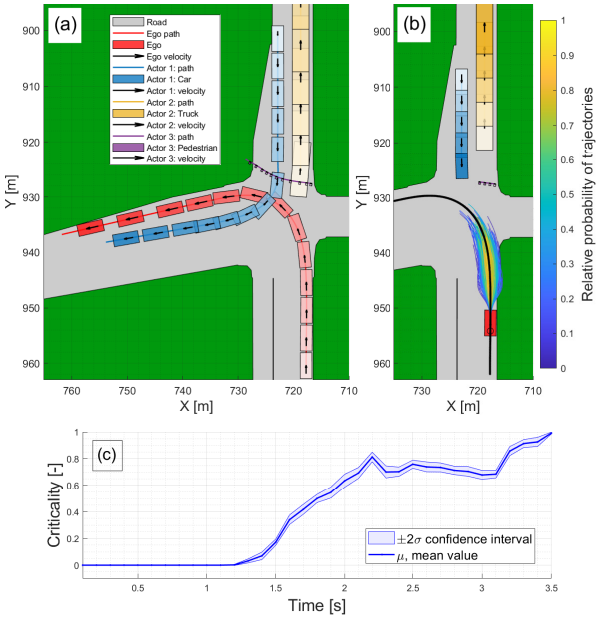

Figure 2: (a)-Scenario, (b)-alternate trajectories at $t = 1s$, and, (c)-Criticality measure vs. time.

3 Conclusion

A Criticality Measure (CM) for scenarios has been developed which indicates the proximity to a collision and is valid for different road networks and scenario types within the ODD. Applying the CM to evaluate recorded scenarios from a database showcases the potential of the CM.

References

[1] S. Riedmaier, "Survey on Scenario-based safety assessment of Automated Vehicles", IEEE Access 2020.

[2] L. Westhofen, "Criticality Metrics for Automated Driving: A review and suitability analysis of the state of the art", Archives of Computational Methods in Engineering 2023.

[3] K. Kurian, "How unsafe was the scenario? A criticality measure for scenario-based testing of automated vehicles", IEEE IV 2025.

Monotone Concave Systems and Riccati Equations

Dongjun Wu
Department of Automatic Control, Lund University
dongjun.wu@control.lth.se

1 Introduction

Positive systems theory has recently gained attention for addressing scalability issue in large-scale control [1] owing to its structural sparsity and tractability of LP formulations. However, positive systems, especially the linear theory of them, do not cover another wide range of systems, which are nonlinear but also exhibit positive behavior, namely, *monotone concave systems* (MCS).

MCS can be viewed as a nonlinear extension of positive system theory. Rockafellar's work provides an early systematic study of this class in an economics context [2], but control problems exhibit distinct features and require different tools.

Here the central object is the Bellman operator

$$T(J)(x) := \inf_{u\in U(x)} \{L(x,u)+J(f(x,u))\} \tag{1}$$

associated with the optimal control problem

$$\min_{u(\cdot)} \sum_{k=1}^{\infty} L(x_k,u_k) \tag{2}$$

subject to

$$x_{k+1} = f(x_k,u_k(x)), \quad x\in X,\ u_k(x)\in U(x) \tag{3}$$

For $L \geq 0$, T is naturally monotone on the cone of positive functions $\mathscr{F}_+(X)$. Moreover, as an infimum of linear functionals, it is concave, a property appears to be underexplored in the control literature. The objective is to show how the Bellman operator induces a broad and structured class of MCS and how this leads to a Riccati-type theory on cones.

2 Results

Let $\mathscr{F}(X)$ be the space of functions on X, and $\mathscr{V} \subset \mathscr{F}(X)$ a finite-dimensional subspace. Then $\mathscr{V} \cap \mathscr{F}_+(X)$ defines a cone. Two important examples are

$$\mathscr{V}_1 = \{f\in\mathscr{V} : f \text{ linear}\},\ \mathscr{V}_2 = \{f\in\mathscr{V} : f \text{ quadratic}\}. \tag{4}$$

In these cases,

$$\mathscr{K}_1 = \mathscr{V}_1 \cap \mathscr{F}_+(\mathbb{R}^n_+) \simeq \mathbb{R}^n_+ = \{x\in\mathbb{R}^n : x\geq 0\},$$
$$\mathscr{K}_2 = \mathscr{V}_2 \cap \mathscr{F}_+(\mathbb{R}^n) \simeq \mathbb{S}^n_+ = \{P\in\mathbb{R}^{n\times n} : P\succcurlyeq 0\}.$$

If $\mathscr{V}$ is invariant under T, then (1) induces a finite-dimensional monotone concave mapping on the cone $\mathscr{K} = \mathscr{V}\cap\mathscr{F}_+(X)$. This restriction leads to an equation of the form:

$$T(J) = J, \quad J\in\mathscr{K}. \tag{5}$$

which we call as algebraic Riccati equation (ARE). The crucial difference between an ARE (5) and a Bellman equation is that the former is constrained to a finite-dimensional cone and is more tractable, while still encoding the underlying optimal control problem.

Existence and uniquess of solutions to AREs are classically tied to stabilizability and observability. We show that these can be precisely stated in the language of cones.

Given vectors $v_1,\cdots,v_m$ in $\mathscr{K}$, define the positive span as: $\text{span}_+\{v_1,\cdots,v_m\} = \{\sum_{k=1}^m \lambda_k v_k : \lambda_k > 0\}$.

Definition 1 (Observability)**.** *We say that the system* (3) *is observable from L if there exist $\ell\in\mathscr{K}$, with $\ell(\cdot)\leq L(\cdot,u), \forall u$, and $m\geq 0$, s.t., for any feedback policy $u(\ell(\cdot))$ with $\ell(f_u^i)\in\mathscr{K}$, there holds*

$$\text{span}_+\{\ell,\ell(f_u),\cdots,\ell(f_u^m)\} \subset \text{Int}\,\mathscr{K}.$$

Definition 2 (Stabilizability)**.** *The system (3) is stabilizable if there exists a policy $u(\cdot)$, a constant $\theta\in(0,1)$ and some $J\in\text{Int}\,\mathscr{K}$ such that*

$$J\circ f_u \leq \theta J. \tag{6}$$

Beyond monotonicity and concavity, a further key property is positive homogeneity: $k\geq 0, x\in\mathscr{K}\Rightarrow T(kx)=kx$, which allows us to use nonlinear Perron-Frobenius theory. The Bellman operator is not positively homogeneous, but it can be homogenized via

$$T_\alpha(J)(x) = \inf_{u\in U(x)} \{\alpha\langle q,J\rangle L(x,u)+J(f(x,u))\}$$

for a $q\in\mathscr{K}^*$. Under this modification, stabilizability implies the existence of a 1-eigenvector of T_α (yields a solution to the ARE) for α small, while observability implies strict monotonicity (yields uniqueness) of T_α and large spectral radius of when α large. These lead to our main result.

Theorem 1. *Suppose that the system (3) is observable and stabilizable on $\mathscr{K}$. Then the ARE associated with $\mathscr{K}$ has a unique solution in* $\text{Int}\,\mathscr{K}$.

The general framework recovers standard AREs as special cases. e.g., the AREs for LQR and for positive systems [3]:

$$P = A^\top PA - A^\top PB(R+B^\top PB)^{-1}B^\top PA + C^\top C, \tag{7}$$

$$p = s + A^\top p - E^\top|r+B^\top p| \tag{8}$$

arise by restricting the Bellman equation to $\mathscr{V}_2$ and $\mathscr{V}_1$, respectively, as in (4). The usual stabilizability of (A,B) and observability of (A,C) then imply the cone-based stabilizability and observability conditions in Definitions 1-2 for the homogenized Bellman operator on $\mathbb{S}^n_+$, while the corresponding conditions for positive systems can be related to [4, (13)] and [4, Corollary 4.1]. This yields a unified Riccati theory for linear, positive, and more general monotone concave systems.

References

[1] A. Rantzer and M. E. Valcher, "Scalable control of positive systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, pp. 319–341, 2021.

[2] R. T. Rockafellar, *Monotone processes of convex and concave type*. American Mathematical Soc., 1967, no. 77.

[3] A. Rantzer, "Explicit solution to Bellman equation for positive systems with linear cost," in *2022 IEEE CDCc*. IEEE, 2022, pp. 6154–6155.

[4] Y. Li and A. Rantzer, "Exact dynamic programming for positive systems with linear optimal cost," *IEEE Transactions on Automatic Control*, 2024.

Where to Place a Heavy Payload on a Multirotor UAV for Best Control Performance

Sander Doodeman[1], Paula Chanfreut[1], Elena Torta[2], Duarte Antunes[1]
[1] Control Systems Technology group, Department of Mechanical Engineering, Eindhoven University of Technology
[2] Robotics group, Department of Mechanical Engineering, Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
Email: s.doodeman@tue.nl

1 Introduction

With advances in aerial manipulation research, more and more tasks can be performed using drones. For different tasks, different payloads need to be attached to a drone. Here, accurate positioning control of, for example, a tool on the drone is essential to complete the tasks at hand, such as camera-based or X-ray inspections. Therefore, we focus on how the position of such a payload relative to the vehicle's Center of Gravity (*CoG*) affects the stability and control performance at an arbitrary point of interest on the UAV, such as the payload position. Moreover, we address how this position can be optimized. Previous research focused only on the stability of a massless payload, without investigating the optimal position of components [1].

2 Zero dynamics analysis

We analyze the stability of the zero-dynamics of a complete nonlinear model of the UAV with payload. We demonstrate that the stability of the zero dynamics depends on the vertical signed distance in the body-fixed frame between the controlled output position and the combined *CoG* of the UAV with payload. Specifically, positioning the output below the *CoG* yields unstable zero dynamics. We also prove that for a planar representation of the system, the system is Lyapunov stable when the payload is above the *CoG*, while below results in an unstable system. Figure 1 illustrates this via numerical simulations for a three-dimensional case, where the zero dynamics are unstable when the output is below the *CoG* in Fig. 1a, and the zero dynamics are marginally stable (oscillating) when the output is above the *CoG* in Fig. 1b.

3 Linear system performance analysis

We further analyze the performance of the linearized UAV model with payload by providing an analytical expression for the $\mathscr{H}_2$-norm as a function of the payload position (*PL*) and the arbitrary controlled output position (*POI*) on the UAV, from which we can quantify the system's attenuation to white noise input disturbances. This allows for deriving the optimal position of the controlled output and the position of the payload with respect to the *CoG* for closed-loop white-noise disturbance rejection capabilities, as a function

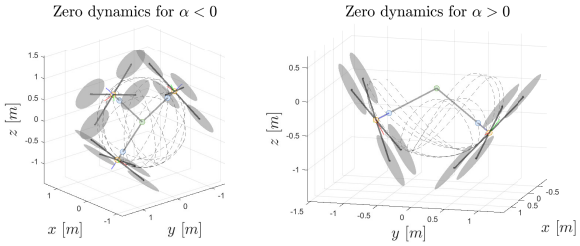

(a) Below-configuration (b) Above-configuration

Figure 1: Zero dynamics trajectories, where the green dot represents the stabilized output position

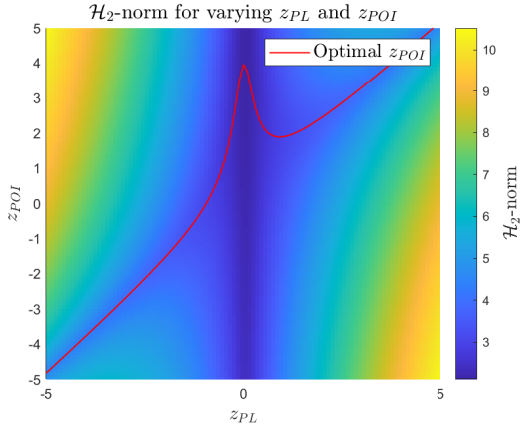

Figure 2: $\mathscr{H}_2$-norm for varying *POI* and *PL*

of a cost term capturing control authority. These optimal positions can be seen in Fig. 2.

We conclude that less control authority leads to a higher optimal position of the controlled output, also when the heavy payload is the controlled output. The results are illustrated through numerical examples.

References

[1] H.-N. Nguyen, C. Ha, and D. Lee, "Mechanics, control and internal dynamics of quadrotor tool operation," *Automatica*, vol. 61, pp. 289–301, Nov. 2015.

ASSISStaNT: Approximating Optimal Controllers for Industrial Applications

Andrea Giusti, Lucian Cristian Iacob, Edward Kikken, Bruno Depraetere
andrea.giusti, cristian.iacob, edward.kikken, bruno.depraetere @flandersmake.be
Flanders Make, Belgium

1 Motivation

Industrial mechatronic systems increasingly demand control strategies beyond simple tracking. Operating under strict constraints, coordinating multiple interacting degrees of freedom, and exploiting future predictions are becoming essential for achieving high performance and efficiency. Model Predictive Control (MPC) offers a comprehensive framework to meet these advanced requirements, yet its industrial uptake remains limited due to practical deployment challenges and elevated computational demands. Moreover, MPC-generated control actions can appear opaque to operators, reducing trust in the system, complicating tuning and day-to-day support. To address these limitations and bring advanced control strategies closer to industry, we developed the ASSISStaNT toolchain to approximate the MPC strategy with a simpler architecture.

2 Approach

The proposed approach starts with the definition of a model, constraints and the control objective. An MPC strategy is designed to solve such a control problem, and multiple simulations are executed to generate a reference dataset. The user then defines the structure of the approximating controller (see Figure 1), by selecting a range of *prediction* horizons, a non-minimal set of *features* (linear and nonlinear functions of state, exogenous signals, ...), and *control blocks* (simple control actions e.g. proportional, linear filters, ...). Finally, sparse linear regression is used to select and tune the resulting control actions to approximate the MPC trajectories.
The ASSISStaNT software streamlines this approach, providing generic model formulation (nonlinear ODE and DAE models), interfaces to automatically extract model description from Simulink/SimScape schematics, libraries of features and control blocks. Finally, the resulting approximate controller can be validated in simulation, and through Simulink it can be compiled and directly deployed on real-time machines for industrial applications.

3 Results

This approach has been validated on various industrial use cases, including autonomous path following for articulated truck-trailer vehicles, multi-input coordination for uniform heating of a slab moving over a conveyor belt, and fast gear synchronization for automatic shifting of heavy vehicles. This last example consists of controlling the speed of an electric motor, and quickly synchronize it to a reference speed, to allow the smooth closure of a dog-clutch. The problem is complicated by the presence of constraints on both the torque and the temperature allowed for the motor. Figure 2 shows one of the example MPC trajectories, and the validation of the resulting controller for a new, unseen, trajectory. Note that the ASSISStaNT controller perfectly reproduces the optimal behavior, while reducing by 100 times the computation time.
Overall, the ASSISStaNT framework brings approximated optimal control to industry in an easily deployable, interpretable and tunable [1] manner.

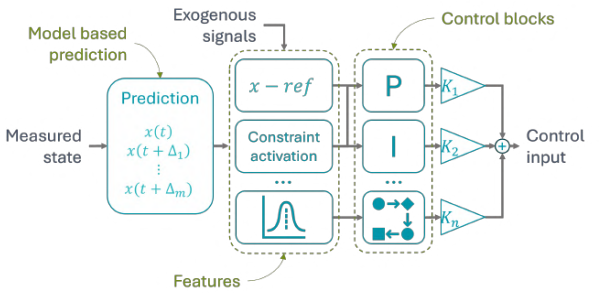

Figure 1: Approximate optimal controller architecture.

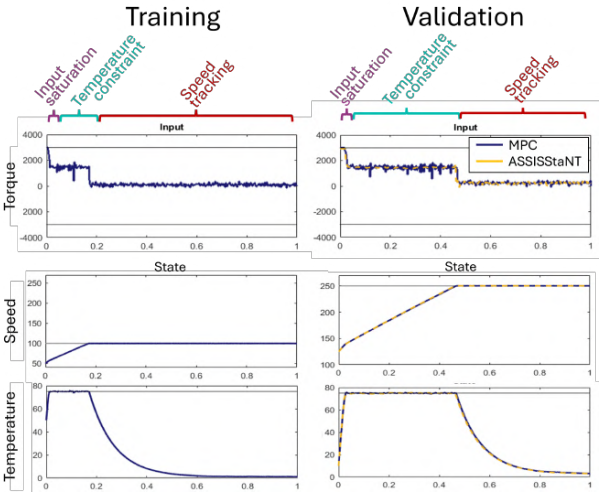

Figure 2: Example trajectories for training and validation of the ASSISStaNT controller for an automatic gear-shift.

References

[1] L. C. Iacob, A. Giusti, E. Kikken, B. Depraetere, "ASSISStaNT: Retuning Approximate Optimal Controllers for Industrial Applications", 45th Benelux Meeting, 2026.

Acknowledgment This research was supported by Flanders Make, the strategic research center for the manufacturing industry, in the framework of the Assisstant_IRVA project.

ASSISStaNT: Retuning Approximate Optimal Controllers for Industrial Applications

Lucian Cristian Iacob, Andrea Giusti, Edward Kikken, Bruno Depraetere
cristian.iacob, andrea.giusti, edward.kikken, bruno.depraetere @flandersmake.be
Flanders Make, Belgium

1 ASSISStaNT Toolchain - Finetuning

Industry faces ever-increasing demands for higher performance, energy efficiency, throughput, and robustness. This leads to more complex machines, and are subject to stricter physical and safety constraints. As a result, well-designed control algorithms are now essential to fully exploit machine capabilities. In many fields, *nonlinear model predictive control* (NMPC) is the state of the art. However, for many industrial applications, the state of practice relies on well-tested and understood linear control solutions, even at the expense of performance. With the ASSISStaNT approach [1], we approximate optimal NMPC control actions via simple and easy to deploy control blocks. In the present work, we explore how we can gain insight into retuning and robustness via leveraging well-known linear control theory tools.

Achieving autonomous operation for these complex nonlinear systems is generally beyond the capabilities of simple PID control. However, it is often a struggle to use and deploy advanced control algorithms due to system complexity, tight performance requirements, and the difficulty of tuning in the field. To address these challenges, we are developing the ASSISStaNT toolchain. We start from an optimal NMPC controller: given the performance requirements, this is the best control action that satisfies the physical, safety, and input constraints. However, this controller is hard to tune and computationally heavy. To solve this, the ASSISStaNT toolchain computes a controller that estimates the optimal solution by leveraging user knowledge of the system (key performance indicators, e.g. lateral error to path) and simple control blocks (PID, feedforward, filters, LUT, switches, etc.). These blocks are easily interpretable and tunable in the field to meet performance requirements in different scenarios.

While the ASSISStaNT toolchain automatically selects the optimal combination between features, control blocks, and the preview time (for the predictive controllers), the optimization is inherently open-loop, as the fitting is performed based on the NMPC control signal. It can be that the real world applications, where disturbances or model mismatch are significant, will benefit from fine-tuning the controller obtained through ASSISStaNT. To simplify the analysis, we consider the following nonlinear system:

$$x_{k+1} = f(x_k, u_k),\ y_k = h(x_k) \tag{1}$$

and the controller is a predictive output feedback controller:

$$u_k = K_p \hat{e}_{k+N} = K_p(r_{k+N} - \hat{y}_{k+N}), \tag{2}$$

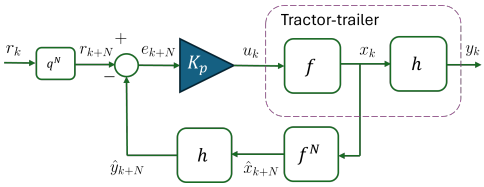

Figure 1: Nonlinear predictor closed loop

where K_p is the static gain and $\hat{e}_{k+N}$ is the predicted error between the reference r_{k+N} and predicted output $\hat{y}_{k+N}$. Here, $N = T_p/T_s$ represents the number of future time steps, T_p is the preview time, and T_s is the sampling time. For the reference, $r_{k+N} = q^N r_k$ is simply a forward shift with N time steps. The predicted output is:

$$\hat{y}_{k+N} = h(x_{k+N}) = h(f^N(x_k, \bar{u})), \tag{3}$$

where f^N is the composition of f with itself N times, and the prediction is performed with a constant input $\bar{u}$. This interconnection is depicted in Fig. 1. As can be observed, this nonlinear closed-loop predictor is complex and it is difficult to analyse stability of the closed-loop system or tune performance. Thus, we perform a linearization on the system dynamics (1), leading to a linear predictive loop. This representation allows to practically assess stability and shape performance through well-known linear control tools, such as eigenvalue analysis, robustness checks via phase and gain margins and sensitivity-based tuning. Preliminary work shows a good correlation between the insights gained from the linearized predictive loop and the simulations performed using the nonlinear tractor-tailer model, as seen in Fig 2.

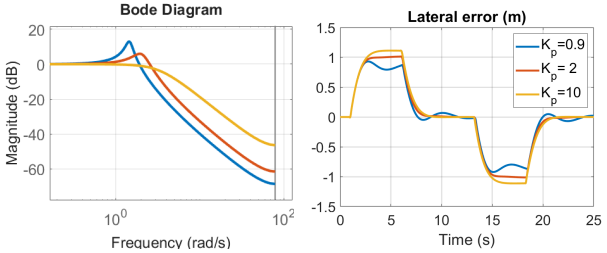

Figure 2: Magnitude of the closed loop frequency response (left) and lateral error in simulation (right).

References

[1] A. Giusti, L. C. Iacob, E. Kikken, B. Depraetere, "ASSISStaNT: Approximating Optimal Controllers for Industrial Applications", 45th Benelux Meeting, 2026.

This work was carried out in the framework of Flanders Make's ASSISStaNT IRVA project. Flanders Make is the Flemish strategic research centre for the manufacturing industry.

Artificial reference Time-Optimal MPC

Alvaro Florez[1*], Jan Swevers[1], Wilm Decré[1]
[1]MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
[1]Flanders Make@KU Leuven, Belgium
[*]Email: alvaro.florez@kuleuven.be

1 Introduction

Model Predictive Control (MPC) is formulated as an optimization problem that is solved online to achieve a given control objective. This objective may vary depending on the application, ranging from point-to-point motion and trajectory tracking to economic optimization under different criteria. One objective that has gained significant interest is the minimization of execution time, aiming to complete tasks as rapidly as possible while improving time efficiency. In this context, Time-Optimal Model Predictive Control (TOMPC) seeks to drive the system to its target in minimum time while satisfying state and actuator constraints. A major challenge of TOMPC lies in solving the resulting MPC problem within a single sampling time, since time-optimal formulations typically require long prediction horizons. This work proposes an artificial reference approach that employs an exponentially weighted ℓ_1-norm with respect to a reference that progressively adapts toward the target. The method constitutes a simplification of time-optimal control strategies, potentially yielding suboptimal trajectories while significantly reducing the computational burden. Nevertheless, the approach provides a practical and computationally efficient solution suitable for a wide range of applications.

2 Artificial reference pseudo-Time-optimal MPC

This technique employs an exponential weighting strategy within a fixed prediction horizon to induce time-optimal behavior. The artificial reference follows a structured strategy in which the system is constrained to reach an intermediate, artificial reference at the end of the prediction horizon. After each optimal control problem (OCP) is solved, this artificial reference is updated and progressively moved toward the true target. This approach enables the use of shorter prediction horizons, thereby significantly reducing the computational complexity of the resulting MPC problem.

$$\begin{aligned} \min_{s,u,s_s} \quad & w_1 \sum_{k=0}^{N-1} \theta^k \|s-s_s\|_1 + w_2 \, \|s_s - s_{\text{end}}\|_2^2 \\ \text{s.t.} \quad & s_0 = s_{\text{ini}}, \\ & s_{k+1} = f(s_k, u_k, T_{s1}) \quad \text{for } k = 0,\dots,N-1, \\ & h(s,u) \le 0 \qquad\qquad\quad \text{for } k = 0,\dots,N-1, \\ & s_N = s_s \end{aligned} \tag{1}$$

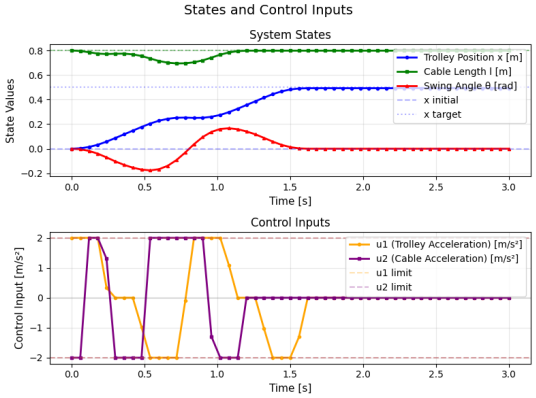

3 Application Example

The proposed methodology is illustrated through an example implementation of the artificial reference time-optimal MPC. The considered system is an overhead crane, whose state variables include the trolley horizontal position and velocity (x, $\dot{x}$), the hoisting cable length and its rate of change (l, $\dot{l}$), and the cable angle and angular velocity (θ, $\dot{\theta}$). The control inputs are the trolley and hoist accelerations ($\ddot{x}$, $\ddot{l}$). The system dynamics are defined as

$$s(t) = \begin{bmatrix} x(t) & \dot{x}(t) & l(t) & \dot{l}(t) & \theta(t) & \dot{\theta}(t) \end{bmatrix}^\top,$$

$$u(t) = \begin{bmatrix} u_1(t) & u_2(t) \end{bmatrix}^\top,$$

$$\dot{s}(t) = \begin{bmatrix} \dot{x}(t) & u_1(t) & \dot{l}(t) & u_2(t) & \dot{\theta}(t) & f_\theta(t) \end{bmatrix}^\top,$$

$$f_\theta(t) = -\frac{g}{l(t)} \sin\theta(t) - \frac{2\dot{l}(t)}{l(t)} \dot{\theta}(t) - \frac{u_1(t)}{l(t)} \cos\theta(t).$$

4 Conclusions

A time-optimal MPC methodology based on an artificial reference was presented in this work. The proposed approach is well suited for practical deployment, as it can be solved more efficiently than traditional time-optimal formulations while retaining a structure compatible with real-time implementation.

Acknowledgement This work was supported by Flanders Make through projects DIRAC_SBO, ASSISSTANT_IRVA, and LearnOpTra_SBO. Flanders Make is the Flemish strategic research centre for the manufacturing industry.

GUESS: GPU Enhanced Synthesis of Supervisory Controllers

Laurens van der Wal[1,*], Anton Wijs[2], Michel Reniers[1]
Control Systems Technology[1], Software Engineering and Technology [2]
Eindhoven University of Technology, The Netherlands
*Email: l.e.v.d.wal@tue.nl

1 Supervisory Control Synthesis

To ensure that a plant satisfies a set of requirements, such as being able to reach a marked state (being non-blocking), we can use Supervisory Control Theory [1] to synthesize a controller for the plant. This controller will then guarantee that all requirements are fulfilled, in addition to preventing the plant from blocking.

This synthesis process consists of a series of computational steps using a model composed of the plant and the requirements. First we perform reachability analysis to identify states that are not able to reach a marked state. Secondly we perform a reachability analysis to find which states are able to reach a blocking state using an uncontrolled event. Lastly we remove all states identified in the previous steps from the plant model. These three steps are repeated until no such states are identified. The goal of this research is to implement these steps into a tool while using the GPU.

2 GPUexplore

GPUexplore is a tool that performs the state-space exploration of a system using a GPU [2]. To achieve this, it uses the internal structure of the GPU, illustrated in Figure 1, by performing the generation of successor states in a parallelized manner on the threads of the Streaming Multiprocessors (SM). The results are stored in the shared memory of the SM. Once the SM are finished, the tool will store the explored states and constructed transition matrix from shared memory in the global memory.

It is important that interactions with global memory are kept to a minimum, as these operations are more expensive compared to interactions with the SM shared memory. To prevent these interactions from being required often, the format in which states and other data from the models are stored is highly optimised through the use of a hash table with compression techniques and the storage of states as binary trees.

The biggest advantage of GPUexplore over contemporary tools is performance, as it outperforms these tools w.r.t. execution time as the size of the state-space increases [3]. The downside of using the GPU is that it is only efficient if threaded homogeneously, where all threads execute the same instruction simultaneously, which can provide a challenge when implementing new methods.

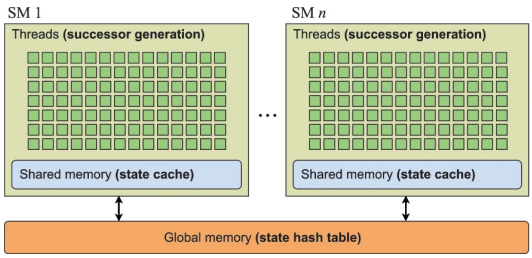

Figure 1: Internal structure of a GPU [3]

3 Implementation

Since contemporary Supervisory Controller Synthesis tools use a CPU, we might be able to increase performance if we perform it on a GPU instead. One method of achieving this is by integrating SCS into GPUexplore.

To achieve this we can use the state-space exploration already included in GPUexplore to index all states in the model and construct the transition matrix. Next, we can create a vector where we assign a 1 to the index of a state that is marked, and a zero everywhere else. We can then use repeated matrix multiplication between this vector and the resulting transition matrix to perform our reachability analysis. Lastly we can edit the transition matrix to disable any transition into bad and blocking states, and then repeat until the matrix stabilizes to synthesize our controller.

This proposed implementation should be achievable in an efficient manner, as both matrix multiplication and matrix editing are easily threaded homogeneously. However, the effectiveness and feasibility are still to be shown in future research.

References

[1] J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete event processes. SIAM Journal on Control and Optimization, 25(1):206–230, 1987.

[2] A. Wijs, M. Osama, GPUexplore 3.0: GPU accelerated state space exploration for concurrent systems with data, in: Proc. 29th International Symposium on Model Checking Software (SPIN 2023), in: LNCS, vol. 13872, Springer, 2023, pp. 188–197.

[3] Heemstra, J., Osama, M., Wijs, A. (2025). Towards End-to-End GPU Acceleration of PCTL Model Checking. In: Jansen, N., et al. Principles of Verification: Cycling the Probabilistic Landscape. Lecture Notes in Computer Science, vol 15261.

Acknowledgement This work is supported by NWO grant OCENW.M.23.403 for the GUESS project.

Simultaneous Exploration of Topological and Control Space in Production Systems

D. Sarsur[1,2,*], P. Pena[1], L. Alves[1], M. Goorden[2], M. Reniers[2]
[1]Universidade Federal de Minas Gerais, Brazil
[2]Eindhoven University of Technology, The Netherlands
[*]Email: danielsc@ufmg.br

1 Introduction

Modern production systems, such as Reconfigurable Manufacturing Systems (RMS), are characterized by a high degree of flexibility in both structure and operation. To cope with the resulting design complexity, Design Space Exploration (DSE) has emerged as a standard methodology. However, existing DSE approaches typically treat the physical topology and the control logic as independent optimization problems. In this work, we argue that these aspects are inherently interdependent and should therefore be optimized simultaneously.

2 Automated Design Space Exploration

We propose a framework for automated design space exploration of the design space, depicted in Figure 1. This process is driven by a Genetic Algorithm (GA). Each candidate design is encoded as a chromosome that specifies the set of active components, their interconnections, and the assignment of control logic. For each chromosome, a corresponding system model is automatically instantiated and evaluated through discrete-event simulation. The resulting performance measures are used as fitness values to guide the evolutionary search toward increasingly effective designs. When the stop criteria is met, the framework returns a recommended system design, detailing the selected components, their topological arrangement, and the associated control strategy.

3 Experiment and Results

An experiment with the proposed framework was conducted using a flexible manufacturing scenario comprising multiple sources, machines with heterogeneous processing times, buffers with varying capacities, and sinks. The total number of feasible instances exceeded several billions, therefore, the reservoir sampling technique was employed and a representative subset of 8,601,012 instances was considered. The performance metric considered was the average stay time of items in the system. In the initial generations of the genetic algorithm (GA), the population exhibited high average stay times, due to the random nature of the initial populations. The framework gradually converged toward a high-performance region of the design space, with the best design identified at the 14th generation. This solution is characterized by a highly parallel topology and the exclusion of the machine with the longest processing time. Moreover, the associated routing logic prioritizes production completion, avoiding cyclic flows of items within the system. To further corroborate the hypothesis that system performance depends critically on the coupling between physical structure and control strategy, a counter-test was conducted using the same optimized topology with an altered control logic. Under this configuration, performance deteriorated substantially, with the average stay time increase of 2.190%.

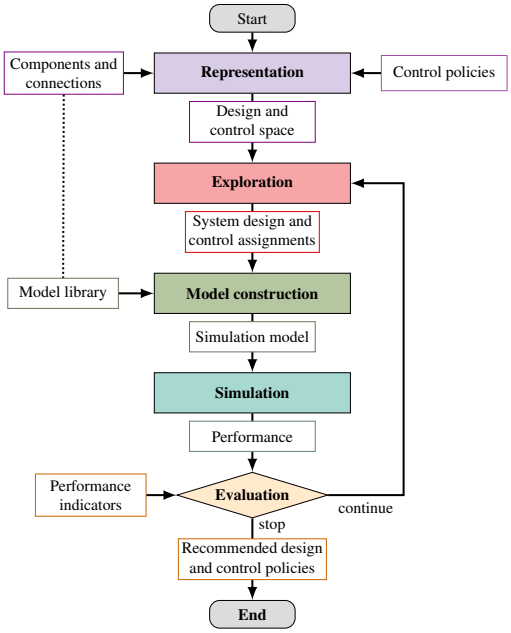

Figure 1: Proposed simultaneous optimization framework. Extended from [1].

References

[1] N. Paape, J. van Eekelen and M. Reniers, "Automated design space exploration for poultry processing systems using discrete-event simulation", *International Journal of Food Engineering*, 21 (11), 783-802, 2025.

Acknowledgment This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001 and Fapemig P&D APQ-06580-24.

A Unified Semantics-Based Framework for Timed Opacity

Zhe Zhang, Martijn Goorden, Michel Reniers

Department of Mechanical Engineering, Eindhoven University of Technology, The Netherlands
Email: z.zhang4@tue.nl, m.a.goorden@tue.nl, m.a.reniers@tue.nl

1 Introduction

In real-time cyber-physical systems, *timed opacity* ensures that a system's secret behaviors remain plausibly deniable to an outside intruder with partial observability of states, events, and precise timing information. The current landscape is fragmented: existing works often define timed opacity based on decidability convenience rather than practical security constraints. In this work, focusing on timed automata, we define a semantic baseline to capture the essence of timed opacity against such a powerful intruder.

2 A Unified Framework and Hierarchy

We formalize two notions to capture timed opacity. Current state timed opacity (CSTO) is a state-based notion where the secret is the current state of the system (e.g., being in secret state s_4), as shown in Fig. 1 (Left). Evolution-based timed opacity (EBTO) protects the secrecy of the entire execution history. As shown in Fig. 1 (Right), a secret is defined over the path itself (e.g., ρ_1, ρ_3), regardless of the final state.

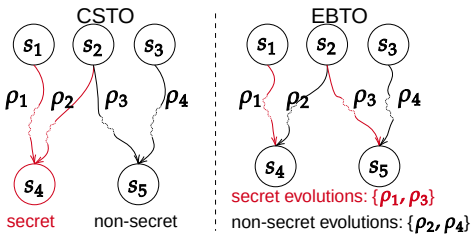

Figure 1: Comparison between CSTO and EBTO

To bridge these notions, we develop a structural transformation from EBTO to CSTO (Fig. 2). This process resolves the ambiguity of history dependence by unfolding, reconstructing, and renaming the underlying semantic graph, while preserving behavioral equivalence.

Finally, we bridge EBTO with existing notions (CLTO, LBTO, ETO [1, 2]), culminating in the unified hierarchy in Fig. 3. This demonstrates that existing definitions are *syntactic restrictions* within our framework.

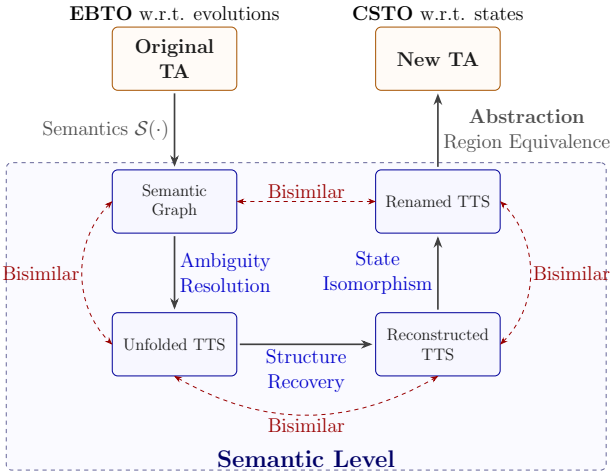

Figure 2: The transformation workflow

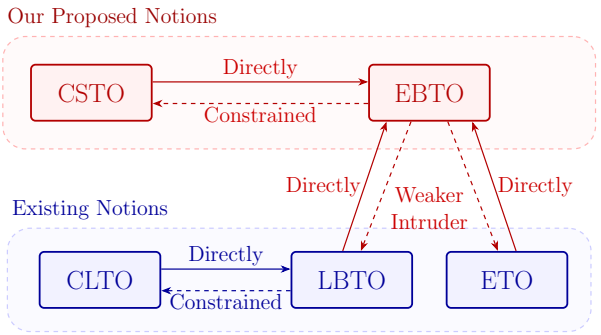

Figure 3: The hierarchy

3 Conclusion

The proposed framework provides a rigorous baseline to systematically analyze the sources of undecidability and unify existing decidable fragments.

References

[1] An J, Gao Q, Wang L, et al. The opacity of timed automata. In *Proc. of Int. Symp. on Formal Methods*. Springer, 2024: 620-637.

[2] André É, Duflot M, Laversa L, et al. Execution-time opacity control for timed automata. In *Proc. of Int. Conf. on Software Engineering and Formal Methods*. Springer, 2024: 347-365.

Supervisor synthesis for multilevel discrete-event systems with local buses

M.M. Baubekova, M.A. Goorden, M. A. Reniers,
J. M. van de Mortel-Fronczak, J. E. Rooda, W. J. Fokkink
Control Systems Technology Section, Eindhoven University of Technology
Email: {m.baubekova, m.a.goorden, m.a.reniers,
j.m.v.d.mortel, j.e.rooda, w.j.fokkink}@tue.nl

1 Introduction

Cyber-physical systems rely on supervisory controllers to guarantee correct and safe system behavior. The design of these controllers is becoming increasingly complex due to higher performance requirements, growing system variability and complexity, and a shortage of skilled engineers. Synthesis-based engineering (SBE), based on supervisory control theory introduced in [1], provides a formal framework to synthesize supervisors for discrete-event systems (DES) that constrain system behavior according to specified requirements. However, synthesizing supervisors is limited by high computational memory demand. Several attempts have been made to address this problem, one of which is multilevel supervisory control synthesis [2]. In multilevel supervisor synthesis, dependency structure matrix (DSM) techniques can be used to construct a tree-structured hierarchical decomposition of the synthesis problem and thus efficiently synthesize local supervisors. A bus component, which has many dependencies across a system, tends to lead to an undesirable clustering of many components in one synthesis subproblem. Prior work in [3] showed how to recognize and properly treat a global bus structure, which can be beneficial for systems that have clear bus components. For the rest of the cases, the multilevel discrete-event system (MLDES) without a global bus produced better results. As a performance measure, the total number of controlled states in the supervisors is used.

2 MLDES with local buses

This abstract leverages this work from global to local bus structures through a novel MLDES architecture. Specifically, the hierarchical system decomposition is revisited by allowing bus detection not only on the top level but at each level of the system hierarchy. Figure 1 depicts an example of the clustered DSM with local buses of a production cell model with a robot *Ro* with two arms *A*1 and *A*2 moving objects from the elevating rotary table *TaV* and *TaH* to the press *Pr*. In the DSM *A*1 is a global bus having the strong dependencies with other components, and *Ro* is a local bus of the cluster with *TaV*, *Pr* and *A*2. Figure 2 illustrates the resulting MLDES, with the hierarchy guided by the clustered DSM. The nodes for which the supervisor is synthesized include plant components *G*, the number of control requirements that reference them $|R|$, and the controlled state-space size of the resulting supervisor *css*. Both the global and local buses are distributed over the tree and placed as low as possible. This propagation of the bus components to the lowest feasible level relieves the upper nodes from excessive synthesis effort, adding a marginally increased effort at lower levels. The total number of states in the controlled MLDES decreased from 125,880 to 5,089 by introducing the local buses, which also outperforms the best result in [3] of MLDES without bus with a total *css* of 18,118.

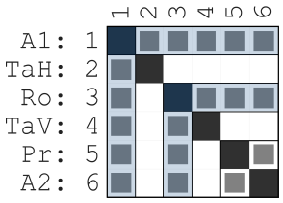

Figure 1: DSM-based clustering of a production cell. The local bus for each cluster is highlighted in blue.

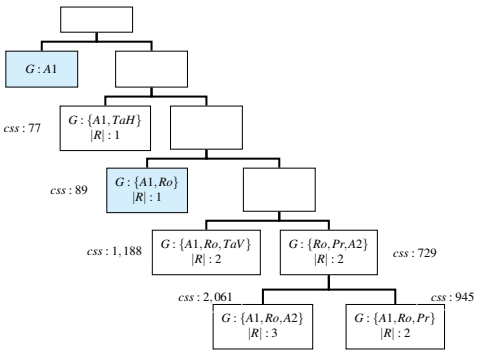

Figure 2: The MLDES with local buses (highlighted in blue).

References

[1] P.J. Ramadge and W.M. Wonham, "Supervisory control of a class of discrete event processes," SICON, 1987.

[2] J. Komenda, T. Masopust, and J.H. van Schuppen, "Control of an engineering-structured multilevel discrete-event system," IFAC WODES, 2016.

[3] M. Goorden, C. Dingemans, M. Reniers, J. van de Mortel-Fronczak, W. Fokkink, and J. Rooda, "Supervisory control of multilevel discrete-event systems with a bus structure," ECC, 2019.

Control-oriented modelling and control with stochastic discrete pellet actuation for fusion reactors

Lennard Ceelen[1,2,†], Christopher Orrico[1,2], Loes Jansen[1,2], Thomas Bosman[1,2], Matthijs van Berkel[1,3]

[1] DIFFER – Dutch Institute for Fundamental Energy Research, Eindhoven, the Netherlands

[2] Eindhoven University of Technology, Control Systems Technology, Eindhoven, the Netherlands

[3] Eindhoven University of Technology, Systems and Control, Eindhoven, the Netherlands

[†] l.ceelen@differ.nl

1 Pellet injection for density control in fusion reactors

Plasma density control is essential for safe and efficient tokamak operation, and pellet injection is a key actuator for core fueling in next-step devices [1]. Pellet injection launches frozen fuel pellets at high speed into the plasma, where they rapidly ablate and deposit particles predominantly in the core; it is therefore expected to provide the primary fueling mechanism for fusion reactors. From a control perspective, however, pellet injection is inherently impulsive: a continuous pellet-rate request can be realised as a sequence of binary events, which complicates feedback design in the presence of safety constraints such as density limits. Rather than approximating pellet injection as a continuous input, one may treat it as a mixed-integer actuation problem, where the controller directly selects injection times [2]. This comes at the cost of increased computational complexity and does not scale well to multiple actuators. In addition, real injectors exhibit uncertainties and constraints: uncertain time-of-flight (ToF) delays, stochastic fueling effectiveness (e.g. pellet loss), and, for centrifuge-based systems, synchronisation and maximum-rate constraints [3]. These features motivate control-oriented models that explicitly account for discrete actuation, uncertainty, and stochasticity while remaining suitable for synthesis and identification.

In this work, we develop a modelling framework for pellet injection systems that captures these actuator effects and yields linear approximate models for controller design and identification experiment planning. The resulting approximations also clarify how actuation limits and uncertainty impact achievable closed-loop performance.

2 Generalised pellet injection system model

We model the pellet injector as a stochastic event-driven system mapping the requested pellet rate u^k [Hz] to delivered fueling events $\delta^k \in \{0,1\}$. The model comprises: (i) a pellet quantiser that converts u^k into discrete requests while enforcing maximum-rate and synchronisation constraints via a countdown period; (ii) a stochastic ToF delay implemented as a shift register to represent variable arrival times and multiple pellets in transit; and (iii) a stochastic fueling model, represented by a Bernoulli process, that captures pellet loss. This composition provides a systematic representation of actuator uncertainties relevant to feedback design and constraint handling.

3 Linear expected and reduced-order approximation

For controller synthesis and identification experiment design, we derive a linear belief model from the generalised description. The mapping from u^k to $\mathbb{E}[\delta^k]$ is obtained by combining three effects: (i) a synchronisation filter induced by the centrifuge quantisation, (ii) a ToF-delay distribution, and (iii) the fueling probability p_{fuel}. To support loop shaping, we also introduce an expected approximation based on moment matching: it preserves the DC gain and mean delay of the belief model, yielding a low-order gain-and-delay representation. We discuss validity regimes in terms of requested-rate bandwidth relative to the centrifuge frequency and the signal-to-quantisation-noise ratio.

4 Identification and control with stochastic pellets

Using the belief model, we design identification experiments that remain informative under discrete and stochastic pellet actuation, enabling estimation of a linear input–output model from pellet-rate request to density diagnostics. Based on the identified dynamics, we design a density feedback controller that regulates core density via the requested pellet injection rate, and integrate a safety filter to enforce edge-density limits with high probability despite actuation uncertainty and stochastic fueling. Overall, the proposed framework provides a practical route to constraint-aware PI density regulation with a pellet safety filter under reactor-relevant conditions.

References

[1] Baylor, L. R., et al. (2007). Pellet fuelling and control of burning plasmas in ITER. Nuclear Fusion, 47(5), 443-448.

[2] Lang, P. T., et al. (2022). Actuator Development Step by Step: Pellet Particle Flux Control for Single- and Multiple-Source Systems. Fusion Science and Technology, 78(1), 1-9.

[3] Orrico, C. A., et al. (2025). Predictive density profile control with discrete pellets, applied to integrated simulations of ITER. Nuclear Fusion, 65(7), 076041.

Aircraft Maintenance Scheduling via Reinforcement Learning

Margot Devillers[†], Marco Saerens[†] and Raphaël M. Jungers[†]
[†] ICTEAM institute, UCLouvain, firstName.lastName@uclouvain.be

1 Introduction

Aircraft maintenance scheduling aims to determine when maintenance activities should be performed to ensure fleet availability, regulatory compliance, and cost efficiency. This problem has regained attention due to the digital transformation of the aviation industry [1], which enables data-driven maintenance strategies to reduce downtime and operational costs.

Maintenance tasks are periodic and must be grouped into projects and scheduled over a planning horizon so as to maximize aircraft utilization while respecting task due dates. The main challenge lies in jointly selecting task groupings and maintenance start times under operational constraints.

Building on the Colored Petri Net (CPN) digital twin of the physical maintenance scheduling operations of an aircraft fleet proposed in [2], this work designs a Reinforcement Learning (RL) agent that interacts with the CPN environment to learn an optimal fleet-level maintenance scheduling policy.

2 System Modeling and Learning Framework

We consider a fleet of aircraft $\mathscr{A} = \{a_1, \dots, a_m\}$ and a set of maintenance tasks $\mathscr{K} = \{k_1, \dots, k_n\}$ over a discrete planning horizon $t \in \{0, \dots, H\}$ (days). For each aircraft a_i and task k_j, $d_{i,j}$ and $f_{i,j}$ denote the task duration and the maximum allowable interval between two consecutive executions, respectively. The number of aircraft simultaneously under maintenance is limited by a capacity L.

At each time t, action variables determine *when* aircraft a_i is sent to maintenance and *which* tasks are executed. The binary variable $u_i(t) \in \{0,1\}$ indicates whether a maintenance project is initiated, while $w_{i,j}(t) \in \{0,1\}$ selects task k_j. The fleet-level control action is $\mathbf{U}(t) = \{\mathbf{u}_i(t)\}_{i=1}^m$, with $\mathbf{u}_i(t) \in \mathscr{U} := \{u_i(t), \mathbf{w}_i(t)\}$. The state of aircraft a_i at time t is

$$\mathbf{x}_i(t) \in \mathscr{X} := \{x_i(t), \mathbf{r}_i(t), \mathbf{q}_i(t), \mathbf{c}_i(t), \delta_i(t)\},$$

where $x_i(t) \in \{0,1\}$ indicates whether a_i is under maintenance ($x_i = 1$) or operating ($x_i = 0$). The vectors $\mathbf{r}_i(t)$, $\mathbf{q}_i(t)$, and $\mathbf{c}_i(t)$ have n components: the binary variable $r_{i,j}(t)$ equals 1 if task k_j is included in the ongoing project; $q_{i,j}(t)$ is the last completion time of k_j; and $c_{i,j}(t) = (1 - x_i(t))(t - q_{i,j}(t))$ is the elapsed operating time since last execution of k_j, satisfying $c_{i,j}(t) \leq f_{i,j}$. The variable $\delta_i(t) \geq 0$ denotes the remaining maintenance duration, equal to zero when $x_i(t) = 0$ and to $\sum_{j \in \mathscr{K}'} d_{i,j}$ when a project starts, with

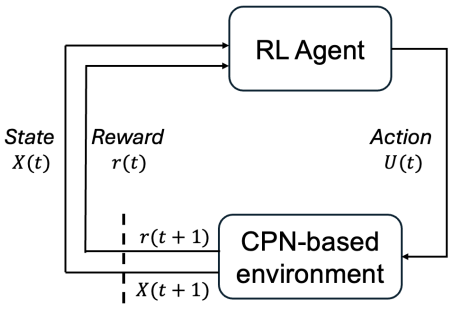

Figure 1: RL framework using CPN-based environment.

$\mathscr{K}' = \{k_j \in \mathscr{K} : r_{i,j}(t) = 1\}$. The objective is to minimize the total grounded time of the fleet over the planning horizon while respecting task interval constraints, leading to the optimal control problem

$$\begin{aligned} \min_{\mathbf{u}} \quad & \sum_{t=0}^{H} \sum_{i=1}^{m} \Big[x_i(t) + \sum_{j=1}^{n} x_i(t) \big(f_{i,j} - c_{i,j}(t) \big) \Big] \qquad (1) \\ \text{s.t.} \quad & \mathbf{x}_i(t+1) = f(\mathbf{x}_i(t), \mathbf{u}_i(t)), \quad \forall i \in \{1, \dots, m\}, \\ & \mathbf{x}_i(t) \in \mathscr{X}, \, \mathbf{u}_i(t) \in \mathscr{U}, \quad \forall i \in \{1, \dots, m\}, \end{aligned}$$

where $f : \mathscr{X} \times \mathscr{U} \to \mathscr{X}$ defines the system dynamics.

To jointly determine when aircraft should be sent to maintenance and how tasks should be grouped into maintenance projects over the horizon, this work proposes a framework that combines the flexibility of RL with the expressive power of Petri-net–based modeling. More precisely, the RL environment is defined as a CPN representing the aircraft maintenance scheduling process. Fig. 1 illustrates the resulting agent–environment interaction loop, where at each time step t, the agent observes the fleet state $\mathbf{X}(t)$, selects a control action $\mathbf{U}(t)$, and transitions to the next state $\mathbf{X}(t+1)$. The agent receives a reward derived from the cost function in (1), which penalizes aircraft downtime and deviation from task due dates.

References

[1] Mariza Tsakalerou, Dauren Nurmaganbetov, and Nurtay Beltenov. “Aircraft Maintenance 4.0 in an era of disruptions”. In: *Procedia Computer Science* 200 (2022).

[2] Lydia Mavraidi et al. “Aircraft Maintenance Scheduling based on Digital Twins”. In: *IFAC-PapersOnLine* 59 (2025).

A Hybrid Method for Multi-Agent Motion Planning on Industrial Magnetic Levitation Platforms

Bavo Tistaert, Stan Servaes, Alejandro Gonzalez-Garcia, Ibrahim Ibrahim,
Louis Callens, Jan Swevers, and Wilm Decré
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium
Email: bavo.tistaert@kuleuven.be

1 Introduction

Multi-agent systems excel in complex and large-scale settings by cooperating and/or distributing tasks across agents. However, safe and scalable motion planning of these systems remains a challenge in many real-time applications, such as transportation on industrial production lines.

While Model Predictive Control (MPC) has proven to be a powerful optimization-based method for motion planning, the classical centralized approach quickly becomes intractable as agents are added to the environment. The Alternating Direction Method of Multipliers (ADMM) [1] offers a decentralized, and hence more scalable, approach for convex decomposable optimization problems. However, due to the non-convexity of the collision avoidance constraints between agents, ADMM cannot guarantee safety for the problem at hand. To address this, we complement decentralized ADMM with a centralized safety filter, based on High-Order Control Barrier Functions (HOCBF) [2]. We implement and validate the resulting hybrid motion planner on the XPlanar Mover System, an industrial magnetic levitation platform from Beckhoff.

2 Methodology

The goal of the motion planner is to steer every agent from its initial state to a given target state without collisions. To this end, we formulate an optimal control problem (OCP) to be repeatedly solved in an MPC loop. We then use ADMM to decouple this OCP and solve its dual in an iterative manner. Here, agents plan trajectories and communicate their results in an alternating fashion, allowing independent parallel calculations. This procedure continues until the agents reach a consensus, i.e., the collision avoidance constraints are satisfied. To avoid repeated calculations, we only perform a fixed number of iterations m per time step, in order to reach a consensus over time.

In practice, due to the non-convexity of the problem, agents do not always reach a consensus. Therefore, we formulate a quadratic program that centrally modifies the input, generated by the m ADMM iterations, in a minimally invasive manner while adhering to a linear HOCBF constraint. This safety filter guarantees collision avoidance while being computationally tractable.

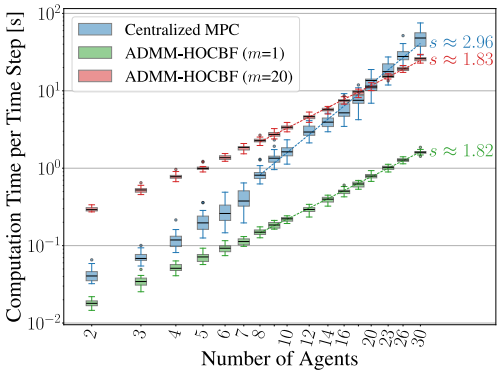

Figure 1: Scalability of ADMM-HOCBF ($m = 1$ and $m = 20$) and centralized MPC.

3 Results and experiments

We perform simulations for an increasing number of agents, using centralized MPC and the hybrid method (ADMM-HOCBF) and plot the average computation time per time step in Fig. 1. It is shown that ADMM-HOCBF reduces scalability with more than an order. The trajectories are computed in real time with a C++ implementation and subsequently deployed on the XPlanar system without collisions.

References

[1] S. Boyd et al., "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers", Foundations and Trends in Machine Learning, 2011.

[2] W. Xiao et al., "High-order control barrier functions", IEEE Transactions on Automatic Control 67.7, 2021.

Acknowledgment: This work was performed within the framework of Flanders Make's SBO projects ARENA and LearnOpTra.

Symbolic–Geometric Task and Motion Planning for Heterogeneous Dual-Manipulator Systems

Mehran Raisi Hatmabadi[1,2], Gaoyuan Liu[1], Mohayad Omer[1], Hamed Firouzipouyaei[1], Ann Nowé[2], Bram Vanderborght[1]

[1]Brubotics, Vrije Universiteit Brussel, [2]AI Lab, Vrije Universiteit Brussel (VUB), Belgium
Email: mehran.raisi.hatmabadi@vub.be

1 Introduction

Heterogeneous dual-manipulator systems combine robots with non-interchangeable capabilities (e.g., a fingered gripper and a suction cup) to enable cooperative manipulation such as handovers and staged insertions (Fig. 1). Compared to homogeneous teams, planning is affected by two properties: (i) actions are bound to specific robots due to capability constraints, and (ii) cooperation induces synchronization requirements, as some actions are only feasible when robots are simultaneously available. Consequently, allocation, precedence, and coordination constraints are embedded in the task structure and must be represented explicitly. In this work, we focus on the modeling layer and argue that PDDLStream [1] provides a symbolic–geometric backbone for capturing these constraints while preserving geometric executability.

2 Problem Formulation and Validation

We consider two manipulators $\mathscr{R} = \{R_1, R_2\}$ operating in a shared workspace and manipulating objects $\mathscr{O} = \{o_1, \ldots, o_N\}$. A symbolic–geometric planner generates a plan $\pi = [a_1, \ldots, a_K]$, where each action a_k is assigned to a robot $r(a_k) \in \mathscr{R}$ and parameterized by continuous variables (grasps, placements, handover configurations) that must satisfy collision-free inverse kinematics. In PDDLStream, such parameters are generated and validated through procedural streams [1]. Robot–action binding is encoded via robot-specific feasibility predicates, while synchronization is expressed by cooperative operators (e.g., *pass*) whose preconditions require stream-certified configurations reachable by both robots.

A central modeling challenge arises in hierarchical assembly tasks (e.g., peg–hole, cylinder–slot), where the feasibility of inserting a child object depends not only on the parent identity but on the *sampled pose* at which the parent was placed. Because PDDLStream may sample valid parent placements, a symbolic parent–child relation can lead to inconsistent plans unless this dependency is enforced. To address this issue, we introduce a geometric consistency predicate SSafe certified by streams. The predicate $\mathsf{SSafe}(o_c, p_c, o_p, p_p)$ holds if the child object o_c at pose p_c is inserted relative to the sampled parent pose p_p used when placing o_p, and the insertion stream validates compatibility at that pose. Propagating SSafe ensures coherence between symbolic assembly relations and sampled continuous configurations while retaining PDDLStream's search.

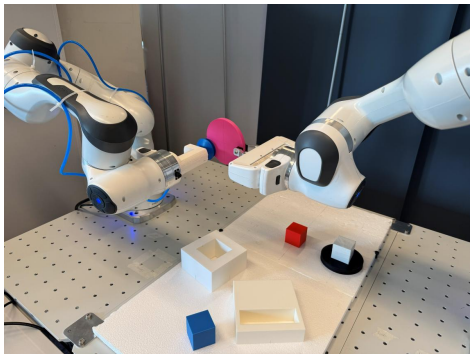

Figure 1: Heterogeneous cooperation: handover and reorientation enable insertion tasks.

Table 1: Sequential PDDLStream execution performance.

# Objects	# Time Slots	Execution Time (s)
2	4	111.8 ± 9.1
4	8	217.6 ± 12.9
6	12	330.5 ± 18.7

The formulation was validated in simulated dual-manipulator scenarios involving end-effectors. In all cases, SSafe prevented inconsistent parent–child insertions while preserving executability. Table 1 reports the execution performance of PDDLStream plans.

Acknowledgment

This work was supported by the Flemish Government under the program *Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen*, and euROBIN (No. 101070596).

References

[1] C. R. Garrett, T. Lozano-Perez, and L. P. Kaelbling, "PDDLStream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning," *Proc. ICAPS*, 2020.

ImpaC++t: A C++ Toolchain for Seamless Deployment of Nonlinear Model Predictive Control with Impact

Mathias Bos, Wilm Decré, Jan Swevers
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium
Email: mathias.bos@kuleuven.be

1 Introduction

The gap between high-level mathematical modeling of Nonlinear Model Predictive Control (NMPC) and its real-time execution on physical hardware remains a bottleneck in robotics and mechatronics research. While the Impact tool [1] automates the step from NMPC formulation to generation of efficient C-code via CasADi and Rockit, the user is often still tasked with writing complex "boilerplate" code to handle real-time scheduling, hardware I/O, and testing the compiled code in simulation. As an extension to the existing Impact artifact generation, which includes C-code, Simulink blocks and ROS2 nodes [2], we present a C++ deployment and simulation template that builds on the generated NMPC C-code artifacts and that provides a standardized interface for real-time Linux environments.

The template provides an abstracted hardware/simulation interface, dynamic loading of MPC artifacts to eliminate excessive recompilation time during prototyping, and a low-jitter threading model optimized for Linux with PREEMPT_RT.

2 Example usage

The following code snippet illustrates the simple usage of the presented tool for a double integrator MPC example. The full example is included in the ImpaC++t repository at `https://gitlab.kuleuven.be/u0110026/impact-cpp`.

```
// [...] Includes, define constants, etc.
// 1. Define Simulation model via CasADi
SimulationModel simulation_model() {
    using namespace casadi;
    MX x = MX::sym("x", 2), u = MX::sym("u", 1);
    Function f = Function("f", {x, u}, {vertcat(x(1), u
        )});
    return {f, Function("h", {x}, {x})};
}

int main() {
    // 2. Thread-safe data exchange via Shared Vectors
    SharedVector measurement{0., 0.};
    SharedVector control_input{0.};
    // 3. Initialize Simulator and Estimator
    Simulator sim(simulation_model(), measurement,
        control_input);
    FullStateFeedbackEstimator estimator(2);
    // 4. Dynamic loading of NMPC artifact and RT
        deployment
    ImpactMPC controller("double_integrator",
        measurement, control_input, estimator,
        CTRL_PERIOD_NS);
    // 5. Start asynchronous low-jitter threads (
        SCHED_FIFO)
    sim.run_sim_thread(RT_ENABLE);
    controller.run_ctrl_thread(RT_ENABLE);
    usleep(10e6); // Execute RT loop for 10s
    return 0;
}
```

3 Example application

The toolchain is illustrated on the example case of a Furuta pendulum, a pendulum mounted on a rotary base. Figure 1 shows the simulation results of a swing-up maneuver to bring the pendulum angle θ_2 from its bottom position at $\theta_2 = 0$ rad to the upright position at $\theta_2 = \pi$ rad. This is achieved with a cost driving the sum of potential and kinetic energy or total energy E (scaled, hence units s^{-2}), to the total energy in the upright position, balanced with a cost on the deviation from the reference state and control input.
By implementing this use case as a hardware experiment on a CANopen over EtherCAT (CoE) driven Furuta pendulum setup, we demonstrate a direct path for researchers to move from Python-based NMPC prototyping to industrial-grade deployment on hardware platforms commonly found in factory automation and advanced robotics.

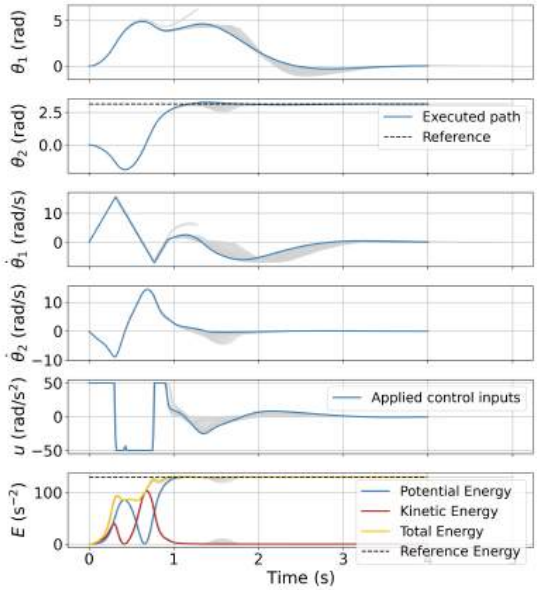

Figure 1: Simulation results of Furuta pendulum swing-up.

References

[1] A. Florez, A. Astudillo, W. Decré, J. Swevers, and J. Gillis, "IMPACT: A Toolchain for Nonlinear Model Predictive Control Specification, Prototyping, and Deployment," IFAC-PapersOnLine, vol. 56, no. 2, pp. 3164–3169, 2023

[2] A. Astudillo, A. Florez, W. Decré, and J. Swevers, "Rapid Deployment of Model Predictive Control for Robotic Systems: From IMPACT to ROS 2 Through Code Generation," in 2024 IEEE 18th International Conference on Advanced Motion Control (AMC), Kyoto (Japan), pp. 1–6

Acknowledgment This work was supported by Flanders Make through INFRA Master Plan 2025 - Intralogistics Physical.

Fast Motion Planning Using Rectangular Corridor Representations of Structured Environments

Alejandro Gonzalez-Garcia, Sonia De Santis, Jan Swevers and Wilm Decré
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium.
Flanders Make@KU Leuven, Belgium
Email: alex.gonzalezgarcia@kuleuven.be, sonia.desantis@kuleuven.be, jan.swevers@kuleuven.be, wilm.decre@kuleuven.be

1 Introduction

Autonomous Mobile Robots (AMRs) are increasingly deployed in large-scale structured environments such as factories, warehouses, hospitals, and agricultural facilities. A central challenge in autonomy is global motion planning, where paths or trajectories must be computed efficiently while respecting kinematic constraints. Common grid-based approaches, including A*, Hybrid-A*, and lattice-based planners, are widely used in practice. However, their computational performance degrades as map size and resolution increase. In this work, we propose a scalable motion planning framework for non-holonomic AMRs that leverages a corridor-based representation to achieve fast, kinematically-feasible planning. By operating on a compact topological abstraction of the environment, the proposed approach decouples planning complexity from map resolution, enabling consistent computation times in large, structured environments.

2 Methodology

The proposed framework decomposes the environment into a set of connected corridors and performs planning on the resulting topological graph. A sequence of corridors is selected via graph search, after which a kinematically-feasible analytical trajectory is generated within the corridor constraints [2] (see Figure 1 for an example of the planning framework).

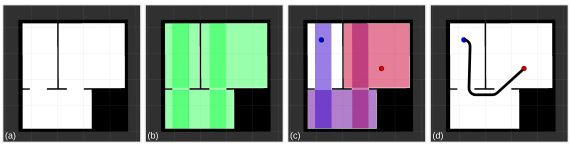

Figure 1: Corridor-based motion planning framework: (a) input occupancy grid, (b) corridor decomposition, (c) planned corridor sequence, and (d) generated analytical trajectory.

3 Results

The proposed framework was validated in simulation against open-source planners from the ROS 2 Nav2 stack, including A*, Hybrid-A*, and the Smac State Lattice Planner [1]. Figure 2 shows an example planning query, while Table 1 reports average planning times for increasing path lengths. The proposed method achieves consistently low planning times across all scenarios, remaining largely independent of map size and path length. These results indicate the proposed approach scales with the complexity of the corridor topological graph, instead of the grid map size or resolution.

Table 1: Average planning time (ms) for increasing path lengths. Short: < 10 m, Medium: 10–25 m, Long: > 25 m.

Planner	Short	Medium	Long
A*	9.4	25.5	61.5
SLP	24	126	389
H-A*	101	210	530
Ours	19.2	22.2	26.7

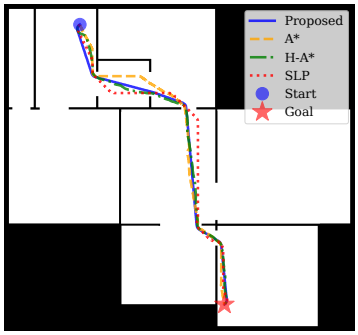

Figure 2: Illustrative motion planning comparison.

References

[1] S. Macenski et al., "The Marathon 2: A Navigation System," 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2020.

[2] S. De Santis et al., "Towards the Analytical Computation of Time-Optimal Trajectories for Unicycle Robots in Corridor Environments," 2024 European Control Conference (ECC), 2024.

Acknowledgment This work was done within the framework of Flanders Make's SBO project ARENA (Agile & Reliable Navigation).

Co-tuning of Path Planner and Controller with Bayesian Optimization

Carolina Vissers, René van de Molengraft, Michelle Chong, Elena Torta
Department of Mechanical Engineering, Eindhoven University of Technology
{c.m.e.vissers, m.j.g.v.d.Molengraft, m.s.t.chong, e.torta}@tue.nl

1 Introduction

Autonomous movement of robotic systems in dynamic environments with humans is a continuing area of research, as unpredictable obstacles require real-time trajectory adaptation to avoid collisions and safely reach goals. In this abstract, we focus on challenges of safe autonomous motion in crowded, dynamic environments with a specific focus on manipulator robots in operating room environments. We approach motion planning in dynamic environments with a manipulator that occupies a large workspace (namely the Image Guided Therapy (IGT) system in Figure 1). The goal is to navigate to a desired end position while avoiding collisions with static and dynamic obstacles.

Given a snapshot of the environment configuration, a path is computed in configuration space via sample based path planning (SBPP). Online trajectory generation is realized through a Model Predictive Controller (MPC) for local adaptation to environmental changes with respect to the static configuration used during the planning phase. In the dynamic environment, the reference path may become infeasible and lead the MPC to become stuck in local minima. This requires a state machine to coordinate the system. The full system architecture is seen in Figure 2, and is tested in simulation.

It is already known that MPC benefits from parameter tuning to result in more comfortable trajectories [1], better computational time, or many other desirable metrics. While this is often done by hand, tuning the parameters can be unintuitive and time-consuming. Automating this process, reduces much of the mental load and guess work involved in tuning.

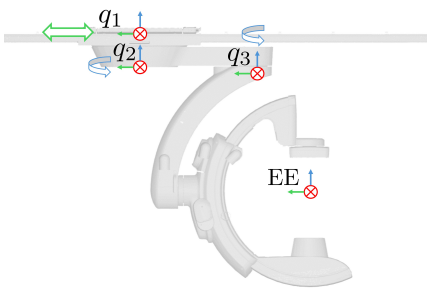

Figure 1: The IGT robot with the prismatic joint q_1, two revolute joints q_2 and q_3, and the end-effector (EE) position.

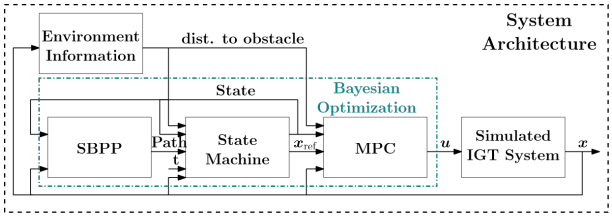

Figure 2: Schematic overview of the integrated system architecture, with modules tuned by Bayesian Optimization.

One such method is by using Bayesian Optimization (BO), where a cost function is defined via performance metrics and used to tune related MPC parameters [1]. The same principle applies to SBPP, where tuning the parameters chosen in SBPP can result in better paths, and faster results [2]. However, this previous work [1, 2] does not utilize a system-level perspective and seeks to only improve one part of a larger framework. This can create limitations where one part of the framework is optimized to the detriment of another module of the system, potentially reducing overall performance.

2 Proposed Approach

We propose a novel, system-level approach to co-tune SBPP, MPC, and the state machine using BO. The parameters have been selected for their effects on multiple modules to be optimized. Parameters from all 3 modules are co-tuned, as seen in the highlighted section of Figure 2. Examples of relevant performance metrics are overall completion time, MPC computation time, and joint zero-crossing acceleration rate. We hypothesize that co-tuning will show an improvement in the metrics, more than by tuning single modules.

References

[1] A. van der Horst et al., "A Bayesian Optimization Framework for the Automatic Tuning of MPC-based Shared Controllers," in 2024 IEEE International Conference on Robotics and Automation (ICRA), May 2024, pp. 11259–11265. doi: 10.1109/ICRA57147.2024.10610655.

[2] M. Moll, C. Chamzas, Z. Kingston, and L. E. Kavraki, "HyperPlan: A Framework for Motion Planning Algorithm Selection and Parameter Optimization," in 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Sept. 2021, pp. 2511–2518. doi: 10.1109/IROS51168.2021.9636651.

Kappa Motion Planner: A Python Framework for Analytical Motion Planning in Structured Planar Environments

Sonia De Santis, Wilm Decré and Jan Swevers
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium
Email: sonia.desantis@kuleuven.be,
wilm.decre@kuleuven.be, jan.swevers@kuleuven.be

1 Introduction

This work introduces *Kappa Motion Planner*, an open-source Python framework[1] for analytical motion planning of mobile robots in structured planar environments. The framework computes near-time-optimal trajectories with low computational effort and supports velocity-controlled unicycle [1] and bicycle [2] kinematic models.

The planner assumes that the environment is modeled as a set of connected rectangular corridors representing the available free space, whose overlaps encode connectivity and allow large-scale maps to be described compactly. This representation enables efficient motion planning in structured environments such as warehouses, industrial facilities, and service areas.

2 Planning Framework

The motion planning problem is formulated by specifying a set of *ingredients*: the environment, the vehicle model, and the planning task. The environment is defined as an ordered sequence of corridors, each represented by a `CorridorWorld` object encoding its width, length, center position, and orientation.

The vehicle model is instantiated by selecting a kinematic abstraction, either a `Unicycle` or `Bicycle`, and specifying its geometric dimensions and velocity bounds, or by selecting a predefined model from the vehicle library; for simplicity, vehicle footprints are currently assumed to be circular. The planning task is defined by an initial and a final planar pose (x, y, θ).

All ingredients are combined through a `MotionPlanner` object, which provides an analytical method for computing kinematically feasible, near-time-optimal trajectories. The resulting solution is expressed as a sequence of motion primitives tailored to the selected vehicle model and can be resampled or visualized using the provided tools.

The package documentation provides several motion planning examples, illustrating alternative ways to specify the problem ingredients. These examples are runnable using the `arena` Python package after installation according to the provided instructions.

[1] https://github.com/soniadesantis/kappa-motion-planner

3 Illustrative Example

A minimal runnable code snippet is reported below, and the corresponding solution path is shown in Figure 1.

```
c1 = CorridorWorld(0.6, 2, [0.0, 0.0], pi/2)
c2 = CorridorWorld(0.6, 2.1, [0.75, 1], 0)
c3 = CorridorWorld(0.6, 2.1, [1.8, 0.25], -pi/2)
c4 = CorridorWorld(0.6, 2.1, [2.55, -0.8], 0)
c_list = [c1, c2, c3, c4]

veh = Unicycle(model='Rosbot circular')
start = [0.145, -0.617, pi]
end   = [3.295, -0.727, 3*pi/4]

mp = MotionPlanner(veh, c_list, start, end)
traj = mp.compute_trajectory_analytical()
```

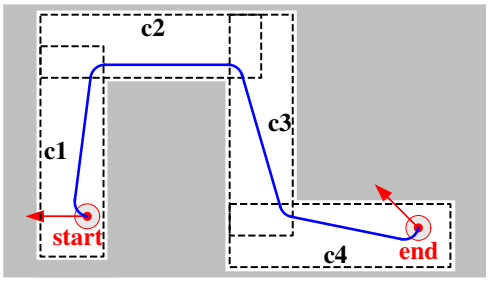

Figure 1: Path within four corridors.

References

[1] S. De Santis, A. Astudillo, W. Decré and J. Swevers, "Towards the Analytical Computation of Time-Optimal Trajectories for Unicycle Robots in Corridor Environments," in *Proc. 2024 European Control Conference (ECC)*, Stockholm, Sweden, 2024, pp. 3239–3246.

[2] S. De Santis, A. Astudillo, W. Decré and J. Swevers, "Towards an Analytical Time-Optimal Motion Planner for Bicycle Robots in Structured Environments," in *Proc. Benelux Meeting on Systems and Control*, Egmond aan Zee, The Netherlands, 2025.

Acknowledgment This work has been carried out within the framework of Flanders Make's SBO project ARENA (Agile & Reliable Navigation). Flanders Make is the Flemish strategic research centre for the manufacturing industry.

Asymmetry and Homogeneity in β-Divergence Recurrence Plots

Carlos Martínez, Martijn Boussé, Philippe Dreesen, Ralf Peeters
Department of Advanced Computing Sciences, Maastricht University, The Netherlands
Email: carlos.martinezgarcia@maastrichtuniversity.nl

1 Introduction

Recurrence Plots (RPs) constitute a powerful tool for revealing structural patterns in complex time series by measuring proximity in a reconstructed state space [1]. Traditionally, these similarities rely on standard geometric norms, such as the Euclidean (L_2) or Manhattan (L_1) metrics. While robust for general applications, these symmetric norms can be limiting when analyzing signals with heavy tails, outliers, or irreversible dynamics. Beyond these classical distances, a variety of alternative similarity criteria have been proposed in recent literature [1]. This diversity naturally motivates the use of β-divergences as a flexible dissimilarity family.

2 Methodology

As described in [2], β-divergences form a parametric family that continuously interpolates between the Itakura-Saito divergence ($\beta = 0$), Kullback-Leibler divergence ($\beta = 1$), and Squared Euclidean distance ($\beta = 2$). The β-divergence $d_\beta(x,y)$ is defined as:

$$d_\beta(x,y) = \begin{cases} \frac{\left(x^\beta + (\beta-1)y^\beta - \beta x y^{\beta-1}\right)}{\beta(\beta-1)}, & \beta \in \mathbb{R}\setminus\{0,1\}, \\ x\ln\frac{x}{y} - x + y, & \beta = 1, \\ \frac{x}{y} - \ln\frac{x}{y} - 1, & \beta = 0. \end{cases} \quad (1)$$

The utility of β-divergences in time series analysis stems from two properties that differ from standard norms.

Asymmetry implies that, for $\beta \neq 2$, $d_\beta(x,y) \neq d_\beta(y,x)$, which, among other features, enables the RP to capture temporal directionality and to discriminate time-reversed signals. **Homogeneity of degree** β ensures that the divergence scales as $d_\beta(\lambda x, \lambda y) = \lambda^\beta d_\beta(x,y)$, providing explicit control over how signal magnitude influences the recurrence structure.

3 Case Study: River Discharge

To illustrate these properties, we analyze river discharge data from the Lobith station, monitoring the Rhine River's entry into the Netherlands. The discharge data at Lobith were obtained from the public Rijkswaterstaat Waterinfo platform.[1] The corresponding signal and RPs are shown in Fig. 1.

[1] https://waterinfo.rws.nl/publiek/waterafvoer/Lobith(LOBI)/details

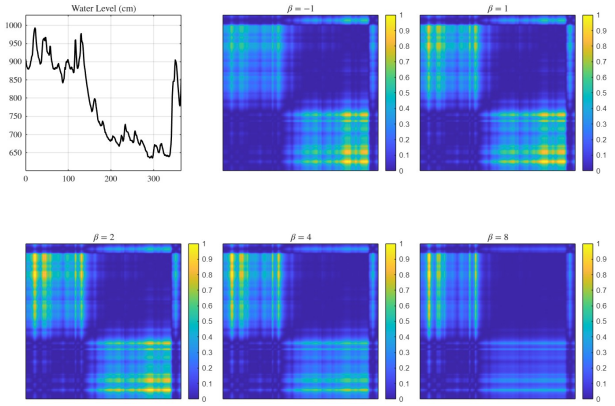

Figure 1: β-RPs of daily Rhine River discharge at Lobith over a 9-month period in 2018, including a drought period.

The **asymmetry** property is clearly visible in the top-left and bottom-right blocks of each RP. For lower values of β, the largest divergence values concentrate in the bottom-right block. Conversely, as β increases, the region of highest divergence shifts to the top-left block.

The **homogeneity** property dictates the visibility of features based on magnitude. As β increases the RP is able to capture some small differences among the high discharge values correspondig to the bottom-left square of the RP. In contrast, the top-right square reflects fluctuations within the low-flow drought period; these features are prominent at low β values and diminish as β increases.

These observations illustrate how the choice of β directly controls the sensitivity of the RP across different amplitudes. This highlights β-RP as a flexible extension of classical RPs, enabling an adaptive characterization of time series by selectively enhancing extreme events or low-amplitude fluctuations in skewed and heterogeneous signals.

References

[1] N. Marwan and K. H. Kraemer, "Trends in recurrence analysis of dynamical systems," *The European Physical Journal Special Topics*, vol. 232, no. 1, pp. 5–27, 2023.

[2] E. Deckert, P. Dreesen, N. Marwan, and M. Boussé, "Beta-Divergence-Based Recurrence Plots for Audio Time-Series Analysis," in *Proc. 33rd European Signal Processing Conference (EUSIPCO)*, 2025, pp. 291–295.

Recurrence-Based Analysis of EEG Signals for Seizure Detection Using β-Divergence

Filip Polakowski-Karol, Carlos Martínez, Martijn Boussé,
Philippe Dreesen, Joël Karel, Pietro Bonizzi, Ralf Peeters
Department of Advanced Computing Sciences, Maastricht University, The Netherlands

1 Introduction

Recurrence Plots (RPs) are a powerful tool for visualizing and quantifying the dynamical structure of time series. Traditionally, RPs rely on fixed metrics such as the Euclidean distance to define recurrence [1]. However, for complex, nonstationary, and noise-prone signals such as electroencephalography (EEG), these metrics may fail to fully capture the underlying dynamics or provide sufficient robustness [2]. In this work, we study β-divergence as a flexible alternative for RP construction [2].

2 Methodology

The β-divergence $d_\beta(x,y)$ is defined as:

$$d_\beta(x,y) = \begin{cases} \frac{\left(x^\beta + (\beta-1)y^\beta - \beta x y^{\beta-1}\right)}{\beta(\beta-1)}, & \beta \in \mathbb{R}\setminus\{0,1\}, \\ x\ln\frac{x}{y} - x + y, & \beta = 1, \\ \frac{x}{y} - \ln\frac{x}{y} - 1, & \beta = 0. \end{cases} \quad (1)$$

EEG recordings from the Bonn University dataset [3] are preprocessed by removing outliers using a Hampel filter. The signals are then decomposed into standard frequency bands (delta, theta, alpha, and beta) using FIR filtering. Finally, the data are normalized and shifted to be strictly positive to ensure compatibility with the β-divergence.

RPs are constructed using a time delay and an embedding dimension selected via Average Mutual Information and False Nearest Neighbors, respectively. To account for variations in distance distributions across different β values, recurrence thresholds are defined using the Median Absolute Deviation of the recurrence matrix. From each frequency band, 13 Recurrence Quantification Analysis features are extracted, yielding 52 features in total.

3 Results and Discussion

A Support Vector Machine (SVM) classifier was trained to distinguish between seizure and non-seizure segments. As shown in Fig. 1, the best seizure detection performance on clean (noise-free) data is achieved with $\beta = 5.2$, reaching an F1-score of 90.0%. This outperforms the squared Euclidean distance ($\beta = 2$), which achieved 86.8%.

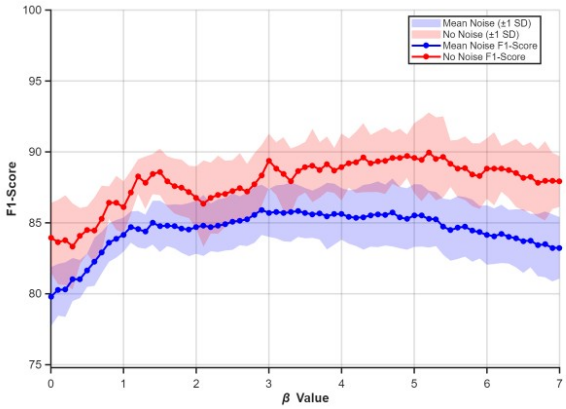

Figure 1: Comparison of F1-scores obtained under noise-free conditions and under additive Gaussian noise.

Under noisy conditions, simulated by additive Gaussian noise, the optimal operating point shifts towards lower values. In this scenario, $\beta = 2.9$ yields the best performance, achieving an F1-score of 85.9%, whereas the Euclidean case ($\beta = 2$) attains 84.7%.

4 Conclusion

The β-divergence RP approach improves seizure detection accuracy and robustness under noisy conditions compared to traditional Euclidean-based RPs. These findings suggest that β-divergences offer a more flexible and effective framework for signal analysis and time-series classification.

References

[1] N. Marwan, M. C. Romano, M. Thiel, and J. Kurths, "Recurrence plots for the analysis of complex systems," *Physics Reports*, vol. 438, no. 5-6, pp. 237–329, 2007.

[2] E. Deckert, P. Dreesen, N. Marwan, and M. Boussé, "Beta-Divergence-Based Recurrence Plots for Audio Time-Series Analysis," in *Proc. 33rd European Signal Processing Conference (EUSIPCO)*, 2025, pp. 291–295.

[3] A. Shankar et al., "Analysis of epileptic seizures based on EEG using recurrence plot images and deep learning," *Biomedical Signal Processing and Control*, vol. 69, 102854, 2021.

Parameter estimation in dynamical systems based on topology of interacting recurrence networks

Tamás Kovács
Eötvös University, Pázmány P.s. 1A, Budapest, Hungary, H-1117
tamas.kovacs@ttk.elte.hu

1 Abstract

The continuously increasing availability of high-resolution observational and experimental data calls for accurate estimation of parameters governing dynamical systems. Traditional statistical approaches and inverse modeling techniques have been widely developed to infer system parameters from measured data, ranging from maximum likelihood estimation and Bayesian inference to gradient-based optimization methods. In parallel, time-domain observations in the form of scalar time series, when transformed into complex network representations, have emerged as a powerful framework for investigating dynamical systems through their topological properties [1]. Recent studies have focused in particular on uncovering causal relationships and coupling directions between interacting dynamical subsystems using network-based methods [2, 3]. These developments have opened new avenues for system identification and parameter inference that complement classical optimization-based approaches, offering advantages particularly when dealing with nonlinear dynamics, non-stationary behavior, or systems where explicit mathematical models are unavailable or incomplete.

In this study, we present a novel parameter estimation approach for interacting dynamical systems based on recurrence network analysis. When subsystems are coupled, their dynamics deviate from those of isolated systems, leading to measurable variations in characteristic time scales and state space trajectories. The strength and nature of coupling between subsystems fundamentally alters the geometrical structure of trajectories in phase space, creating signatures that can be systematically exploited for parameter inference. These modifications manifest not only in the global properties of the reconstructed attractors but also in the local recurrence structure, which captures the temporal patterns of state space revisitations. By exploiting these coupling-induced effects, it is possible to reconstruct state vectors that evolve within a shared or entangled phase space, thereby capturing the essential features of the coupled dynamics. Transforming the reconstructed trajectories of the interacting subsystems into recurrence-based network representations, we demonstrate that the topological asymmetries and coupling directions between the resulting subnetworks encode information about the relative system parameters [4, 5]. Specifically, we show that network-based measures derived from the recurrence analysis, including clustering coefficients, transitivty, and path-length statistics, can be systematically related to the parameter ratios characterizing the coupled system, thereby enabling qualitative parameter estimation without requiring explicit dynamical equations or detailed knowledge of the underlying governing mechanisms. The method is particularly suited to scenarios where observational constraints limit direct access to all state variables, as is commonly encountered in real-world applications. As a contemporary application, we illustrate the method by estimating planetary mass ratios in multi-planet exoplanetary systems from observed time-domain variations. This application demonstrates the practical utility of the approach in astrophysical contexts where direct measurements are challenging and observational data are limited to indirect signatures of gravitational interactions, such as transit timing variations or radial velocity perturbations.

References

[1] Donner, R.V., et al., International Journal of Bifurcation and Chaos, vol. 21, issue 04, p. 1019 (2011)

[2] Feldhoff, J.H., et al., Physics Letters A, Volume 376, Issue 46, p. 3504-3513. (2012)

[3] Feldhoff, J.H., et al., Europhysics Letters, Volume 102, Issue 3, article id. 30007 (2013)

[4] Kovács, T., Chaos: An Interdisciplinary Journal of Nonlinear Science, Volume 29, Issue 7, id.071105 (2019)

[5] Kovács, T., Monthly Notices of the Royal Astronomical Society, Volume 491, Issue 3, p.3137-3154 (2020)

Recurrence analyses of event time-series

Norbert Marwan
Research Deptartment of Complexity Science
Potsdam Institute for Climate Impact Research (PIK), Member of the Leibniz Association
P.O. Box 60 12 03, 14412 Potsdam, Germany
Email: marwan@pik-potsdam.de

The analysis and modelling of event time-series are of broad interest in data science, as they are manifestations of the dynamics of the underlying system. Their relevance spans diverse scientific fields – from financial transactions and customer interactions to cardiac events, system failures, and natural phenomena – making them a ubiquitous and valuable subject of study. Event series may consist of discrete or binary events, or events with varying amplitudes, such as extremes, anomalies, or events derived from heavy-tailed distributions. Research questions for event data are often similar to those for continuous time-series, including data comparison, classification of underlying dynamical behaviour, detection of regime changes, and use in simulations or predictions. The critical challenge addressed here is the limitation of conventional time-series analysis methods when applied to sparse, irregular, or discrete event data (although several approaches address event-data analysis [2, 3]), hindering the ability to effectively compare these data with continuous records or reliably characterise their underlying spectral properties and nonlinear dynamics [5]. In this context, recurrence-based methods provide a promising alternative [5, 6].

A recurrence plot (RP) visualises pairwise similarity between all states $\vec{x}_i$ and $\vec{x}_j$: $R_{i,j} = \Theta(\varepsilon - d(\vec{x}_i,\vec{x}_j))$, where $d(\cdot,\cdot)$ is a distance- or cost-based similarity measure and $\Theta(\cdot)$ the Heaviside function, yielding $R_{i,j} = 1$ if the similarity (e.g., defined by a spatial distance $d_{L2} = \|\vec{x}_i - \vec{x}_j\|$) falls below a threshold ε [6]. By selecting a similarity measure suitable for discrete data, RPs can also be constructed for event data. A natural choice is the edit distance, an extension of the Levenshtein distance [11]. For two event sequences $S_i = \{t_k^{(i)}\}_{k=1}^{N_i}$ and $S_j = \{t_k^{(j)}\}_{k=1}^{N_j}$ the edit distance d_{ED} is the minimal total cost of transforming S_i into S_j using deletions, insertions, and temporal shifts of events t_k. This metric has been successfully combined with recurrence analysis [1, 10] and allows many advanced data analysis tasks offered by the RP-based framework, thereby extending the powerful tools of nonlinear time-series analysis and phase space characterisation to discrete and event-based processes.

For example, it enables a direct comparison of event with continuous data. One important recurrence-based approach is based on joint recurrence: $JR_{i,j}^{x,y} = d_{ED}(S_i^x, S_j^x) \cdot \Theta(\varepsilon - d_{L2}(\vec{y}_i,\vec{y}_j))$, where x_i denotes an event series and y_i a continuous time-series. It quantifies simultaneous recurrences in both signals and, thus, provides a powerful indicator of generalised synchronisation [9]. This capability is essential for analysing coupled dynamical systems where different components may be measured using heterogeneous data types (discrete events vs.continuous variables). This would also allow future extensions of the JRP approach for detecting coupling directions or even causal links between systems [8]. Nevertheless, the method faces practical challenges such as parameter selection, dimension mismatch, alternative similarity measures, and sensitivity to noise, which will be discussed briefly. Edit-distance-based RPs also enable power-spectrum estimation for spiky signals [4] and facilitate analysis of irregularly sampled data [7], a common challenge, e.g., in astrophysics and palaeoclimate research.

These extensions are demonstrated on climate and palaeoclimate data, where extreme event-like records are analysed together with continuously recorded observations. The results demonstrate that recurrence-based techniques can uncover hidden synchronisation patterns and produce reliable spectral estimates despite the sparsity of the event data.

References

[1] A. Banerjee et al. In: *Nonlin. Proc. Geophys.* 28 (2021), pp. 213–229.

[2] J. Beirlant et al. *Statistics of Extremes: Theory and Applications*. Chichester, UK: Wiley, 2006, p. 512. ISBN: 978-0-470-01237-6.

[3] M. Ciba et al. In: *Neur. Comp.* 32.5 (May 2020), pp. 887–911. ISSN: 0899-7667.

[4] K. H. Kraemer et al. In: *Entr.* 24.11 (2022), p. 1689.

[5] N. Marwan. In: *Front. Appl. Math. Stat.* 9 (2023), p. 1129105.

[6] N. Marwan et al. In: *Phys. Rep.* 438.5–6 (2007), pp. 237–329.

[7] I. Ozken et al. In: *Phys. Rev. E* 91 (2015), p. 062911.

[8] A. M. T. Ramos et al. In: *Phys. Rev. E* 95 (2017), p. 052206.

[9] M. C. Romano et al. In: *Phys. Lett. A* 330.3–4 (2004), pp. 214–223.

[10] S. Suzuki et al. In: *Int. J. Bif. Chaos* 20.11 (2010), pp. 3699–3708.

[11] J. D. Victor and K. P. Purpura. In: *Netw. Comp. Neur. Syst.* 8.2 (Jan. 1997), pp. 127–164. ISSN: 0954-898X.

Exploring Collective Synchronisation in Football Matches Using Joint Recurrence Plots

Sebastian Hermann[1,2], Hendrik Meth[2] and Martin Lames[1]

[1]Technical University of Munich, Am Olympiacampus 11, 80809 Munich, Germany
[2]Stuttgart Media University, Nobelstr. 10, 70569 Stuttgart, Germany

Email: hermannsebastian@protonmail.com

1 Introduction

Recurrence Analysis (RA) is a common method to analyse complex dynamical systems [1]. In sports science, football is considered complex, dynamic systems shaped by players' continuously evolving interaction processes [2]. As such, RA is useful for investigating these processes, as it identifies recurring states in a system's phase space and provides insights into its temporal structures. The recurrence plot (RP), serving as its methodological foundation, reveals pairs of time points at which the system returns to a previous state [3]. Although researchers have mostly used RPs for individual or aggregated trajectories in football contexts [2], the broader literature on RA offers extensions addressing direct subsystem interactions [1]. Of particular value is the joint recurrence plot (JRP), which is created from the Hadamard product of two different RPs and thus shows when two subsystems exhibit recurrences simultaneously. Applied to football, the JRP can be used to capture and quantify collective synchronisation among football players. The aim of this work is to explore JRP's applications for identifying and analysing collective synchronisation in football matches.

2 Methods

Player-tracking data from the 2021/2022 German Bundesliga serve as the empirical basis for this study. To illustrate the use of JRPs, we analyse several player combinations from a single match. Each player's two-dimensional position at time t_i is represented by the state vector $\mathbf{x}(t_i) = (x(t_i), y(t_i)) \in R^2$. The RP is constructed as a binary matrix $R \in \{0,1\}^{N\times N}$, where an entry is defined as

$$R_{i,j} = \{\, 1 \,, \|\mathbf{x}(t_i) - \mathbf{x}(t_j)\| \leq \varepsilon, 0, otherwise, \qquad (1)$$

with ε denoting the spatial recurrence threshold. To analyse synchronisation between two players, the JRP is constructed as the element-wise product of their individual RPs, $JR_{i,j} = R^{(1)}_{i,j} R^{(2)}_{i,j}$ for $i, j = 1, \ldots, N$.

3 Results

Figure 1 shows a JRP that captures periods of simultaneous recurrence for two players over an entire match. Overall, the JRPs display a fine-grained, relatively uniform texture with short diagonal and vertical line structures. The diagonal lines indicate brief periods during which both players exhibit simultaneous, recurrent movement patterns. Vertical lines indicate short intervals in which both players maintain their positions across consecutive time steps. Arc-like line structures reflect simultaneous dynamics with variable temporal scaling, indicating recurring movement sequences executed at different speeds.

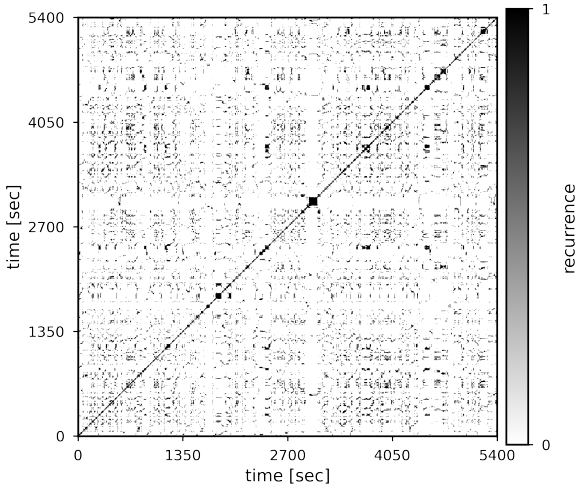

Figure 1: JRP of two players over a full match (Bayern Munich vs. Borussia Dortmund, 23 April 2022)

References

[1] N. Marwan, M. C. Romano, M. Thiel, and J. Kurths, "Recurrence plots for the analysis of complex systems," *Physics Reports*, vol. 438, no. 5–6, pp. 237–329, 2007.

[2] M. Lames, S. Hermann, R. Prüßner, and H. Meth, "Football match dynamics explored by recurrence analysis," *Frontiers in Psychology*, vol. 12, article 4125, 2021.

[3] J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle, "Recurrence plots of dynamical systems," *Europhysics Letters*, vol. 4, no. 9, pp. 973–977, 1987.

Stability Properties of Minimal Gated Unit Neural Networks[1]

Stefano De Carli, Davide Previtali, Mirko Mazzoleni
Dept. of Management, Information and Production Engineering
University of Bergamo, Dalmine (BG), Italy stefano.decarli@unibg.it

1 Introduction

Recurrent neural networks, such as Gated Recurrent Unit (GRU) networks, are widely adopted in system identification due to their ability to model nonlinear dynamical behaviors. However, their complexity poses barriers for embedded applications requiring low memory footprint and fast inference. Furthermore, standard training does not guarantee Input-to-State Stability (ISS) or Incremental ISS (δISS). These properties are vital not only for safe control deployment [3], but also to ensure that the identified model is physically consistent, preventing unbounded predictions in long-term simulations. In this work, we analyze *Minimal Gated Unit (MGU)* networks [1], lightweight architectures simplifying GRU networks. We extend the stability analysis framework from [3] to the MGU network, deriving sufficient parametric conditions for δISS, and propose a stability-promoting training strategy. Validation shows this yields safe models without sacrificing accuracy.

2 The MGU network Model

We consider an MGU network of L stacked layers. For each layer $l \in \{1,\dots,L\}$, let $n_h^{(l)} \in \mathbb{N}$ denote the state dimension and $n_{\tilde{u}}^{(l)} \in \mathbb{N}$ the input dimension. We denote the external input as $u_k \in \mathbb{R}^{n_u}$, where $n_u \in \mathbb{N}$ is the input size. The l-th layer (Figure 1) represents a discrete-time dynamical system where, for each time step $k \in \mathbb{N}$, the state $h_k^{(l)} \in (-1,1)^{n_h^{(l)}}$ evolves via a forget gate $f_k^{(l)} \in (0,1)^{n_h^{(l)}}$ and candidate state $c_k^{(l)} \in (-1,1)^{n_h^{(l)}}$:

$$f_k^{(l)} = \sigma\left(W_f^{(l)}\tilde{u}_k^{(l)} + R_f^{(l)}h_k^{(l)} + b_f^{(l)}\right),$$
$$c_k^{(l)} = \phi\left(W_c^{(l)}\tilde{u}_k^{(l)} + R_c^{(l)}\left(f_k^{(l)} \circ h_k^{(l)}\right) + b_c^{(l)}\right),$$

where $\sigma(\cdot)$ and $\phi(\cdot)$ are the element-wise sigmoid and tanh functions. The network parameters are $W_{\{f,c\}}^{(l)} \in \mathbb{R}^{n_h^{(l)} \times n_{\tilde{u}}^{(l)}}$, $R_{\{f,c\}}^{(l)} \in \mathbb{R}^{n_h^{(l)} \times n_h^{(l)}}$, and biases $b_{\{f,c\}}^{(l)} \in \mathbb{R}^{n_h^{(l)}}$. The layer input $\tilde{u}_k^{(l)} \in \mathbb{R}^{n_{\tilde{u}}^{(l)}}$ is defined as the external input u_k for $l = 1$ ($n_{\tilde{u}}^{(1)} = n_u$), and as the previous layer's state $h_{k+1}^{(l-1)}$ for $l > 1$ ($n_{\tilde{u}}^{(l)} = n_h^{(l-1)}$). The state update is a convex combination:

$$h_{k+1}^{(l)} = \left(1_{n_h^{(l)}} - f_k^{(l)}\right) \circ h_k^{(l)} + f_k^{(l)} \circ c_k^{(l)}.$$

[1]This work was partially funded by the EU European Defence Fund Project KOIOS (EDF-2021-DIGIT-R-FL-KOIOS). KOIOS is a project Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Defence Fund. Neither the European Union nor the granting authority can be held responsible for them.

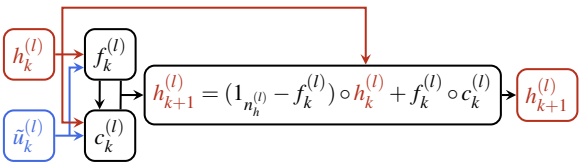

Figure 1: MGU l-th layer architecture at time k.

MGU layers require $\sim 33\%$ fewer parameters than GRUs.

3 Stability Analysis

We focus on the δISS property, which ensures that the distance between state trajectories is bounded by input differences. Since a cascade of δISS systems is δISS, we derive conditions for a single MGU layer. For notational simplicity, in the following we omit the layer index (l) when referring to generic layer quantities.

Theorem 1 (δISS of MGU layers). *Assume bounded inputs such that $\|u_k\| \le 1, \forall k \in \mathbb{N}$. The l-th MGU layer is δISS if:*

$$\bar{\sigma}_f + \bar{\sigma}_f^2 \|R_c\| + \frac{1}{4}\|R_f\|\left(\bar{\sigma}_f \|R_c\| + \bar{\phi}_c + 1\right) < 1, \quad (1)$$

where $\|\cdot\|$ is the infinity norm, and $\bar{\sigma}_f, \bar{\phi}_c$ are defined as:

$$\bar{\sigma}_f := \sigma\left(\left\|\begin{bmatrix} W_f & R_f & b_f \end{bmatrix}\right\|\right), \quad \bar{\phi}_c := \phi\left(\left\|\begin{bmatrix} W_c & R_c & b_c \end{bmatrix}\right\|\right).$$

Note that δISS analytically implies ISS. To train stable models, we employ a *loss augmentation* strategy, adding a penalty term to the objective function proportional to the violation of Condition (1). During training, we save the model that yields the best performance on a validation dataset while satisfying (1) for each layer $l \in \{1,\dots,L\}$.

4 Experimental Results

We evaluated the stable MGU networks on three benchmarks, following the protocols and architectures of the referenced papers. On the *pH Reactor* [2] and *4-Tank* [3], the stable MGU networks matched the test performance and noise robustness of stable GRU networks [3], despite using fewer parameters. Finally, on the *Silverbox* dataset [4], the stable MGU networks achieved a 6.88 mV RMSE on the *Arrow (full)* test set. Although higher than unconstrained baselines ($\sim 3 \div 5$ mV) reported in [4], the MGU networks ensure theoretical stability guarantees that standard models lack. Moreover, the stable MGU networks significantly outperform existing stable GRU networks (24.38 mV), which failed to converge under strict stability enforcement.

References

[1] G.B. Zhou et al., "Minimal gated unit for recurrent neural networks," *Int. J. Autom. Comput.*, vol. 13, no. 3, 2016.
[2] E. Terzi et al., "pH reactor dataset," *Tech. Rep.*, 2020.
[3] F. Bonassi et al., "On the stability properties of Gated Recurrent Units neural networks," *Syst. Control Lett.*, vol. 157, 2021.
[4] M.D. Champneys et al., "Baseline Results for Nonlinear System Identification Benchmarks," *arXiv:2405.10779*, 2024.

Distributed Deadlock-Avoidance Control for Low-Latency Industrial Multi-Robot Systems

Bin-Bin Hu[1], Brazil de Vries[1], Taco Hacquebord[2], Jorick van der Vlag[2]
Ming Cao[1], Bayu Jayawardhana[1]
[1]University of Groningen, 9747 AG Groningen, The Netherlands
[2]Philips, 9203 ZN Drachten, The Netherlands
Email: {b.hu@, b.de.vries.22@student., b.jayawardhana@, m.cao@}rug.nl,
{taco.Hacquebord@, jorick.van.der.vlag@}philips.com

1 Introduction

Over the years, developing effective deadlock-resolution strategies has remained an important yet challenging problem in industrial applications [1]. Although this problem has been extensively studied in multi-robot systems, most existing solutions rely on optimization that are ill-suited to industrial platforms requiring high-frequency updates and low-latency communication. Moreover, the lack of standardized optimization toolboxes for nonlinear optimization, together with bounded platform, further exacerbates these challenges [2]. To tackle this issue, we propose a distributed deadlock-avoidance controller for low-latency industrial multi-robot systems on the constrained platform.

2 Detailed Controller Design

We consider a team of N robots indexed by $\mathcal{V} = \{1, \cdots, N\}$. Each robot i is moving with the single integrator dynamics,

$$\dot{\mathbf{x}}_i = \mathbf{u}_i, i \in \mathcal{V}, \tag{1}$$

where $\mathbf{x}_i := [x_{i,1}, x_{i,2}]^\intercal \in \mathbb{R}^2, \mathbf{u}_i := [u_{i,1}, u_{i,2}]^\intercal \in \mathbb{R}^2$ are the positions and inputs, respectively. We also consider a team of N target points indexed by $\mathcal{T} = \{1, \cdots, N\}$, where the position of each static target i is denoted by $\mathbf{x}_i^d := [x_{i,1}^d, x_{i,2}^d]^\intercal \in \mathbb{R}^2$. For the bounded platform, we consider M inequalities $\mathcal{M} = \{1, \cdots, M\}$ below,

$$A_k^\intercal \mathbf{x} - b_k \leq 0, j \in \mathcal{M}, \tag{2}$$

According to the robot dynamic in (1), the target $\mathbf{x}_i^d$ and the boundaries in (2), we can design the distributed deadlock-avoidance controller below

$$\mathbf{u}_i = k_{i,a}\mathbf{F}_i^a + k_{i,r}\sum_{j\in\mathcal{N}_i}\mathbf{F}_{i,j}^r + k_{i,b}\sum_{k\in\mathcal{M}}\mathbf{F}_{i,k}^b, \tag{3}$$

where $k_{i,a} \in \mathbb{R}^+, k_{i,r} \in \mathbb{R}^+, k_{i,b} \in \mathbb{R}^+$ are the control gains, $\mathcal{N}_i$ is the sensing neighbor set, and $\mathcal{M}$ is the boundary set in (2). Note that $\mathbf{F}_i^a(t) \in \mathbb{R}^2, \mathbf{F}_{i,j}^r(t) \in \mathbb{R}^2, \mathbf{F}_{i,k}^b(t) \in \mathbb{R}^2$ represent the target-attraction, neighboring-repulsion, and boundary-repulsion forces, respectively.

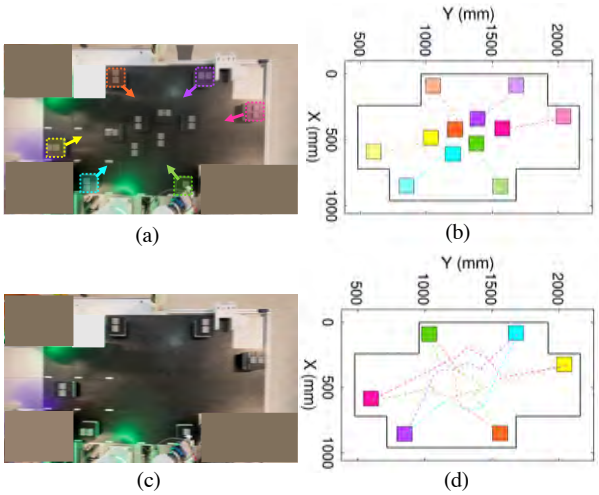

Figure 1: Experiments of six movers in the XPlanar platform. (a) Snapshots at $t = 2$s. (b) Trajectories at $t = 2$s. (c) Snapshots at $t = 7$s. (d) Trajectories at $t = 7$s.

3 Main results

We conduct the deadlock-avoidance experiments with six movers on the XPlanar platform at Philips. The system operates at a control frequency of 4k Hz to ensure the platform's safety. Figure 1 presents experimental snapshots together with the corresponding trajectories of the six movers, demonstrating that all movers successfully avoid deadlock and navigate to their designated target points.

References

[1] M. P. Fanti and M. Zhou, "Deadlock control methods in automated manufacturing systems," IEEE Trans. Syst., Man, Cybern. A: Syst. Humans, **34**(1), 5–22, 2004.

[2] J. Grover and C. Liu and K. Sycara, "The before, during, and after of multi-robot deadlock," Int. J. Robot. Res., **42**(6), 317–336, 2023.

Online error compensation of redundant resolvers

Louise Massager and Michel Kinnaert
SAAS, Université libre de Bruxelles
Av. Franklin Roosevelt 50, 1050 Brussels, Belgium
Emails: louise.massager@ulb.be, michel.kinnaert@ulb.be

1 Introduction

A specific sensor fault diagnosis problem is considered for an electromechanical actuator (EMA) aimed at aeronautic and aerospace applications. This EMA is equipped with two redundant sensors for measurement of the angular position of the shaft. More precisely, resolvers are used –namely electromechanical transformers– that provide sine and cosine AC voltages in the stator windings, from which the angular position θ can be recovered. Our aim is to detect and identify three types of faults, namely imperfect quadrature β (due to stator coils misalignment), amplitude imbalance α (due to insulation fault in the stator) and offsets $\Delta V_{sin,cos}$ (induced by the data acquisition system). As soon as the fault is identified, it can be compensated online.

2 Modeling

The sampled output signals of a wound resolver $R1$ at time instant k including the above mentioned faults are written:

$$\begin{aligned} V_{cos}^{R1}[k] &= M(1+\alpha)\cos(\theta[k]+\beta)+\Delta V_{cos}+\xi_{cos}^{R1}[k] \\ V_{sin}^{R1}[k] &= M\sin(\theta[k])+\Delta V_{sin}+\xi_{sin}^{R1}[k] \end{aligned} \quad (1)$$

In these equations, the amplitude M is known; the fault parameters (α,β and $\Delta V_{sin,cos}$) are zero in healthy operation; $\xi_{cos}[k]$ and $\xi_{sin}[k]$ are assumed to be zero mean Gaussian white noise sequences with known identical standard deviation σ_ξ. The angular position $\theta[k]$ of resolver $R1$ will be considered known and approximated by the measurement provided by resolver $R2$. Note that the fault parameters are assumed constant, which covers slowly evolving faults. Besides, only one resolver is assumed to present an uncompensated fault at a time.

3 Fault detection and identification

The fault detection and identification system for resolver $R1$ relies on the following fault indicator signal, called residual:

$$r^{R1}[k] = (V_{cos}^{R1})^2[k]+(V_{sin}^{R1})^2[k]-(M^2+2\sigma_\xi^2) \quad (2)$$

It was demonstrated in [1], through data analysis, that r^{R1} follows a Gaussian distribution of expected value μ_f:

$$\begin{aligned} \mu_f[k,\Phi] =& \Phi_1+\Phi_2\cos(2\theta[k])+\Phi_3\sin(2\theta[k]) \\ &+\Phi_4\cos(\theta[k])+\Phi_5\sin(\theta[k]) \end{aligned} \quad (3)$$

where the parameters Φ_i are function of the fault parameters and are zero as well in the absence of fault. Two scenarios are considered: ($\mathscr{H}_0$) no fault occurred up to the present time k; ($\mathscr{H}_1$) a fault occurred at a time $k_f < k$ causing a non-zero expected value of the residual (Eq. 3). The most likely scenario is estimated by a Generalized Likelihood ratio (GLR) algorithm based on a window of samples of fixed length. Error compensation consists of the following steps:

A The most likely profile parameters $\hat{\Phi}$ and the fault occurrence time $\hat{k}_f$ are estimated under hypothesis $\mathscr{H}_1$.

B A fault is then detected if the log likelihood ratio between $\mathscr{H}_1$ and $\mathscr{H}_0$ is higher than a fixed threshold.

C If a fault is detected in (B), fault parameter estimates $\hat{\alpha},\hat{\beta},\Delta\hat{V}_{sin},\Delta\hat{V}_{cos}$ are retrieved from $\hat{\Phi}$ obtained in (A).

D Based on the estimates of fault parameters in (C), error compensation is performed by exploiting model 1.

4 Results

Error compensation on a simulated study case is presented in figure 1. the scenario consists of a fault $\alpha^{R1} \neq 0$ after 0.6s in resolver $R1$ and another fault $\Delta V_{sin}^{R2} \neq 0$ emerging at 2.6s in resolver $R2$. One can observe that both faults are detected (i.e. h^{R1} and h^{R2} exceed the threshold on figures 1.a and 1.c) and compensated with parameters estimates (* on figures 1.b and 1.d). Error compensation is done a window length after the fault is detected to ensure that all samples in the window used for parameter estimation correspond to $\mathscr{H}_1$.

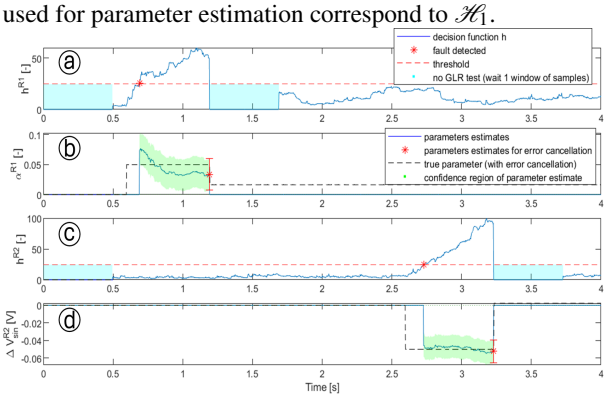

Figure 1: Error compensation of $\alpha^{R1} = 5\%$ & $\Delta V_{sin}^{R2} = -5\%$.

This research is part of the BLU-ULB Space Research Center and the Space4ReLaunch project, which is supported by the SPW Economie Emploi Recherche of the Walloon Region, under grant agreement no. 2210181.

References

[1] Massager, L., Meulenaer, J., Alexandre, P. & Kinnaert, M. Online fault detection, isolation and identification of redundant resolvers, submitted to the IFAC World Congress. *Available At SSRN 5701415*.

Towards off-road motion planning on elevated terrain

Patrick Scheffe, Wilm Decré, and Jan Swevers
MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium
Flanders Make@KU Leuven, Belgium
Email: patrick.scheffe@kuleuven.be

1 Introduction

Heavy-duty vehicles are piloted by humans in construction, forestry, agriculture, and mining. Substituting these human-driven vehicles with autonomous robots promises to reduce human workload and to increase productivity and safety.

Motion planning for autonomous robots through rough terrain with large elevation changes is challenging, since collisions with obstacles must be avoided and the robot-environment interaction needs to be accounted for to prevent hazardous situations such as rollover or getting stuck in nontraversable terrain. It is impossible to obtain perfect knowledge about this robot-environment interaction, which results in uncertainties in the motion planning task. These challenges render motion planning for robots on rough terrain an area of active research [1]. The problem setting of this article is to plan a motion for a robot with Ackermann steering geometry to safely and efficiently move from a start position to a goal position. We assume the elevation map of the robot's surroundings to be known.

2 Methodology and Results

Planning a motion from a start position to a goal position is intractable, as it is near impossible to model the robot-ground interaction accurately, and computational effort is immense. Therefore, we propose a hierarchical control structure. On the upper layer, a path planner first determines a path which avoids steep elevation changes based on the elevation map. There are many state-of-the-art algorithms for path planning, which trade off computation time with solution quality [1]. We realize path planning with an optimal rapidly-exploring random tree (RRT*) on a gridded traversability map, in which grid cells are flagged as non-traversable if their finite elevation difference is larger than a threshold. An example is given in Figure 1. On the left, an example elevation map is displayed. On the right, a planned path from the origin to the lower right that avoids large elevation changes is shown. This path planner disregards system dynamics to be computationally efficient.

On the lower layer of the hierarchical control structure is a path follower. Its foremost goal is to reach the goal safely and efficiently, i.e., without getting stuck or rolling over and in short time, rather than tracking the path closely. We created a simulation environment in MuJoCo, displayed in Figure 2, to develop and test such a path follower. The simulation environment features a robot with Ackermann steering geometry and suspension on all wheels, and supports friction. For low vehicle speed and elevation gain, the robot can reach the goal with a pure-pursuit algorithm. Our future work will look into traditional and learning-based approaches to control the robot in more realistic conditions.

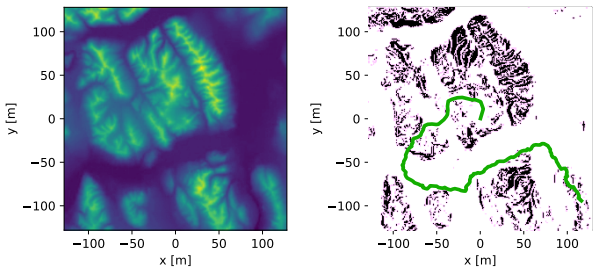

(a) Example elevation map with elevation changes of up to 10 m. (b) Path planned with RRT* [2] disregarding system dynamics.

Figure 1: Example elevation map and planned path.

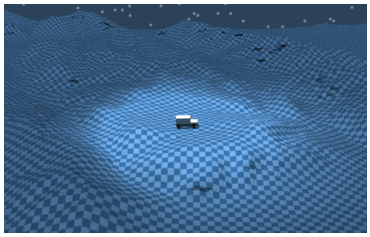

Figure 2: Off-road simulation environment in MuJoCo.

References

[1] Xiao, Xuesu, Bo Liu, Garrett Warnell, and Peter Stone. "Motion Planning and Control for Mobile Robot Navigation Using Machine Learning: A Survey." Autonomous Robots 46, no. 5 (2022).

[2] https://github.com/ai-winter/python_motion_planning

Acknowledgement: This work has been carried out within the framework of Flanders Make's SBO project LearnOpTra (Learning Meets Optimisation for Robust and Multimodal Trajectory Planning and Control).

Picking strategy for automated apple harvesting

Marissa Jonker[†,‡,⋆], Jan de Jong[⋆], Wesley Roozing[‡], and Wouter Hakvoort[†]

[†] Control and Mechatronics , [‡] Robotics and Mechatronics , [⋆] Precision Engineering,
University of Twente, Enschede, The Netherlands
email: {l.m.jonker, j.j.dejong, w.roozing, w.b.j.hakvoort}@utwente.nl

1 Introduction

Most existing works on automated harvesting of hanging fruits such as apples, consider a picking motion either as a pull in one direction [1], or as a rotation around the apple center [2]. However, these methods require a larger force applied to the apple surface, in order to break at the abscission joint (detachment point, a in Figure 1a), compared to a rotation around this point [3]. This could cause damage to the apple and/or its pedicle. A rotation around just a is enough to cause detachment, or will result in requiring less pulling force for detachment, thus decreasing the risk of causing damage. Since an undesired detachment can also occur at the attachment point of the pedicle to the apple, p in Figure 1a, a suitable picking strategy is needed to ensure this point stays intact. Although learning from demonstration has been used to perform the rotation motion around a, the motion execution was purely kinematics based, and the detachment was not detected [4].

We formulate a bending-stress steered motion strategy for the harvesting of apples using a three-fingered gripper, while keeping into account possible movement of the branch.

2 Picking strategy

Suppose we take a simplified 2D case of one apple hanging unobstructed from a branch, as shown in Figure 1. The 'ideal' rotation of the apple around its abscission joint, goes *against* the stem's growing direction [3]. In order to prevent detachment at p, we propose to move the apple such that internal moment and thus bending stress is zero at p but high at a. This implies a net resultant force through point p, as indicated by f_p in Figure 1b.

Assume we kinematically know the motion path. When we actually want to move the apple around a, we first account for the gravitational forces f_g on the apple, and get f_p by applying end-effector force f_e. However, due to the gravity compensation of the apple, it is likely that the main branch slightly moves as well, changing the position of our axis of rotation. This might actually benefit the environment, as otherwise the detachment would cause a sudden movement of the branch, possibly damaging the other apples. Adding force sensors to our gripper allows us to apply the right forces, and by implementing impedance control, we can follow the movement of the rotation axis and perform the right rotation motion around point a.

3 Future work

The formulated strategy requires additional knowledge of the environment. We need to be able to select a suitable apple s.t. the ideal picking motion can be kinematically executed (that is, without collisions). Thus we require a collision-aware world model with e.g. 5D apple poses, pedicle growing directions, and abscission joint locations.

References

[1] Velasquez et al., *IEEE RA-L*, vol. 10, pp. 11030–11037, Oct. 2025.
[2] Zhang et al., *Mechatronics*, vol. 79, no. 102644, Nov. 2021.
[3] Li et al, *J. Comp. and Elec. in Agr.*, vol. 127, pp. 633–640, Sept. 2016.
[4] Van de Ven et al., *J. S. Agric. Techn.*, vol. 11, no. 101006, Aug. 2025.

Acknowledgment
1. Funding for this work is provided by Dutch Research Council (NWO) under the research programme P17-01 "FlexCRAFT"
2. The Synergia project is organized and led by Wageningen University and Research in close cooperation with Next Food Collective as well as the Universities of Delft, Twente, Eindhoven, and Nijmegen. The authors have declared that no competing interests exist in the writing of this publication. Funding for this research was obtained from the Netherlands Organisation for Scientific Research (NWO grant 17626), IMEC-One Planet and other private parties.

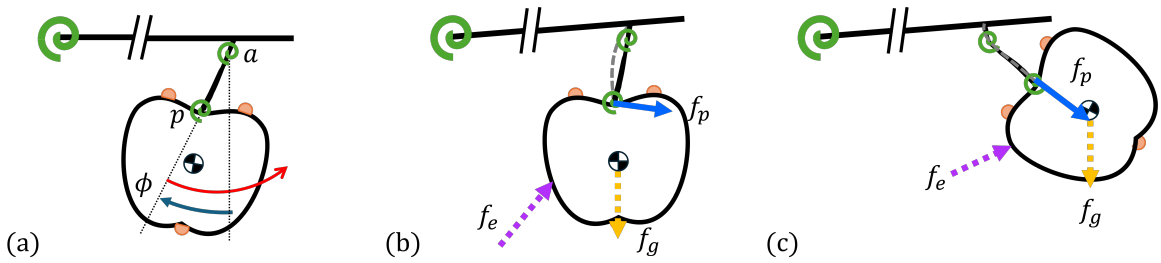

Figure 1: (a): simplified drawing of apple hanging from a branch. Detachment point a, pedicle attachment point p, growing angle ϕ. The rotation points are indicated by springs. The apple is grasped with a three-fingered gripper indicated by the orange half-circles. Ideal harvest motion is indicated by the red arrow. (b): we obtain resultant force f_p of gravitational force f_g and end-effector force f_e, resulting in a rotation around a. (c): once the apple is rotated (up to the kinematic limit), we pull the apple along its pedicle-calyx axis.

Energy Management and Scheduling for Regional Electric Aviation Networks

Finn Vehlhaber and Mauro Salazar
Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands
{f.n.vehlhaber,m.r.u.salazar}@tue.nl

The inception of electric aircraft in the near future will introduce unprecedented power and energy demands to regional airports. In preparation, airports willing to participate in regional electric aviation networks will have to install charging infrastructure, but also potentially upgrade their grid connection and consider a more integrated local energy ecosystem that includes renewables and energy buffers like stationary batteries. Research in adjacent fields has shown that such energy systems benefit from control schemes that address the intermittent nature of renewables when managing charging tasks, which will be of even greater importance in far-reaching interconnected aviation-energy networks.

In this work we propose an energy management scheme for a fleet of electric aircraft operating in a regional aviation network. Specifically, we consider local conditions at every airport in the network to inform charging decisions that reduce energy costs while guaranteeing flight operations.

Overview of the Methodology We model the energy system at every airport $h \in \mathscr{H}$ as depicted in Fig. 1. The power to charge aircraft, P_a, is supplied from renewable sources P_rnw and the grid, P_gr. The sytem also includes an energy buffer in form of a battery energy storage system (BESS), but other buffers, such as electric vehicle fleets parked in the parking lot providing vehicle-to-grid (V2G), are possible.

We optimize the charging schedule of every aircraft $k \in \mathscr{P}$ in the fleet based on (projected) availability of solar power and the flight schedule with respect to the time-varying prices of energy, $c^h(t)$, as sketched below.

Problem 1 (Flight Network Energy Management)**.**

$$
\begin{aligned}
\min_{\{P_\mathrm{c}^k(t)\}_{k\in\mathscr{P}}} \quad & \int_0^T \sum_{h\in\mathscr{H}} c^h(t) P_\mathrm{gr}^h(t)\,\mathrm{d}t \\
\text{s.t.} \quad & P_\mathrm{a}^h(t) = P_\mathrm{gr}^h(t) + P_\mathrm{rnw}^h(t) + P_\mathrm{b}^h(t) && \forall h \in \mathscr{H}, \forall t \\
& P_\mathrm{a}^h(t) = \sum_{k\in\mathscr{A}^h(t)} P_k(t) && \forall h \in \mathscr{H}, \forall t \\
& \frac{\mathrm{d}}{\mathrm{d}t} E_\mathrm{b}^\kappa(t) = -P_\mathrm{b}^\kappa(t) && \forall \kappa \in \{\mathrm{h},\mathrm{k}\}, \forall t \\
& \textit{operational constraints}, \textit{limits}, && \ldots
\end{aligned}
$$

We discretize Problem 1 with an Euler Forward scheme and solve it with slight reformulations for two scenarios; (i) charge scheduling for a day of operations for a fixed flight schedule with V2G at the hub airport, and (ii) a model predictive control (MPC) for rescheduling and charging.

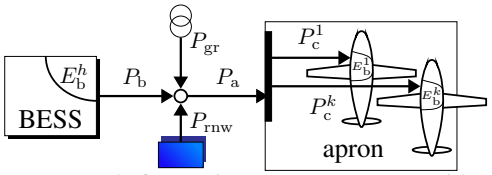

Figure 1: Example for an airport energy system with renewables, a stationary battery (BESS) and grid connection.

(i) V2G Support for Electric Aviation A case study for KLM flight schedules for inter-European shuttle flights to Amsterdam finds reductions in energy costs of up to 30 % for fleet charging when capitalizing on potential V2G capabilities of electric vehicles (EVs) parked in the airport parking lot, depending on the amount of EVs used.

(ii) Charge Scheduling MPC Case studies for a flight network in the American North East in a custom agent-based simulation environment show significant reduction in grid energy requirements when adapting the charging and flight assignment to the availability of renewable energy [1].

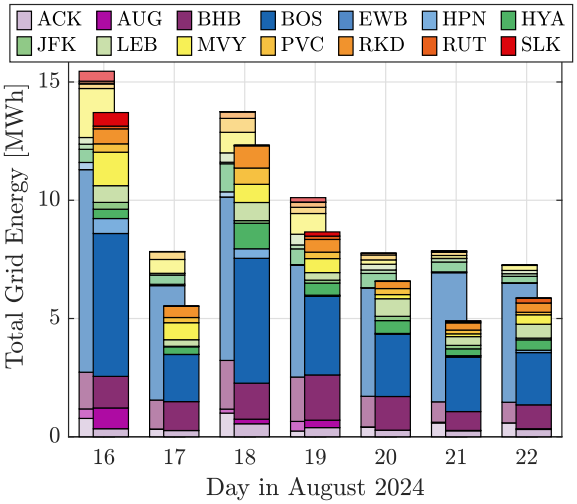

Figure 2: Comparison of energy requirements for a baseline (muted, background) and with the MPC for the daily flight schedules for a week of operations in New England at airports in the network presented by their IATA identifiers.

References

[1] F. Vehlhaber and M. Salazar, "A Model Predictive Control Scheme for Flight Scheduling and Energy Management of Electric Aviation Networks," in *European Control Conference*, 2026, under review. arXiv preprint arXiv:2404.09282v2.

Stability and Congestion Control of Information-Based Routing in Dynamic Transportation Networks

Shaya Garjani, Ashish Cherukuri, Bayu Jayawardhana, Nima Monshizadeh
ENTEG, University of Groningen, 9747 AG Groningen, The Netherlands
Email: {s.garjani, a.k.cherukuri, b.jayawardhana, n.monshizadeh}@rug.nl

1 Motivation

Recent work has highlighted that the growing use of route-recommendation technologies can adversely affect transportation networks, giving rise to congestion and oscillatory flow patterns [1]. This is due to the fact that these technologies allow users to continuously adjust their route choices to minimize individual travel times, which can lead to inefficient outcomes when many users respond to the same information. At the same time, these technologies provide a promising opportunity for non-intrusive traffic control, as network operators can influence traffic evolution by shaping the information communicated to users. A key challenge is therefore to identify information signals that lead to equilibria with desirable stability and convergence properties. This issue becomes particularly relevant in settings where traffic densities and routing decisions evolve simultaneously [1], [2], as is increasingly the case with real-time navigation and digital traffic information systems.

2 Problem Formulation

Consider p parallel path connecting a pair of origin-destination nodes. For each path $j \in \mathcal{P}$, the joint macroscopic density and routing dynamics are given by

$$\dot{x}_j(t) = -f_j(x_j(t)) + \lambda r_j(t), \tag{1a}$$

$$\dot{r}_j(t) = -r_j(t) + \frac{e^{-\eta u_j(t)}}{\sum_{i\in\mathcal{P}} e^{-\eta u_i(t)}}. \tag{1b}$$

The density dynamics in (1a) are governed by a mass conservation law, where $x_j(t)$ denotes the traffic density on path j, $f_j : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ represents the outflow from that path, and $\lambda > 0$ is the constant inflow to the network. At each time t, this inflow is split among the available paths according to the routing ratios $r_j : \mathbb{R}_{\geq 0} \to [0, 1]$. The routing ratios evolve according to the logit dynamics in (1b). The signal $u_j(t) \in \mathbb{R}_{\geq 0}$ represents traffic information communicated for path j and is assumed to depend continuously on the corresponding density, i.e., $u_j(t) \equiv u_j(x_j(t))$. The parameter $\eta \geq 0$ captures the responsiveness of the driver population to the communicated information. Each path $j \in \mathcal{P}$ has a critical density $\bar{x}_j$ and a corresponding critical flow $\bar{f}_j := f_j(\bar{x}_j)$. The interval $\mathcal{X}_j := [0, \bar{x}_j]$ represents the free-flow regime of path $j \in \mathcal{P}$, while outside this region corresponds to the congested regime. Inside the free-flow region, f_j is μ_j-strongly monotone.

Problem 1 *Specify a class of density-dependent information signals u, denoted by $\mathcal{U}$, such that for any $u \in \mathcal{U}$, (1) admits a unique asymptotically stable equilibrium $(x^{\mathrm{eq}}, r^{\mathrm{eq}}) \in \mathrm{int}(\mathcal{X}) \times \mathcal{R}$, where $\mathcal{R} := \{r \in \mathbb{R}^p_{\geq 0} \mid \mathbb{1}^\top r = 1\}$.*

3 Main Result

Theorem 1 *Let u belong to*

$$\mathcal{U} = \left\{ u : \mathcal{X} \to \mathbb{R}^p_{\geq 0} \middle| \lambda \leq \bar{f}_j e^{\eta \underline{u}_j} \sum_{i\in\mathcal{P}} e^{-\eta \bar{u}_i}, \left\| \frac{\partial u}{\partial x} \right\| < \frac{2\mu_m}{\eta\lambda} \right\},$$

where $\underline{u}_j$ and $\bar{u}_j$ are, respectively, the lower bound and upper bound of u_j, and $\mu_m := \min_{j\in\mathcal{P}} \mu_j$. Then, there exists a unique equilibrium $(x^{\mathrm{eq}}, r^{\mathrm{eq}}) \in \mathrm{int}(\mathcal{X}) \times \mathcal{R}$ that is asymptotically stable. Furthermore, the set $\mathcal{X} \times \mathcal{R}'$, where $\mathcal{R}' := \{r \in \mathcal{R} \mid r \leq \bar{f}/\lambda\}$ is positively invariant and solutions starting in this set asymptotically converge to the unique equilibrium.

The results of Theorem 1 are illustrated in Figure 1. As can be seen, all trajectories $x(t)/\bar{x}$ and $r(t)/\bar{f}$ starting in $\mathcal{X} \times \mathcal{R}'$ remain within this set and converge to the unique equilibrium.

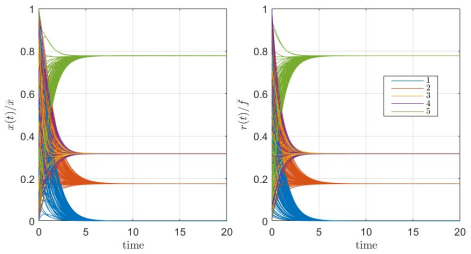

Figure 1: 100 trajectories of $x(t)/\bar{x}$ and $r(t)/\bar{f}$ starting from $x^0 \in \mathcal{X} \times \mathcal{R}'$ under a designed $u \in \mathcal{U}$.

References

[1] G. Bianchin and F. Pasqualetti, "Navigation Systems May Deteriorate Stability in Traffic Networks," IEEE Open Journal of Control Systems, 2024.

[2] G. Como and R. Maggistro, "Distributed Dynamic Pricing of Multiscale Transportation Networks," IEEE Transactions on Automatic Control, 2022

Real-time Human-in-the-loop Energy Management Policy for Electric Endurance Racing Cars

Erik van den Eshof, Jorn van Kampen, Mauro Salazar
MOVEMENT Research Group, Control Systems Technology Section
Eindhoven University of Technology
{r.c.p.v.d.eshof,j.h.e.v.kampen,m.r.u.salazar}@tue.nl

1 Abstract

In recent years, there is a shift towards battery electric vehicles (BEVs) within the automotive industry. This trend is observed in both road and racing vehicles, whereby the latter acts as an experimental platform for the former. However, both types of vehicles are still facing challenges related to driving range and charging time. Specifically within the context of electric racing, making optimal use of the battery energy is of paramount importance. To this end, optimization methods are required to compute the optimal powertrain operation (PO), whilst adhering to battery energy constraints [1]. Yet such an offline race strategy cannot be applied directly during a race, since we have to deal with many disturbances, such as competitor interactions and model errors. Furthermore, tracking energy-limited PO trajectories is challenging for human drivers, as these profiles are typically dynamic [1] and following them increases cognitive load. This calls for algorithms that adapt the optimal energy management strategy such that it is safe in operation alongside a human driver, provide disturbance rejection, and are computationally efficient enough to operate in real-time.

In previous work, we proposed a robust real-time capable energy management framework that can be efficiently tracked by a human driver [2]. Hereby, we used previously developed convex models [1] and adapted the solution to a trajectory consisting solely of full-throttle, coasting and braking sections. We used the kinetic co-state to calculate the coasting points whenever a threshold is exceeded, whereby this threshold value is found through a bisection algorithm, and could be adapted online for disturbance rejection.

Our current work further refines these methods. Instead of adapting an optimized trajectory, we choose a model structure that is guaranteed to result in an optimal solution that a driver can safely follow. Moreover, we analytically reconstruct the optimal trajectory using the optimizer co-states. This eliminates the need for a bi-section algorithm and means that in addition to the coasting phase, we identify a regenerative braking phase. Fig. 1 shows how the driver is informed to coast and regenerate as the kinetic to battery co-state ratio passes thresholds based on propulsive and regenerative powertrain efficiencies.

Overall, this greatly simplifies online implementation and the addition of a regenerative braking phase improves lap-time potential. Fig. 2 shows the significant lap-time improvements of our active, co-state-based energy management strategy (EMS) over a baseline strategy that lowers the maximum propulsive power as an energy-saving measure.

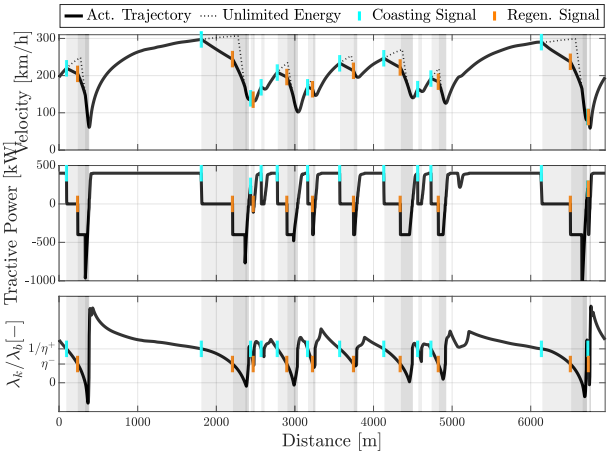

Figure 1: When the kinetic to battery co-state ratio passes the thresholds, a signal is given to the driver to coast or regenerate. This co-state ratio trajectory is loaded on the car as a look-up-table, enabling real-time use.

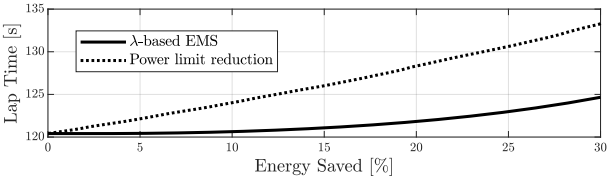

Figure 2: Demonstrating the significance of co-state-based energy management over a more straightforward method that saves energy by lowering the car's peak power.

References

[1] J. van Kampen, T. Herrmann, T. Hofman, and M. Salazar, "Optimal endurance race strategies for a fully electric race car under thermal constraints," *IEEE Trans. on Control Systems Tech.* , 2024.

[2] E. van den Eshof, J. van Kampen, and M. Salazar, "Human-in-the-loop energy and thermal management for electric racing cars through optimization-based control," *European Journal of Control*, p. 101293, 2025.

A Competitor-aware Race Management Policy: a Game Theoretical Approach

Wytze de Vries, Erik van den Eshof, Jorn van Kampen, Mauro Salazar
MOVEMENT Research Group, Control Systems Technology Section
Eindhoven University of Technology
w.a.b.d.vries@student.tue.nl,
{r.c.p.v.d.eshof,j.h.e.v.kampen,m.r.u.salazar}@tue.nl

In fully electric endurance racing, limited battery capacity and aerodynamic interactions make race management a complex strategic challenge. The presence of multiple competitors turns the problem into a multi-agent setting, where each vehicle's actions directly influence the energy usage, feasible strategies, and finishing position of others. Traditional single-agent approaches fail to capture these competitive dynamics.

In this work, we present a bi-level framework, shown in Fig. 1, for competitor-aware race management that integrates game-theoretic optimization with reinforcement learning. At the lower level, a multi-agent optimal control problem is formulated to simultaneously optimize powertrain operation and spatial trajectories for multiple agents over a single lap. We explicitly model aerodynamic interactions and incorporating asymmetrical collision-avoidance constraints inspired by real-world motorsport rules. This ensures that the following vehicle is responsible for avoiding contact while the leading vehicle can maintain its preferred trajectory, improving the realism of the resulting trajectories as shown in Fig. 2.

To enable computationally tractable long-horizon optimization, a neural-network-based surrogate model of the single-lap problem is developed. Symmetry and monotonicity constraints are enforced within the surrogate to guarantee adherence to fundamental game-theoretic and optimization principles, providing a reliable environment for higher-level decision-making. At the upper level, reinforcement learning is used to train agents to make robust per-lap decisions on battery energy allocation, pit-stop timing and charging times, with the objective of maximizing finishing position over the full race horizon.

The framework is demonstrated in a two-agent scenario over 45 laps at the Zandvoort circuit. The results show that exploiting aerodynamic interactions is decisive for race outcomes, revealing that the agent starting in the lead is ultimately disadvantaged due to the strategic utilization of slipstreaming and overtaking. These results highlight how finishing-position-focused strategies differ fundamentally from minimum-time strategies and emphasize the importance of integrating competitive dynamics into race management.

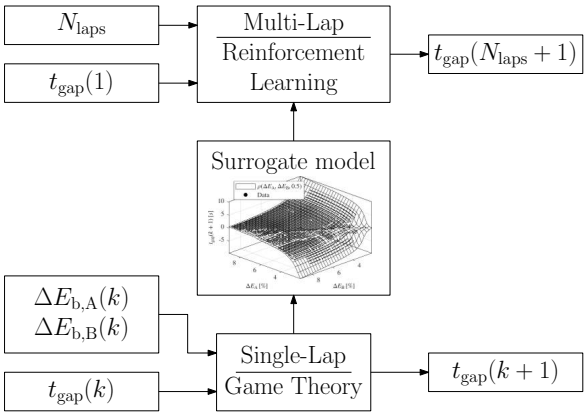

Figure 1: Proposed bi-level game-theoretic framework for competitor aware race management. The lower level solves the single-level multi-agent problem for a given initial time gap $t_{\text{gap}}(k)$ and battery energy allocation ΔE_{b}, resulting in a time gap at the end of the lap of $t_{\text{gap}}(k+1)$. The higher level trains a policy to solve the multi-lap multi-agent problem using reinforcement learning for a given race length of N_{laps} and a time gap at the start of the first lap $t_{\text{gap}}(1)$. The output the result of the race in the form of the terminal time gap $t_{\text{gap}}(N_{\text{laps}}+1)$.

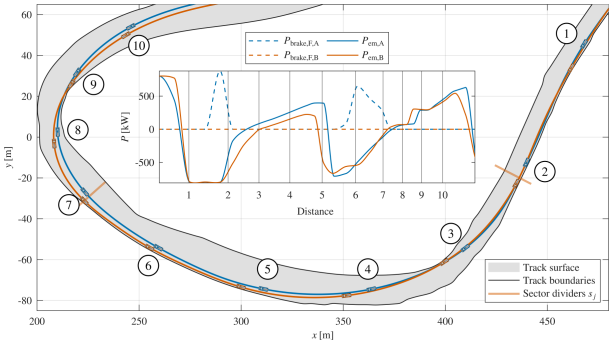

Figure 2: Driven racing lines, electric motor power, and front braking power during the overtaking maneuver through the combination of Turns 9 and 10. Agent A brakes more aggressively into Turn 9 (2), resulting in a reduced entry velocity that enables earlier acceleration (3) and a tighter exit line (5). By spending additional energy, agent A moves ahead of its opponent at the next sector divider (7), thereby gaining the right to select the preferred racing line and complete the overtake (8).

On the Engineering Trap and Responsible Innovation: Challenges and Opportunities for Systems and Control

Mauro Salazar and Alessandro Piazza
MOVEMENT Research Group, Control Systems Technology Section
Eindhoven University of Technology
{m.r.u.salazar,a.p.piazza1}@tue.nl

Our societies face pressing challenges across multiple domains: congestion in critical infrastructure, the sustainable procurement and distribution of energy, pollution and traffic-related environmental impacts, and persistent social inequalities, including barriers to accessibility. At the same time, advances in technologies such as the Internet of Things, AI, Big Data, automation, smart grids, and renewable energy systems create genuine opportunities to tackle these challenges. Addressing these challenges requires more than technical fixes; it demands a careful consideration of what society truly needs and how technology can serve those needs: Without critical reflection, such innovations risk reinforcing what we call the "engineering trap": on the one hand, an uncritical techno-optimism that treats technology deployment as a societal solution in itself; on the other hand, a narrow focus on optimal answers without questioning whether we are addressing the right societal questions. Ride-hailing platforms like Uber and Lyft offer a cautionary example: initially introduced to reduce congestion, pollution, and inequality in urban mobility, they have often worsened all three. In line with Langdon Winner's argument that (technological) "artifacts have politics", now more than ever (control) engineers have the double responsibility of not only devising effective technological solutions but also ensuring that they are catering to the needs of society and the environment in terms of wellbeing, justice and sustainability in the first place.

This talk will focus on how the field of systems and control can address these challenges. First, I will review the current status quo and the "engineering trap," reflecting on our societal and environmental responsibilities as engineers and the limits of purely technical solutions. Next, I will discuss how transdisciplinary efforts can offer a pathway to responsible innovation. Drawing on philosophical theories of justice and their application to domains such as energy, water, and transport, I will show how these principles can be operationalized in engineering through a transdisciplinary nexus that leverages optimization as a systematic advisory tool. Finally, I will present recent results on optimization models for designing and operating mobility systems according to the capabilities approach. I will demonstrate how future intermodal mobility systems can be managed to provide sufficient accessibility to each citizen, rather than optimizing solely for efficiency (Fig. 1), and how cycling infrastructure can be upgraded to ensure a sufficient level of safety for all users (Fig. 2). These examples will show that improving justice and accessibility is not only compatible with, but can actually co-benefit, traditional efficiency metrics.

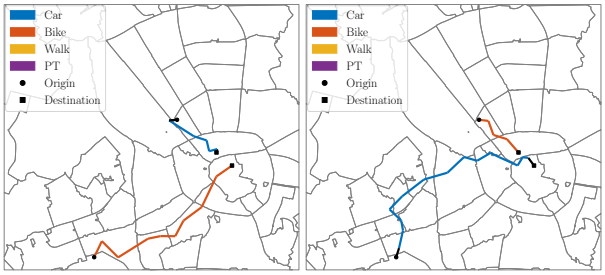

Figure 1: Optimal routes in line with conventional utilitarian approaches on travel time minimization (left), and sufficientarian accessibility paradigms (right) [1].

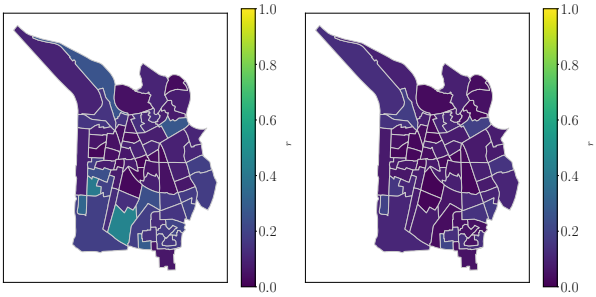

Figure 2: Safety deficit levels for Grenoble, FR, when the cycling infrastructure is designed with utilitarian (left) and sufficientarian metrics (right) [2].

References

[1] M. Salazar, S. Betancur Giraldo, F. Paparella, L. Pedroso, and K. Martens, "Mobilizing transport justice: A sufficientarian optimization framework for intermodal mobility systems," *NPJ Sustainable Mobility and Transport*, vol. 2, no. 43, 2025.

[2] M. Campero Jurado, C. Canudas de Wit, G. De Nunzio, and M. Salazar, "On transport justice and safety in bicycle network design optimization," in *Proc. IEEE Conf. on Decision and Control*, 2025, in Press.

Linear Time-Delay Control by Shaping Delay Distribution

Jasper Visterin and Wim Michiels
KU Leuven,
Department of Computer Science, NUMA Research Unit, 3001 Leuven, Belgium,
jasper.visterin@kuleuven.be, wim.michiels@kuleuven.be.

1 Introduction

Time delays appear in many engineering, economic, and biological systems, as the transfer of material, energy, or information is rarely instantaneous [1]. They arise through communication or computational latencies, transport phenomena, hereditary effects, or modeling simplifications [2]. Because feedback is based on past information, delays can complicate control and may degrade performance or even destabilize a system. Nevertheless, several studies also show that delays can have a stabilizing influence under suitable conditions [3]. We consider Time-Delay Systems modeled by functional differential equations. A basic example is a delay-differential equation with a single discrete delay,

$$\dot{x}(t) = F(x(t), x(t-\tau)), \quad F : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n. \tag{1}$$

More general models include distributed delays (2), which capture weighted effects of past states and often describe biological processes more accurately [4].

$$\dot{x}(t) = \int_{-\tau}^{0} g(\theta) x(t+\theta) d\theta, \quad g : [-\tau, 0] \to \mathbb{R}^n \times \mathbb{R}^n. \tag{2}$$

2 Controller Design

The stability framework in [3] extends to systems of the form

$$\begin{aligned} \dot{x}(t) &= \sum_{k=0}^{m} A_k x(t-\tau_k) + \sum_{k=0}^{m} B_k u(t-\tau_k) + \int_{-\tau_I}^{0} G(\theta) x(t+\theta) d\theta, \\ y(t) &= \sum_{k=0}^{m} C_k x(t-\tau_k). \end{aligned} \tag{3}$$

with $A_k \in \mathbb{R}^{n\times n}$, $B_k \in \mathbb{R}^{n\times n_u}$, $C_k \in \mathbb{R}^{n_y \times n}$, $G : [-\tau, 0] \to \mathbb{R}^{n\times n}$ and the delays satisfying $0 = \tau_0 \leq \cdots \leq \tau_m < \infty$.

Static output feedback of the form $u(t) = Ky(t)$ can be synthesized, or a dynamic controller of order n_c,

$$\begin{cases} \dot{x}_c(t) = A_c x_c(t) + B_c y(t), \\ u(t) = C_c x_c(t) + D_c y(t), \end{cases} \tag{4}$$

with real matrices of appropriate dimensions. The design minimizes the spectral abscissa, whose negativity is necessary and sufficient for stability. Since the spectral abscissa is typically non-convex and non-smooth, specialized optimization and multiple initializations may be required to obtain satisfactory solutions. The same framework can be extended to address robustness and performance objectives expressed in terms of $\mathscr{H}_2$ and $\mathscr{H}_\infty$ norms.

3 Shaping the Delay Distribution

The framework further enables controllers using distributed delays, which is natural since the state of a Time-Delay System is a function segment rather than a finite-dimensional vector. Distributed delays also arise in predictor-based control. A controller combining static feedback and a distributed delay is given by

$$u(t) = K_1 y(t) + K_2 \int_{-\tau_I^C}^{0} g(\theta) y(t+\theta) d\theta, \tag{5}$$

where g is parameterized by Chebyshev polynomials,

$$g(\theta) = \sum_{k=0}^{K} c_k T_k\left(2\frac{\theta}{\tau_I^C} + 1\right), \tag{6}$$

with optimization variables $K_1, K_2 \in \mathbb{R}^{n_y \times n_u}$ and $c_k \in \mathbb{R}$. Here $\tau_I^C > 0$ denotes the integration horizon. This formulation provides a flexible way to shape the delay kernel and incorporate distributed memory into the control law.

References

[1] T. Erneux, *Applied delay differential equations*, 1 2009. [Online]. Available: https://doi.org/10.1007/978-0-387-74372-1

[2] W. Michiels, *Theory of Control of Linear Time-Delay Systems*, 1 2024. [Online]. Available: https://doi.org/10.1016/b978-0-443-14081-5.00032-5

[3] P. Appeltans, H. Silm, and W. Michiels, "TDS-CONTROL: a MATLAB package for the analysis and controller-design of time-delay systems*," *IFAC-PapersOnLine*, vol. 55, no. 16, pp. 272–277, 1 2022. [Online]. Available: https://doi.org/10.1016/j.ifacol.2022.09.036

[4] R. M. May, "Time-Delay Versus Stability in Population Models with Two and Three Trophic Levels," *Ecology*, vol. 54, no. 2, pp. 315–325, 3 1973. [Online]. Available: https://doi.org/10.2307/1934339

System-Level Disturbance Control

G. Van Meerbeeck [1,*], M. van der Hulst[1], N. Dirkx[2], R.A. González[1], K. Tiels[1],
J. van de Wijdeven[2], T. Oomen[1]
[1]Dept. of Mechanical Engineering, Eindhoven University of Technology, The Netherlands
[2]ASML, Veldhoven, The Netherlands
[*]Email: g.v.meerbeeck@student.tue.nl

1 Background

High-performance mechatronic systems commonly consist of multiple subsystems that must jointly achieve a goal, yet are typically controlled in a decentralized manner that addresses disturbances at a submodule level.

2 Approach

The main idea is to develop a control design framework in which a feedthrough (FT) add-on addresses disturbances. The add-on structure augments the existing decentralized controllers with interconnections (feedthroughs) between subsystems, specifically aimed at suppressing disturbances at the system level.

3 Framework

Consider the closed-loop control scheme in Figure 1 that generalizes the framework in [1]. Herein, P is a system-of-systems, V is the disturbance model, and W is the weighting matrix that shapes the system-level error e. The FT controller $\bar{C}$ aimed at reducing e is expressed as

$$\bar{C} = \left(C - \hat{P}^{-1}ZH\right)(I + ZH)^{-1}, \tag{1}$$

where C is the initial decentralized controller, $\hat{P}$ is the plant model, Z is a filter matrix for causal and stable inversion of $\hat{P}$, and the FT matrix H is the free design parameter for performance. The resulting system-level error is given by

$$e = W(I + ZH)SVd, \tag{2}$$

where $S = (I + PC)^{-1}$ is the original sensitivity function, and d is the exogenous disturbance vector.
The problem considered is to minimize the 2-norm of (2) through design of the FT matrix H, given the a priori knowledge of the disturbances embedded in the model V.

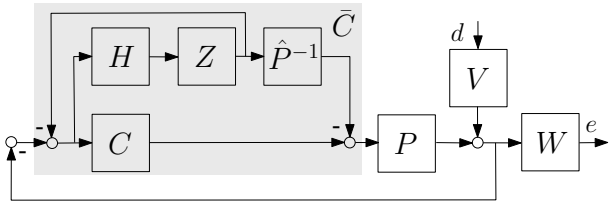

Figure 1: Control structure.

4 Suppressing disturbances

The system-level performance of FT-controlled systems is characterized by their ability to obtain a zero error in (2) at zero frequency and by the well-posedness of the controller in (1). Existence conditions on these characteristics are formulated under two classes of disturbance models: *full*-rank and *reduced*-rank models, which correspond to *uncorrelated* and *correlated* disturbance structures, respectively. The key distinction is that reduced-rank models induce a left-nullspace in V which can be exploited in the design of H to minimize (2).
The existence conditions reveal that uncorrelated disturbance structures impose fundamental limits on achievable performance, whereas incorporating knowledge of correlated disturbances in the FT design overcomes these limitations. Since conventional FT approaches assume uncorrelated disturbances, the explicit exploitation of disturbance correlation enables significant improvement in system-level disturbance suppression.

5 Results

Figure 2 illustrates the simulated system-level error of two rigid-body subsystems perturbed by one shared disturbance. Incorporating the corresponding reduced-rank disturbance model into the FT design reduces system-level errors in the low-frequency range compared to conventional FT control and decentralized control.

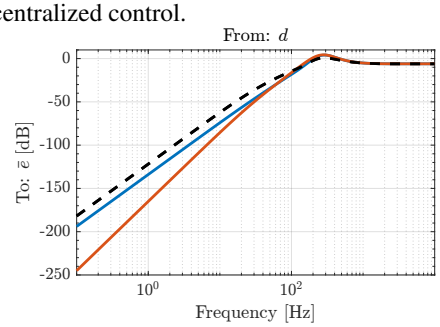

Figure 2: Bode magnitude plot of the transfer function from d to e in (2) for disturbance-based FT control (—), conventional FT control (—), and decentralized control (- -).

References

[1] E. Evers, M. v.d. Wal, and T. Oomen, "Beyond decentralized wafer/reticle stage control design: A double-Youla approach for enhancing synchronized motion," *Control Engineering Practice*, vol. 83, pp. 21–32, 2019.

A Density-Based Approach to the Dual Koopman Observer

Judicaël Mohet[a], Alexandre Mauroy[a] and Joseph J. Winkin[a]
a: Namur Institute for Complex Systems (naXys) and Department of Mathematics, University of Namur, Rue de Bruxelles 61, B-5000 Namur, Belgium
Email: judicael.mohet@unamur.be, alexandre.mauroy@unamur.be, and joseph.winkin@unamur.be

1 Introduction

In [1], we introduce the theory of the dual Koopman observer to estimate the state of a nonlinear finite-dimensional system through an infinite-dimensional linear observer. This framework is based on the (dual) Koopman operator, i.e., the composition operator associated with the flow of the nonlinear dynamics.

However, the analysis in [1] assumes that the initial condition of the observer is an evaluation functional at a guessed initial state of the nonlinear system, which may be hard to select. In the present work, we extend the Koopman observer to accommodate initial conditions associated with bounded density functions representing a probability distribution of the true initial state. This approach offers greater flexibility in selecting the observer's initial condition, as it distributes the initial uncertainty over a range of possible states rather than relying on a single point estimate.

2 Dual Koopman System

Consider the nonlinear system on $\mathbb{R}^n$

$$\begin{cases} \dot{x}(t) = F(x(t)),\ x(0) = x_0 \\ y(t) = h(x(t)), \end{cases} \tag{1}$$

with its solution $x(t) = \varphi^t(x_0)$, where $(\varphi^t)_{t\geqslant 0}$ is a flow (semigroup). It is well known that the Koopman operator, i.e. the composition operator with the flow φ^t, provides a linear representation of the nonlinear system (1), but in an infinite dimensional setting, see e.g. [2].

Let $\mathbb{H}$ be a reproducing kernel Hilbert space on a compact set $S \subseteq \mathbb{C}^n$, i.e. an Hilbert space on which each evaluation functional is continuous, and let us denote its inner product by $\langle \cdot, \cdot \rangle$. The dual Koopman system of (1) on $\mathbb{H}$ is defined by

$$\begin{cases} \dot{f}(t) = A_F^* f(t),\ f(0) = k_{x_0} \\ y(t) = C_h f(t), \end{cases} \tag{2}$$

where $f(t) \in \mathbb{H}$, for all $t \geqslant 0$, A_F^* is the adjoint operator of the Liouville operator $A_F g = \nabla g \cdot F$, k_{x_0} denotes the evaluation functional at $x_0 \in \mathbb{S}$, i.e. $\langle f, k_{x_0} \rangle = f(x_0), \forall f \in \mathbb{H}$, and C_h is defined by

$$C_h g = \begin{pmatrix} \langle h_1, g \rangle & \cdots & \langle h_m, g \rangle \end{pmatrix}',$$

with $h = (h_1, \cdots, h_m) \in \mathbb{H}^m$. The solution $f(t)$ of this system, when it exists, satisfies the identity $\langle g, f(t) \rangle = g(\varphi^t(x_0))$, for all $g \in \mathbb{H}$, i.e. $f(t) = k_{\varphi^t(x_0)}$ in the week sense.

3 Density approach

For a given class of dynamics, we proved in [1] that the (infinite dimensional) Luenberger observer

$$\dot{\hat{f}}(t) = A_F^* \hat{f}(t) + L\left(C_h \hat{f}(t) - y(t)\right),\ \hat{f}(0) = k_{\hat{x}_0} \tag{3}$$

where L is a linear bounded operator stabilizing the error dynamics, provides an estimate $\hat{x}(t)$ that converges weakly and exponentially fast to $x(t)$.

Now, we define the bounded linear operator $T\colon L^2(S) \to \mathbb{H}$ such that, for all $g \in L^2(S)$ and all $x \in S$,

$$Tg(x) = \int_S k(x,y) g(y) dy, \tag{4}$$

where $k\colon S^2 \to \mathbb{C}$ is the kernel of $\mathbb{H}$, see for instance [3, Theorem 4.26]. Take a bounded density $\rho_0 \in L^2(S)$. For a given class of dynamics, we prove that the observer (3) with $\hat{f}(0) = T\rho_0 \in \mathbb{H}$ provide an estimate $\hat{x}(t)$ that converges weakly and exponentially fast to $x(t)$. In particular, it appears that this estimate is the expectation of the estimate obtained with an evaluation functional associated with an initial guess randomly distributed according to the law ρ_0. More specifically, if the solution to the dual Koopman observer is denoted by $\hat{K}(t)T\rho_0$, then we have, for all $g \in \mathbb{H}$,

$$\langle g, \hat{K}(t) T \rho_0 \rangle = \langle g, \int_S \rho_0(x) \hat{K}(t) k_x dx \rangle.$$

This approach is illustrated by numerical simulations.

References

[1] Mohet, J. and Mauroy A. and Winkin J.J., A Dual Koopman Approach to Observer Design for Nonlinear Systems (arXiv:2503.08345).

[2] Mauroy, A. and Mezić I. and Susuki Y., The Koopman Operator in Systems and Control : Concepts, Methodologies, and Applications. 2020

[3] Steinwart, I. and Christmann, A., Support Vector Machines, Information Science and Statistics Series, Springer, New York, 2008.

Nonlinear Bandwidth and Bode Diagrams based on Scaled Relative Graphs

Julius P. J. Krebbekx[1], Roland Tóth[1,2], Amritam Das[1]
[1]Control Systems group, Department of Electrical Engineering, Eindhoven University of Technology Eindhoven, The Netherlands.
[2]Systems and Control Laboratory, HUN-REN Institute for Computer Science and Control, Budapest, Hungary.
E-mail: {j.p.j.krebbekx, r.toth, am.das}@tue.nl

1 Introduction

In the case of *Linear Time-Invariant* (LTI) systems, graphical analysis using the Nyquist diagram and the Bode diagram forms the cornerstone of control engineering. *Scaled relative graphs* (SRGs), introduced in [1] for system analysis, provide a novel graphical frequency-domain method for the analysis of *nonlinear* (NL) systems.

Inspired by the success of analyzing performance through Bode diagrams in the LTI case, there have been multiple attempts to generalize the Bode diagram to the NL case: the describing function, sinusoidal input L_2-gain, and even input-specific gain, see [2] for an overview. All these methods are either approximative, consider only sinusoids, or require input-specific computations.

In our work [3], we go beyond sinusoidal signal analysis and develop a method to compute frequency-dependent gain bounds not only for sinusoidal signals, but also periodic signals and periodic signals with subharmonic content, for which SRGs are the main computational tool. Plotting these bounds as function of frequency yields a *nonlinear Bode diagram*, based upon which we can define the *nonlinear bandwidth* analogously to the LTI case. We also illustrate how these nonlinear Bode diagrams can correctly predict the gain at different frequencies, and nonlinear effects such as a nonzero DC gain due to a nonlinearity [3].

2 Main Results

2.1 Frequency-Dependent Gains

Consider a SISO NL system $y = Ru$, where R is an operator and the input u comes from one of the following spaces.

$$\mathscr{U}_\omega := \{u \in L_{2e} \mid u(t) = a\sin(\omega t + \phi), a, \phi \in \mathbb{R}\}, \quad (1a)$$

$$\overline{\mathscr{U}}_\omega := \{u \in L_{2e} \mid u(t) = \textstyle\sum_{n\in\mathbb{Z}} \hat{u}_n e^{j\omega n t}\}, \quad (1b)$$

$$\underline{\mathscr{U}}_\omega := \{u \in L_{2e} \mid u(t) = \textstyle\sum_{0\neq n\in\mathbb{Z}} \hat{u}_n e^{j(\omega/n)t}\}, \quad (1c)$$

with $L_{2e} = \{u : [0,\infty) \to \mathbb{R} \mid \int_0^T |u(t)|^2 \mathrm{d}t < \infty, \forall T \geq 0\}$. Note that $\mathscr{U}_\omega$ contains all sinusoids of frequency ω, so it can be used to compare with approaches reported in [2]. Additionally, $\overline{\mathscr{U}}_\omega$ contains all periodic inputs with frequency ω and $\underline{\mathscr{U}}_\omega$ contains all signals with only subharmonics of ω.

We write R as an interconnection of LTI and static NL systems, a *Linear Fractional Representation* (LFR), which reads $R = P^R_{\mathrm{ew}}(\phi^{-1} - P^R_{\mathrm{zw}})^{-1}P^R_{\mathrm{zr}} + P^R_{\mathrm{er}}$, where $P^R_{\mathrm{ew}}, P^R_{\mathrm{zw}}, P^R_{\mathrm{zr}}$ and P^R_{er} are stable. Using SRG calculus from [1] on the LFR, and the SRG of LTI systems restricted to (1) from [3], we can compute the frequency-dependent gain bounds

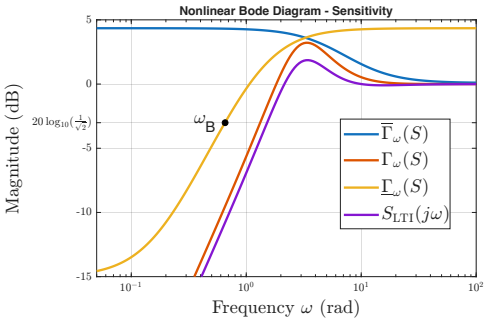

Figure 1: Nonlinear Bode diagram for $S = (1+L)^{-1}$.

$$\Gamma_\omega(R) = \sup_{u_1,u_2\in\mathscr{U}_\omega} \frac{\|\tilde{y}_1 - \tilde{y}_2\|_{\mathrm{RMS}}}{\|u_1 - u_2\|_{\mathrm{RMS}}}, \quad (2a)$$

$$\overline{\Gamma}_\omega(R) = \sup_{u_1,u_2\in\overline{\mathscr{U}}_\omega} \frac{\|\tilde{y}_1 - \tilde{y}_2\|_{\mathrm{RMS}}}{\|u_1 - u_2\|_{\mathrm{RMS}}}, \quad (2b)$$

$$\underline{\Gamma}_\omega(R) = \sup_{u_1,u_2\in\underline{\mathscr{U}}_\omega} \frac{\|\tilde{y}_1 - \tilde{y}_2\|_{\mathrm{RMS}}}{\|u_1 - u_2\|_{\mathrm{RMS}}}, \quad (2c)$$

where $\tilde{y}_i$ is the asymptotically periodic response to u_i.

2.2 Bode Diagram and Bandwidth

Next, we consider the simple feedback system $y = Le, e = r - y$ with loop-transfer L and sensitivity $S = (1+L)^{-1}$. The **closed-loop bandwidth** is the *smallest* value ω_B such that $\underline{\Gamma}_{\omega_B}(S)$ crosses $1/\sqrt{2}$ from below. The **open-loop bandwidth** is the *smallest* value ω_c such that $\overline{\Gamma}_{\omega_c}(L)$ crosses 1 from above. An example NL Bode diagram is plotted in Fig. 1, where it is clear that $\overline{\Gamma}$ or $\underline{\Gamma}$ is used to probe the high- or low-frequency behavior. Also, all gains in (2) are upper-bounded by the incremental L_2-gain. More relations between (2), L_2-gain and bandwidth are discussed in [3].

References

[1] Thomas Chaffey, Fulvio Forni, and Rodolphe Sepulchre. Graphical Nonlinear System Analysis. *IEEE Transactions on Automatic Control*, 68(10):6067–6081, October 2023.

[2] David Rijlaarsdam, Pieter Nuij, Johan Schoukens, and Maarten Steinbuch. A comparative overview of frequency domain methods for nonlinear systems. *Mechatronics*, 42:11–24, April 2017.

[3] Julius P. J. Krebbekx, Roland Tóth, and Amritam Das. Nonlinear Bandwidth and Bode Diagrams based on Scaled Relative Graphs. *arXiv:2504.01585*, April 2025. Accepted and presented at CDC 2025.

Distributed Analysis of Cone-Preserving Variational Dynamics

E. Vacchini[†‡] e.vacchini@rug.nl
Y. Kawano[§] ykawano@hiroshima-u.ac.jp
M. Cucuzzella[†] m.cucuzzella@rug.nl
A. Ferrara[‡] a.ferrara@unipv.it
A.J. van der Schaft[‡] a.j.van.der.schaft@rug.nl

† Jan C. Willems Center for Systems and Control, University of Groningen, the Netherlands
‡ Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Italy
§ Graduate School of Advanced Science and Engineering, Hiroshima University, Japan

1 Introduction

Cone invariance of a linear system enables the derivation of stability conditions through vector-type Lyapunov functions. When applied to the stability analysis of networks composed of cone-invariant systems, this property yields significantly simplified stability conditions. In particular, the stability of positive networks reduces to verifying exponential dissipativity of each subsystem with respect to vector-type storage functions and supply rates, along with a condition on the associated weighted out-degree. This approach is made possible by Perron–Frobenius theory, which guarantees a well-defined ordering and spectral structure for systems that leave cones invariant. Vector-type Lyapunov analysis has been extended to nonlinear systems within two main frameworks: classical Lyapunov analysis and contraction analysis. The latter has proven especially effective for networked systems, establishing incremental stability results via Finsler-Lyapunov functions. Moreover, contraction analysis is particularly well suited for order-preserving and cone-invariant systems. Owing to its differential-geometric nature, contraction theory relies on the analysis of variational systems, a perspective shared with many classical results on monotone systems (e.g., the Kamke conditions). We propose a unified framework demonstrating how networks of systems satisfying generalized forms of the Perron–Frobenius theorem, collectively referred to as *differentially positive systems*[1], admit distributed stability conditions, thereby consolidating and extending existing results across positive, cooperative, monotone, and cone-invariant dynamics.

2 Problem formulation

Similarly to how a vector field X defines for each point p of a manifold $\mathscr{M}$ an element X_p in the tangent space $T_x\mathscr{M}$, a cone field $\mathscr{X}$ defines a cone $\mathscr{X}_p \subset T_x\mathscr{M}$ for each $p \in \mathscr{M}$. Consider the prolongation of the system Σ

$$\Sigma : \begin{cases} \dot{x} = f(x,u) \\ y = h(x) \end{cases}, \quad \delta\Sigma : \begin{cases} \delta\dot{x} = \frac{\partial f(x,u)}{\partial x}\delta x + \frac{\partial f(x,u)}{\partial u}\delta u \\ \delta y = \frac{\partial h(x)}{\partial x}\delta x \end{cases}, \tag{1}$$

with $(x,\delta x) \in T\mathscr{M}^x$, $(u,\delta u) \in T\mathscr{M}^u$ and $(y,\delta y) \in T\mathscr{M}^y$ being the state, input and output of $(\Sigma,\delta\Sigma)$, respectively, and define the triplet of cone fields $(\mathscr{X},\mathscr{U},\mathscr{Y})$ over the state, input and output manifolds, respectively. Σ is said to be differentially positive [1] with respect to $(\mathscr{X},\mathscr{U},\mathscr{Y})$ if

$$\delta x_0 \in \mathscr{X}_{x_0},\ \delta u(t) \in \mathscr{U}_x,\ \forall t \in \mathbb{R}_{\geq 0} \implies \delta x(t) \in \mathscr{X}_x,\ \forall t \in \mathbb{R}_{\geq 0}.$$

In this work, we consider networks Σ of $N \in \mathbb{N}$ interconnected systems Σ_i of the form (1), with interconnection given by $u_i = W_i(y_1,...,y_N)$, where W_i are generic nonlinear maps. We are interested in characterizing incremental stability of Σ defined as follows.

Definition 1 [*Incremental Stability*]: Σ is incrementally stable with respect to a distance function $d : \mathscr{M}^x \times \mathscr{M}^x \to \mathbb{R}_{\geq 0}$ if there exist a class $\mathscr{K}$ map α such that, given any $x_0, x_0' \in \mathscr{M}^x$, it holds $d(x(t),x'(t)) \leq \alpha(d(x_0,x_0'))$. Moreover, Σ is incrementally asymptotically stable if $\lim_{t\to\infty} d(x(t),x'(t)) = 0$. Finally, Σ is incrementally exponentially stable if $d(x(t),x'(t)) \leq ke^{-\lambda t}d(x_0,x_0')$, for some $k,\lambda > 0$.

Problem 1: Consider a network Σ of $N \in \mathbb{N}$ sub-systems Σ_i differentially positive with respect to the cone fields $(\mathscr{X}_i,\mathscr{U}_i,\mathscr{Y}_i)$, for all $i \in \{1,...,N\}$. Derive incremental stability conditions for Σ exploiting differential positivity of of the sub-systems.

3 Main Result

We show that differential positivity plays a role in simplifying stability analysis. In particular, if the interconnections between sub-systems preserve differential positivity, i.e., $\mathscr{Y}_i$ is mapped into $\mathscr{U}_j$ whenever Σ_i and Σ_j are interconnected, stability can be certified in a distributed way. Indeed, it can be shown that Σ is incrementally stable of Σ_i is differentially dissipative and satisfies a condition on the weighted in- and out-degree, for all $i \in \{1,...,N\}$.

References

[1] F. Forni and R. Sepulchre, "Differentially positive systems," IEEE Transactions on Automatic Control, vol. 61, no. 2, pp. 346–359, 2016.

[2] D. Angeli and E. Sontag, "Monotone control systems," IEEE Transactions on Automatic Control, vol. 48, no. 10, pp. 1684–1698, 2003.

On-the-fly Surrogate Modelling with Automatic Grid-Point Allocation

E. Javier Olucha[1], Amritam Das[1] and Roland Tóth[1,2]
[1] Control Systems Group, Eindhoven University of Technology, Eindhoven, The Netherlands.
[2] Systems and Control Laboratory, HUN-REN SZTAKI, Budapest, Hungary.
{e.j.olucha.delgado, am.das, r.toth}@tue.nl

1 Introduction

In this work, we propose an automatic grid-point allocation strategy for the *on-the-fly surrogation* (OTFS) framework introduced in [1]. The OTFS leverages *automatic differentiation* to compute collections of system linearizations reliably and at low computational cost, even for complex high-fidelity models. This collection of local system information is then used to recover the global nonlinear behaviour without explicitly constructing a model. However, the practical performance of OTFS strongly depends on the set of linearization points used in the recovery process. Systematic strategies for grid-point allocation for OTFS remain a key open challenge, as uniform or manually chosen grids may lead to unnecessary computational effort or insufficient accuracy in critical regions of the operating domain. To address this issue, we propose the following automatic grid-point allocation strategy.

2 Approach

Let Σ denote a nonlinear system with state transition function f and let Π be an OTFS surrogate model of Σ described by $\hat{f}(\bullet;\Theta)$, constructed from a given set of linearization points $\Theta = \{\theta_n\}_{n=1}^{N}$. The quality of a candidate grid point θ^* is evaluated by comparing trajectories generated by the surrogate to given reference trajectories x_k generated by f. The grid-point allocation is formulated as

$$\theta^* = \arg\min_{\theta} \frac{1}{n_{\text{sim}} T_i} \sum_{i=1}^{n_{\text{sim}}} \sum_{k=0}^{T_i-1} \|x_{k+1}^i - \hat{x}_{k+1}^i\|_2, \qquad (1)$$

subject to

$$\hat{x}_{k+1}^i = \hat{f}(\hat{x}_k^i, u_k^i; \Theta \cup \{\theta\}), \qquad \hat{x}_0^i = x_0^i,$$

where n_{sim} denotes the number of simulated trajectories, n_x the state dimension, and T_i the length of the i-th trajectory. Then, *Bayesian optimization* [2] is employed to iteratively minimize (1). At each iteration, the resulting θ^* is appended to the grid Θ, yielding an adaptive allocation of linearization points concentrated in the most critical dynamic regions.

3 Results

Consider an autonomous Van der Pol oscillator described by

$$f(x) = \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 - 1.2(1-x_1^2)x_2 \end{bmatrix}, \qquad (2)$$

This work is supported by the European Space Agency (grant number: 4000145530) and The MathWorks Inc. Opinions, findings, conclusions or recommendations expressed in this abstract are those of the authors and don't necessarily reflect the views of The MathWorks Inc. or the European Space Agency.

Table 1: Average RMSE and grid cardinality

Surrogate	Grid type	Avg. RMSE	$\lvert\Theta\rvert$
$\hat{f}_1$	Uniform	0.1690	9
$\hat{f}_2$	Uniform	0.0151	25
$\hat{f}_3$	Allocated	0.0049	15

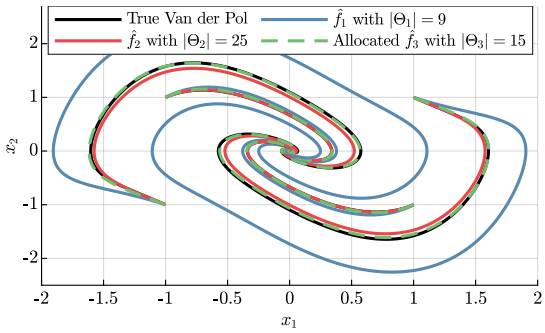

Fig. 1: State trajectories of the Van der Pol oscillator obtained from the true system and the OTFS surrogates.

and the reference trajectories obtained by simulating f using RK4 for 14 seconds with a sampling time of 0.01 seconds, starting from the initial conditions $(-1,-1)$, $(-1,1)$, $(1,-1)$, and $(1,1)$. Next, the surrogates described by $\hat{f}_1$ and $\hat{f}_2$ are obtained with the OTFS approach using the grids $\Theta_1 = \{(x_1,x_2) \mid x_1,x_2 \in \{-1,0,1\}\}$ and $\Theta_2 = \{(x_1,x_2) \mid x_1,x_2 \in \{-1,-0.5,0,0.5,1\}\}$, respectively. Now, starting from $\hat{f}_1$, six additional linearization points are allocated by iteratively minimizing (1), producing a third surrogate described by $\hat{f}_3$ with grid Θ_3. The simulations are repeated using $\hat{f}_1$, $\hat{f}_2$ and $\hat{f}_3$ instead of the true model of the system (2). The results, displayed in Table 1 and Fig. 1, show that $\hat{f}_3$, obtained with the automatic allocation, achieves a higher accuracy than the uniformly sampled surrogate $\hat{f}_2$ despite using fewer linearization points. These results illustrate the potential of the proposed grid-point allocation strategy to improve both simulation accuracy and computational efficiency of grid-based surrogation methods.

References

[1] E.J. Olucha, R. Singh, A. Das and R. Tóth, "On-the-fly Surrogation for Complex Nonlinear Dynamics," Proc. of the 64th Conf. on Decision and Control, 2025.

[2] E. Brochu, V.M. Cora, and N. Freitas, "A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning," arXiv:1012.2599, 2010.

Incremental stability of projection-based control loops

Riccardo Bertollo, Sebastiaan van den Eijnden, W.P.M.H. (Maurice) Heemels
TU Eindhoven, Dept. of Mechanical Engineering, Control Systems Technology section,
{r.bertollo, s.j.a.m.v.d.eijnden, w.p.m.h.heemels}@tue.nl

1 Abstract

Projection-based controllers can overcome fundamental limitations of linear time-invariant controllers, by modifying the controller's input-output behavior via projection. A notable example is the so-called hybrid integrator-gain system (HIGS), where the input-output pair of an integrator is constrained to belong to a sector, significantly improving its performance (the phase lag is reduced from 90 degrees to about 38) [1]. While prior work on analysis and design of projection-based control systems has primarily focused on the continuous-time setting and non-incremental analysis, in this work we analyze the incremental stability properties of discrete-time control systems involving projections, which better reflect the actual digital implementation of such controllers, and addresses their (robust) performance properties.

Our interest in incremental properties is due to the fact that incrementally stable systems are inherently robust, in the sense that small perturbations in initial conditions or inputs only lead to small changes in states and outputs. Moreover, incrementally stable and convergent systems asymptotically forget the effect of initial conditions over time, and admit a steady-state response that solely depends on the input. This opens up possibilities for accurately characterizing system performance through LTI-inspired tools such as nonlinear frequency-response functions [2] and Bode-like plots; an example of this last tool is depicted in Fig. 1, for the interconnection of a linear plant and a HIGS controller.

We focus on control loops where a projected controller is connected in feedback to a nonlinear plant, and we obtain sufficient conditions to prove incremental properties using two different approaches: *(i) quadratic incremental stability* preservation and *small-gain* approach; *(ii)* closed-loop *incremental Lyapunov function* approach.

Quadratic incremental stability preservation and small-gain approach. In this approach, we initially focus on the controller as an open-loop system, where the error e coming from the plant is treated as a general exogenous signal. We start by showing that the projection preserves the quadratic incremental stability properties (i.e., incremental stability certified by a quadratic Lyapunov function) of the unprojected controller, if the norm of the projection is selected appropriately. Building up on this result, we are able to identify sufficient conditions to prove incremental stability of the closed-loop interconnection with an incrementally stable plant, resorting to small-gain arguments. To this end, we extend the well-known small-gain results for (non-incremental) stability of interconnected discrete-time nonlinear systems [3] to our incremental setting.

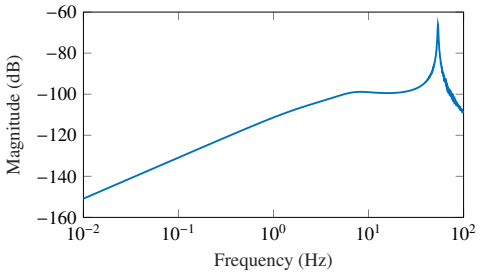

Figure 1: Nonlinear Bode-like plot for the closed-loop interconnection of a discrete-time linear plant and a HIGS.

Closed-loop incremental Lyapunov function approach. This approach relies on explicitly solving the constrained optimization problem associated with a projection. Specifically, given a set S, and a positive definite matrix Q, the projection of x on S with metric Q is given by

$$P_S^Q(x) := \arg\min_{y \in S} (y-x)^\top Q(y-x). \tag{1}$$

By explicitly solving the constrained quadratic program (1), the projected dynamical system can be described by continuous, piecewise-smooth dynamics (assuming that the unprojected dynamics is smooth). Using this characterization, we prove incremental stability of the closed loop by only checking the property for each individual mode (extending [2, Theorem 2] to the case of input-dependent switching), which, in some cases, can be done through simple LMIs.

References

[1] D.A. Deenen et al. (2021). *Projection-based integrators for improved motion control: Formalization, well-posedness and stability of hybrid integrator-gain systems*. Automatica (133).

[2] A. Pavlov, N. van de Wouw (2008). *Convergent discrete-time nonlinear systems: The case of PWA systems*. American Control Conference.

[3] Z.P. Jiang, Y. Wang (2001). *Input-to-state stability for discrete-time nonlinear systems*. Automatica (37).

Acknowledgment: Funded by the European Union (ERC Advanced Grant, Proacthis, no. 10105538).

RKHS method for computing Koopman-based Lyapunov functions

François-Grégoire BIERWART and Alexandre MAUROY
Namur Institute for Complex Systems (naXys), Department of Mathematics, UNamur, Belgium
Email : {francois-gregoire.bierwart, alexandre.mauroy}@unamur.be

1 Introduction

It is well-known that global stability of nonlinear dynamical systems is fully described by Lyapunov functions whose design is challenging. Recently, global stability analysis has been investigated in the Koopman-based framework [1] where Lyapunov functions are obtained through the eigenfunctions of the Koopman operator associated with eigenvalues of the Jacobian matrix at the origin. Usually, the eigenfunctions are unknown and only approximations can be computed, which might not correspond to the true spectrum (i.e. Jacobian eigenvalues).

In this work, we design a kernel-based method to compute the eigenfunctions of the Koopman operator while preserving the spectrum of the Jacobian matrix. The resulting Lyapunov candidate is further validated through scenario-based [2] techniques.

2 Koopman operator framework for global stability

Let $\varphi^t : \mathbb{X} \times \mathbb{R}^+ \to \mathbb{X}$ be the flow map of the dynamical system $\dot{x} = F(x)$, where F is an analytic vector field over $\mathbb{X} \subset \mathbb{R}^n$ and the origin is an hyperbolic equilibrium. Let $\mathscr{F} = C^1(\mathbb{X})$. The linear semi-group of Koopman operators $\{K_t\}_{t\geq 0}$, $K_t : \mathscr{F} \to \mathscr{F}$ is defined as $K_t f = f \circ \varphi^t$ for $f \in \mathscr{F}$. For any differentiable f, we can also define the Koopman generator

$$Lf \triangleq \lim_{t\to 0^+} \frac{K_t f - f}{t} = F \cdot \nabla f.$$

In [1], it is shown that $V = \sum_{i=1}^n \left|\phi_{\lambda_i}\right|^2$ is a Lyapunov function of the system where $\{\phi_{\lambda_i}\}_i$ are the eigenfunctions of L associated with the Jacobian eigenvalues $\lambda_i \in \sigma(\mathbf{J}_F(0))$ such that $\Re(\lambda_i) < 0$.

3 RKHS approximation of the eigenfunctions

The eigenfunctions used to compute V are usually unknown. This issue can be addressed by computing an approximation $L_N := \Pi L_{|\mathscr{F}_N} : \mathscr{F}_N \to \mathscr{F}_N$ of L where $\mathscr{F}_N \subset \mathscr{F}$ is a finite dimensional subspace and $\Pi : \mathscr{F} \to \mathscr{F}_N$ is a well-suited projection. We denote by $\mathbf{L}_N$ the matrix representation of L_N in $\mathscr{F}_N$. Here, we design an approximation of L through a Reproducing Kernel Hilbert Space (RKHS). Let H be a RKHS associated with the analytic kernel k such that its Taylor expansion does not contain monomials of degree one. We then define $\mathscr{F}_N = \text{span}\{x_i\}_{i=1}^n \oplus \text{span}\{k(q_i,\cdot)\}_{i=1}^m$ where $\{q_i\}_{i=1}^m \subset \mathbb{X}$ for some $m \in \mathbb{N}_0$ and Π as

$$(\Pi g)(x) := \nabla^\top g\,|_0\, x + \Psi(x)^\top (\mathbf{G}^{-1} v),$$

where $\mathbf{G}_{ij} = k(q_i,q_j)$, $\Psi(x) = [k(q_1,x),\dots,k(q_m,x)]^\top$, and $v_i = g(q_i) - \nabla^\top g\,|_0\, q_i$. In this case, we can show that $\mathbf{L}_N$ has the following structure

$$\mathbf{L}_N = \begin{bmatrix} \mathbf{J}_F(0)^\top & 0 \\ \mathbf{G}^{-1}\mathbf{N} & \mathbf{G}^{-1}\mathbf{A} \end{bmatrix},$$

where $\mathbf{N}_{ij} = N_j(q_i)$, $\mathbf{A}_{ij} = (Lk(q_i,\cdot))(q_j)$, and $N(x) = F(x) - \mathbf{J}_F(0)x$ is the nonlinear part of F. Since the matrix $\mathbf{L}_N$ is block-triangular (with the upper block $\mathbf{J}_F(0)$), the spectrum of the approximated generator contains the same eigenvalues as $\mathbf{J}_F(0)$.

4 Scenario-based validation

We finally derived scenario-based techniques to validate the Lyapunov candidate with probabilistic guarantees. Figure 1 shows the ROA approximation obtained through our method for the 10-dimensional system given by

$$\begin{cases} \dot{x}_{i1} &= -x_{i2}, \\ \dot{x}_{i2} &= x_{i1} - \alpha_i(1-x_{i1}^2)x_{i2} + \sum_{j\neq i}\beta_{ij}x_{i1}x_{j2}, \end{cases} \quad (1)$$

with $i = 1,\dots,5$, $\mathbb{X} = [-4,4]^{10}$. The parameters α_i and β_{ij} are randomly sampled over $[0.5, 2.5]$ and $[-0.1, 0.1]$, respectively.

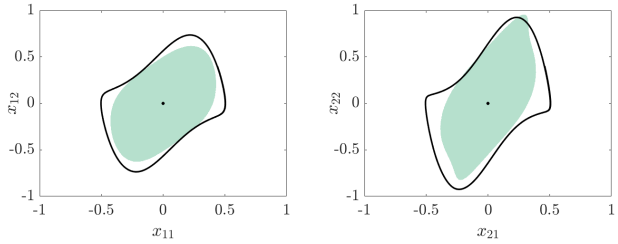

Figure 1: ROA approximation (green area) of system (1). The true ROA is delimited by the black curves.

References

[1] Mauroy, A., & Mezić, I. (2016). Global stability analysis using the eigenfunctions of the Koopman operator. IEEE Transactions on Automatic Control, 61(11), 3356-3369.

[2] M. C. Campi, S. Garatti, & F. A. Ramponi. (2015). Non-convex scenario optimization with application to system identification. In 54th IEEE Conference on Decision and Control (CDC), pages 4023–4028. IEEE, 2015.

Excitable Center–Surround Recurrent Filtering for Event-Based Object–Background Segmentation

Fatemeh Zahed Kachaee
Bât. B28, Quartier Polytech 1, Liège, Belgique
fzahed@uliege.be

Alessio Franci
Bât. B28, Quartier Polytech 1, Liège, Belgique
afranci@uliege.be

Event-based cameras are cutting-edge visual sensors that respond asynchronously to changes in light intensity rather than capturing frames at a fixed rate. They enable fast, data-efficient machine vision for robotics, industrial monitoring, and autonomous driving, offering high temporal resolution, reduced motion blur, wide dynamic range, and low power consumption [4, 3]. However, because event cameras react to any intensity change, they can generate large volumes of data. In dynamic scenes, separating moving objects from moving backgrounds or distinguishing self-motion from object motion becomes difficult when much of the event stream is irrelevant to the task.

We propose an **excitable center–surround recurrent filtering architecture** for event-based object–background segmentation. Our approach combines a retina-inspired spatial **center–surround** mechanism with nonlinear **excitable dynamics**. The center–surround interaction enhances locally distinctive motion patterns, while the excitable mechanism selectively amplifies coherent activity and suppresses weak or background-like responses.

We model the sensor as a grid of decision cells, one per pixel. Each cell maintains a fast decision state $z_i(t)$ that summarizes recent event activity. Cells interact through a fixed center–surround connectivity matrix W, constructed as a difference of Gaussians (DoG), producing local excitation within a small radius and broader inhibition over a larger surrounding region [1, 2]. This spatial coupling emphasizes local contrast in the event stream.

To introduce selectivity and adaptive behavior, we extend the recurrent center–surround model with a two-timescale excitable mechanism. The fast subsystem integrates spatial activity and event input:

$$\tau_f \dot{z}_i = -z_i + S\left((\zeta_0 - \zeta)\sum_{j=1}^{N} W_{ij} z_j\right) + \sum_{j=1}^{N} W_{ij} u_j, \quad i = 1,\ldots,N. \tag{1}$$

A slower adaptive variable $\zeta(t)$ regulates the gain:

$$\tau_s \dot{\zeta} = K_s \frac{1}{N}\sum_{i=1}^{N} \tanh(z_i^4) - \zeta. \tag{2}$$

The fast timescale τ_f enables rapid response to incoming events, while the slower timescale τ_s implements negative feedback that reduces sensitivity after strong activations. This mixed fast-positive and slow-negative feedback structure produces **excitable behavior**: locally coherent motion triggers strong transient responses, whereas background activity remains suppressed. Segmentation decisions therefore emerge as sparse, event-driven activations rather than continuous filtering. The proposed excitable center–surround architecture transforms the raw event stream into a temporally structured and spatially selective response map, improving robustness to noise and ego-motion while preserving sensitivity to salient object motion.

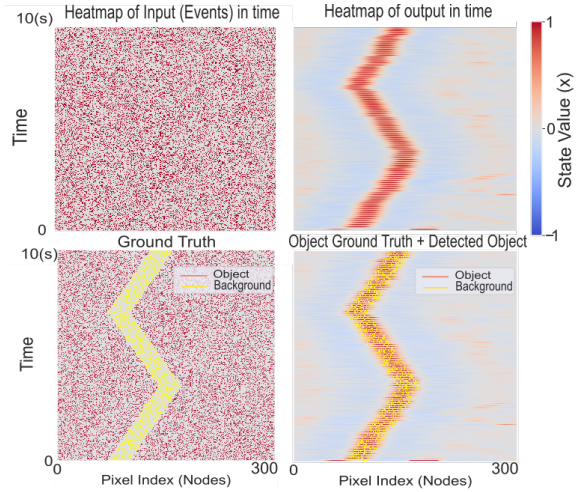

Figure 1: One-dimensional object motion example. **Top left:** raw event input generated by object and background motion; no visually identifiable structure. **Top right:** filtered output suppressing background activity and enhancing coherent object motion. **Bottom:** ground truth object trajectory compared with model output over time.

References

[1] Victoria Clerico, Shay Snyder, Arya Lohia, Md Abdullah-Al Kaiser, Gregory Schwartz, Akhilesh Jaiswal, and Maryam Parsa. Retina-inspired object motion segmentation for event-cameras. In *2025 Neuro Inspired Computational Elements (NICE)*. IEEE, 2025.

[2] Arturo Di Girolamo, Christian Metzner, and Özcan Urhan. An event-based line sensor with configurable antagonistic center surround. In *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2025.

[3] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *CoRR*, abs/1904.08405, 2019.

[4] Alejandro Linares-Barranco, Fernando Perez-Peña, Diederik Paul Moeys, Francisco Gomez-Rodriguez, Gabriel Jimenez-Moreno, Shih-Chii Liu, and Tobi Delbruck. Low latency event-based filtering and feature extraction for dynamic vision sensors in real-time fpga applications. *IEEE Access*, pages 134926–134942, 2019.

On Unwanted Nonlinearity Mitigation and Performance Analysis in Reset Control Systems

Ali Hosseini, Hassan HosseinNia
Department of Precision and Microsystems Engineering, Delft University of Technology, The Netherlands
Email: S.A.Hosseini@tudelft.nl, s.h.hosseinniakani@tudelft.nl

1 Introduction

Reset elements are nonlinear filters that improve control performance beyond linear time-invariant (LTI) limits but introduce higher-order harmonics that complicate design. Although frequency-domain tools like describing functions (DFs) and higher-order sinusoidal-input describing functions (HOSIDFs) analyze reset control systems (RCS), no direct method yet quantifies the impact of higher-order harmonics on the error signal without time-domain simulations.

2 Approach

Despite advancements in HOSIDFs' observations in RCSs [1], a gap remains in the development of a higher-order harmonics quantifier capable of directly revealing the effects of all HOSIDFs on the error signal. Such an observer would enhance our understanding of the reliability of DF-based analyses. Therefore, in this technical note, a robustness factor is considered that integrates higher-order harmonics with the fundamental harmonic using the 2-norm. This robustness factor ($\sigma_2(\omega)$) provides a metric for estimating the additional error introduced by the reset element's nonlinearity during the design process. Importantly, it can be computed entirely in the frequency domain using closed-loop HOSIDFs, thereby significantly reducing the computational cost of predicting the time-domain behavior of reset control systems.

In all HOSIDFs reduction methods, the reduction of the HOSIDFs was achieved by modifying reset control systems with predefined pre- and post-filter structures. However, the reduction obtained with these filters was often insufficient, requiring multiple iterations to identify a suitable filter and consequently achieve the desired performance. Thanks to the introduction of the robustness factor, this process is now streamlined. Thus, a method is proposed for designing pre- and post-filters without imposing a fixed filter structure by constraining the robustness factor to be below a specified threshold ($\sigma_{2,\max}(\omega)$). This ensures that the nonlinearity of the error signal remains within acceptable limits while preserving the DF characteristics and explicitly reducing the HOSIDFs.

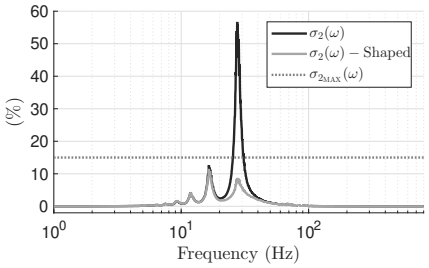

Figure 1: The values for $\sigma_2(\omega)$ (in percent) for both C_{NL} and filtered C_{NL} cases.

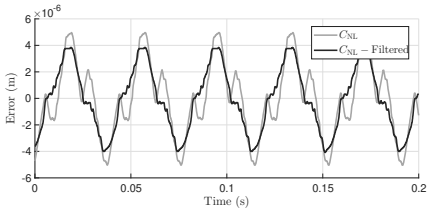

Figure 2: (a) Measured error signals for C_{NL} and filtered C_{NL}.

3 Experimental Results

A reset controller C_{NL} has been designed to improve floor-vibration rejection. Fig. 1 depicts the robustness factor of this controller, indicating that the higher-order harmonics can contribute more than 50% of the fundamental-harmonic contribution to the error. Using the nonlinearity-shaping method presented in [2], we designed pre- and post-filters to reduce these higher-order harmonic contributions to below 15%. Fig. 2 shows the measured error signal for both the filtered and unfiltered cases; the filtered case yields a substantial reduction of higher-order harmonics in the error signal solely by applying these pre- and post-filters.

References

[1] N. Saikumar, K. Heinen, H. HosseinNia, Loop-shaping for reset control systems: A higher-order sinusoidal-input describing functions approach, Control Engineering Practice, 2021.

[2] A. Hosseini, D. Kostić, and H. HosseinNia, "Robust Performance Analysis and Nonlinearity Shaping for Closed-loop Reset Control Systems," *arXiv:2509.02143*, 2025.

Robust performance bounds for component-to-assembly uncertainty propagation in structural dynamics models

R.E.W.L. Thijssen TU/e[1] R.H.B. Fey TU/e[1] V.S. Dolk TU/e[2] L. Hazeleger TU/e[1] N. van de Wouw TU/e[1]

PO Box 513, 5600 MB Eindhoven, the Netherlands, r.e.w.l.thijssen@tue.nl

1 Introduction

As high-tech machinery advances, the requirements on the structural design of those machines become more stringent. Engineers heavily rely on structural dynamics models to meet design specifications. However, models are inherently an approximate representation of the true structure and thus model-based predictions possess uncertainty. This limited predictive capability results in a costly design process, due to expensive design iterations and overengineering to ensure that requirements are met. Hence, the key to improving the cost-effectiveness of the design process lies in better understanding of the uncertainties in structural dynamics models.

2 Damping in interconnected systems

A major source of uncertainty in structural dynamics models stems from damping. Damping originates from various physical mechanisms, most of which are not fully understood, which makes first-principles modeling infeasible. Therefore, damping in structural dynamics models is typically identified by means of experimental modal analysis. As such, a viscous damping model is fitted on these experimental data. Despite all modeling efforts, however, model uncertainty always remains.

High-tech systems are complex assemblies that consist of an interconnection of many components, which are modeled as described above. For engineering design, the dynamic behavior of this complete assembly is of interest, rather than only that of the individual components. Therefore, it is highly relevant to know how the component's model uncertainties propagate through the interconnection to the assembly model, that is, to know to what extent the effects of component uncertainties are amplified or suppressed in the assembly dynamics.

3 Component uncertainty propagation

We propose to leverage methods from the field of robust performance analysis to quantify this model uncertainty propagation. As such, we utilize the modular model reduction framework, as established by Janssen et al. [1]. It describes a mathematical relation between model reduction errors in component FRFs and resulting errors in the FRFs of coupled reduced component models. To this end, it allows the user to specify a norm-bounded region in the components' FRFs, which contains the components' error dynamics. Subsequently, it finds the norm-bounded region in the assembly's FRF, where the corresponding assembly's error dynamics are guaranteed to be contained in.

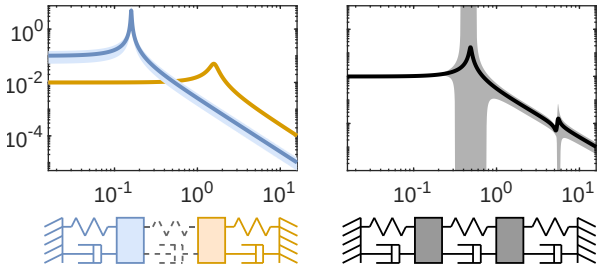

Figure 1: Systems with corresponding magnitude ($\frac{\mathrm{N}}{\mathrm{m}}$) - frequency (Hz) plots, padded areas show uncertainty bounds.

In this work, we apply the robust performance approach to predict how modeling uncertainties in components, as induced by, for instance, a mismatch between modeled and real component damping behavior, affect the uncertainty in predicted assembly dynamics. We illustrate this by the analysis of a use case.

4 Use case

The use case consists of two mass-spring-damper systems, representing two components, which are connected via a spring-damper interface, as shown in Figure 1. A generic uncertainty bound is defined for the FRF of one component, whereas the other component and the interface are assumed to be perfectly modeled. The assembly-level FRF uncertainty bound is computed and compared to one derived by Monte Carlo simulations. This showcases both the effectiveness of the approach and provides novel insight into possibilities for improving conservatism.

5 Conclusions

Robust performance analysis is a promising tool for predicting the propagation of modeling errors in the structural dynamics of components to ward resulting errors in the assembled system dynamics. The approach is demonstrated by a use case, which reveals the effectiveness, as well as points for further improvement, by a comparison with Monte Carlo simulations.

References

[1] L.A.L. Janssen et al., *Automatica*, 160, 111423, 2024.

[1]Eindhoven University of Technology, Department of Mechanical Engineering, Dynamics and Control group

[2]Eindhoven University of Technology, Department of Mechanical Engineering, Control Systems Technology group

Generalized Multilevel Coverage using Hypergraphs

Ashutosh Jindal
Discrete Technology Production and Automation(DTPA)
University of Groningen
Nijenborgh 4,
9747 AG Groningen
The Netherlands
Email: a.jindal@rug.nl

Ming Cao
DTPA
University of Groningen
Nijenborgh 4,
9747 AG Groningen
The Netherlands
Email: m.cao@rug.nl

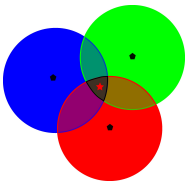

Figure 1: Three mobile sensing units (shown in black) triangulating at a heat source (shown in red) using radial observations

Abstract

Coverage Control problems [1] involve using mobile robots equipped with sensors to monitor a large area. The robots cooperate and obtain information about physical phenomena (static or dynamic) over a bounded domain. The problem objective is to drive robots into a configuration such that, under a predecided metric, optimal coverage is achieved. In the literature, most coverage control problems include scenarios where each point in the region of interest is covered by exactly one robot, and each robot is capable of extracting the information from a given point. However, in certain situations, one requires every point to be covered by more than one robot simultaneously. For instance, in order to triangulate a heat source, it needs to be simultaneously covered by three robots (see Figure 1).

In this work, we generalize the notion of a sensing unit from an individual robot to a team of robots. The robots are grouped into teams (of varied sizes), and these teams are then capable of obtaining information from the environment. This notion of teams further generalizes the idea of k-order coverage [2] where every point is covered by exactly k sensors. We use hypergraphs to model our sensing teams and formulate the coverage control problem. Let $\mathscr{D} \subset \mathbb{R}^2$ be a closed and bounded region. Consider a team of N agents indexed by $[N] := \{1,\dots N\}$. For each $i \in [N]$, let $p_i \in \mathscr{D}$ be the position of robot i and $p := (p_1,\dots,p_N) \in \mathscr{D}^N$ be the sensor configuration.

Sensing Hypergraph: Consider a hypergraph $\mathscr{H} := ([N],\mathscr{E})$, with $\mathscr{E} = \{e_k = (i_1,\dots,i_{l_k}) \mid i_1,\dots,i_{l_k} \in [N]\}$, where $(i_1,\dots,i_{l_k})$ is an edge of length l_k, and with each edge e_k, define $w_k = (w_{k1},\dots,w_{kN})$ to be a vector of length N with nonnegative entries and $w_{ki} > 0$ if and only if $i \in e_k$. Associated with $\mathscr{H}$ define an incidence matrix $E \in \mathbb{R}^{M\times N}$, given by

$$E_{ki} = \begin{cases} w_{ki} & i \in e_k, \\ 0 & \text{otherwise}, \end{cases} \tag{1}$$

where M is the number of unique edges[1]. A sensing unit is defined as a collection of agents $e = \{i_1,\dots,\imath_k\}$ such that $e \in \mathscr{E}$. For each sensor unit k, the sensing perfomance function is define by $F_k(q,p) = 0.5\sum_{i\in e_k} w_{ki}\|q-p_i\|^2$. Let $\mathscr{D} \ni q \longmapsto \Phi(q) \in [0,\infty[$ be a given density function such that $\Phi(q)$ is bounded above. Then, for a given configuration p, the coverage metric is given by

$$H(p) = \int_{\mathscr{D}} V(q,p)\cdot E \cdot F(q,p)\Phi(q)\mathrm{d}q$$

where $V(q,p) = [\dots,V_k(q,p),\dots] \in \mathbb{R}^M$ with $V_k(q,p) = 1$ for all q such that $F_k(q,p) \le F_{k'}(q,p)$ for all $1 \le k' \le M$, $k' \neq k$, and $F(q,p) = [\dots,\|q-p_i\|^2,\dots] \in \mathbb{R}^N$.

We consider a kinematic model for robots $\dot{p}_i = u_i \in \mathbb{R}^2$. The coverage control problem is then finding a control input u such that from any given initial configuration $p(0)$, agents converge to a p^* such that p^* is a local minimizer of H.

Result: Let $m_i := \int_{\mathscr{D}} V(q,p)\cdot E\cdot \pi_i\Phi(q)\mathrm{d}q$ and define $c_i = (\int_{\mathscr{D}} V(q,p)\cdot E\cdot\pi_i\Phi(q)q\mathrm{d}q)/m_i$ as the generalized centroid, where $\pi_i = [\dots,1,\dots] \in \mathbb{R}^N$ is the canonical basis vector with the ith entry 1 and 0 otherwise, then $u_i = -\gamma_i(p_i - c_i)$, $\gamma_i > 0$ drives agents to a locally optimal configuration and such a configuration coincides with $p_i^* = c_i$.

References

[1] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[2] Q. Li and D. Spinello, "Higher-order non-autonomous optimal area coverage control," *IEEE Control Systems Letters*, 2024.

[1] Two edges e_k, e_m are unique if k^{th} and m^{th} rows of I differ in at least one entry

Safe control for flexible link manipulators using Velocity-bounding Proxy-based Sliding Mode Control

Hafsa Nouhi, Chaoyue Fei, Thierry Hubert, Greet Van de Perre
Brubotics, Vrije Universiteit Brussel, and IMEC, Belgium
Email: Hafsa.Nouhi@vub.be

1 Introduction

Modern collaborative robots are designed to operate safely alongside humans, but their high inertia and structural rigidity present significant safety challenges. Reducing inertia and incorporating flexible structures can enhance safety; however, these design choices introduce complexities in control. We propose a control approach, called Flexible-link Velocity-bounding Proxy-Based Sliding Mode Control (FL-VB-PSMC) adapted for flexible-link manipulators by taking into account the flexural effects within the controller. The goal is to ensure accurate and safe control, maintaining precision and addressing the challenges posed by flexibility and safety.

2 Control

FL-VB-PSMC is a control strategy that incorporates a virtual proxy object into the control framework [1]. The proxy is controlled using the principles of sliding mode control (SMC) and is connected to the actual system through a virtual coupling, conceptualized as a damped spring. The PSMC approach addresses the limitations of traditional SMC, particularly the chattering effect. The control law is formulated as follows :

$$\tau_{\mathrm{vb-psmc}} = T\,\mathrm{sat}\left[\frac{1}{T}\left(K_d V\,\mathrm{sat}\left(\frac{\phi_p+\ddot{a}+\dot{\theta}}{V}\right) + K_i a + K_p \dot{a} - K_d \dot{\theta}\right)\right]. \quad (1)$$

With the sliding surface, s, defined as :

$$s = V\,\mathrm{sat}\left(\frac{\phi_p+\ddot{a}+\dot{\theta}}{V}\right) - \ddot{a} - \dot{\theta}. \quad (2)$$

Where $\phi_p = A\left[(\theta_\mathrm{d}-\theta)+H\left(\dot{\theta}_\mathrm{d}+v_\mathrm{c}-\dot{\theta}\right)-\dot{a}-H\ddot{a}\right]$, with θ_d and $\dot{\theta}_d$ the desired angular position and velocity, θ and $\dot{\theta}$ are the angular position and velocity of the controlled object. A, an arbitrary constant value while a represents an internal state, defined as $a = \int(\theta_p-\theta)\,dt$ with θ_p the proxy's angular position. The particularity of this sliding surface is the presence of a magnitude limit to ensure constrained velocities. This constraint prevents excessively large velocities that may arise when the positional error is high and the system has not yet reached the sliding manifold. As long as the actuator torque limit T is sufficiently large to counteract disturbances, the controller drives the proxy to the sliding surface. Once the proxy reaches this surface, it exhibits smooth, exponential convergence toward the target trajectory, governed by the time constant H and a constrained velocity. In the presence of an external disturbance exceeding T, the virtual coupling transfers this to the proxy, causing it to deviate from the sliding surface. Consequently, the end effector, constrained by the virtual coupling, reacts compliantly to the disturbance. A vibration control is integrated by incorporating a strain-feedback loop (v_c) at the velocity level to effectively dampen the oscillations caused by the inherent flexibility of the link [2].

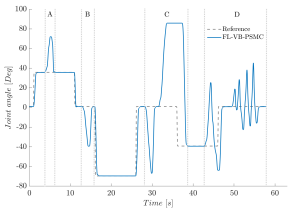

Figure 1: Joint position upon physical interaction at random locations along the link.

3 Results

FL-VB-PSMC was tested experimentally on a 1-DOF flexible-link manipulator. To demonstrate the controller's response to human interactions, the manipulator was instructed to follow a fixed reference while its motion was perturbed by a human at random locations along the link. Fig. 1 visualizes the desired and actual joint position. Regions A, B and C correspond to interactions applied at different locations along the link, while in region D a cyclic interaction was induced. It can be seen that the FL-VB-PSMC reacts compliantly to interactions induced while converging smoothly to the reference when the contact is released. Even under large positional, (reaching up to $\pm 120^\circ$ in region C), the flexible link smoothly returns to the desired trajectory, with an overdamped and vibration-free recovery motion. Future work includes extending FL-VB-PSMC to an adaptive controller.

Acknowlegment: This work was funded by FWO research grant G0A9623N and 12APZ24N.

References

[1] Nouhi, H., Fei, C., Hubert, T. & Perre, G. Flexible-Link Velocity-Bounding Proxy Based Sliding Mode Control. *IEEE Robotics And Automation Letters*. **11** pp. 145-152 (2026)

[2] Perre, G., Hubert, T., Verstraten, T. & Vanderborght, B. Investigating the Potential of Flexible Links for Increased Payload to Mass Ratios for Collaborative Robotics. *IEEE Access*. **11** pp. 15981-15995 (2023)

A Learning Semi-Active Tuned Mass Damper Design

Thijs Romberg[1] Kees Verbaan[2] Hans Butler[3,4] Marcel Heertjes[1,4]

[1]Dynamics & Control, Mechanical Engineering, Eindhoven University of Technology
[2]Control Systems Technology, Mechanical Engineering, Eindhoven University of Technology
[3]Control Systems, Electrical Engineering, Eindhoven University of Technology
[4]ASML, Veldhoven, The Netherlands
PO Box 513, 5600 MB, Eindhoven, The Netherlands, Email: T.L.Romberg@tue.nl

1 Introduction

In the semiconductor industry there is a drive to increase the computing power density per chip. For this to become a reality, the size of features projected on a silicon disc, called a wafer, by a lithography machine has to become smaller. For this process to remain economically viable, the production throughput has to increase simultaneously.
A lithography machine typically consists of multiple motion systems each inducing vibrations. Classical examples are the wafer and reticle stage. These systems influence the throughput and accuracy of the system.

2 Tuned Mass Damper

The bandwidth of a motion system is typically limited by the frequency and magnitude of the first resonance peak created by the internal stiffness and damping of the system [1]. To increase performance, it is possible to reduce the magnitude of this resonance peak by implementing a Tuned Mass Damper (TMD). A schematic diagram is illustrated in Figure 1, where a primary structure is shown representing the mechatronic system, which consists of a mass, m_1, stiffness, c_1, and damper, d_1. A TMD is connected to the primary structure which can be resembled by a mass, m_2, stiffness, c_2, and damper, d_2. A TMD is tuned such that it counteracts the limiting resonance. In doing so, it moves in anti-phase with the primary structure.

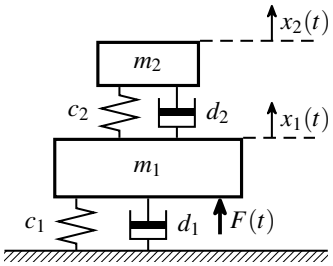

Figure 1: Schematic diagram of a TMD, where m_1, c_1 and d_1 resemble the parameters of the primary structure and m_2, c_2 and d_2 of the TMD.

3 Problem Description

To determine the best parameters for a TMD, analytical formulas were derived by J.P. Den Hartog [2]. These formulas determine the best stiffness and damping for the given mass ratio, $\mu = \frac{m_2}{m_1}$. To show the result of implementing a TMD, Figure 2 shows the results with these optimal parameters, in which the black line corresponds to the transfer function of the system shown in Figure 1 without m_2, c_2 and d_2 from the force input $F(t)$ to the position output $x_1(t)$. By implementing an optimal TMD using the formulas presented in [2] the resonance is damped almost 2 orders of magnitude, shown by the red line in Figure 2.
If the primary structure changes due to for example wear, the resonance frequency could move and the TMD is not optimal anymore. This phenomenon is showcased by the blue line in Figure 2, in which the stiffness c_1 is lowered by 10 % which results in a magnitude increase of 6 dB. This increase could be problematic for stability and performance of the controlled system.
To solve this problem a learning Semi-Active Tuned Mass Damper will be designed, which is able to alter its parameters based on the structure it is attached to. For this device to operate, machine learning techniques will be used to optimize the TMD parameters.

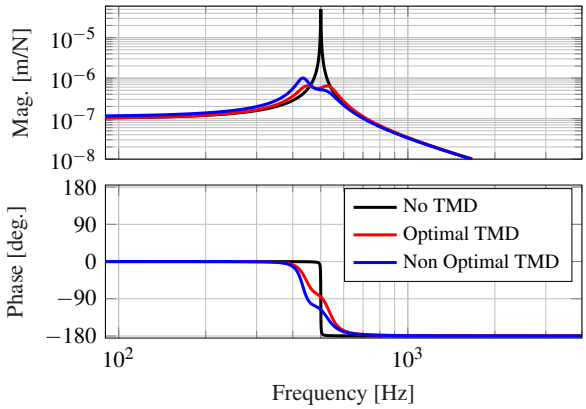

Figure 2: Bode diagram of a structure without TMD, ideally tuned and out of tune TMD.

4 Acknowledgements

This project is co-financed by Holland High Tech, top sector High-Tech Systems and Materials, with a PPP innovation subsidy for public-private partnerships for research and development.

References

[1] R. M. Schmidt, G. Schitter, A. Rankers, and J. van Eijk, *The design of high performance mechatronics*, 3rd ed. Amsterdam: IOS Press, 2020.

[2] J. P. Den Hartog, *Mechanical vibrations*, 4th ed. New York u.a.: McGraw-Hill, 1968.

Current control challenges for airborne laser communications

Mauricio Sotelo González[1,2], Gert Witvoet[1], Lex Meijer[2], Tom Oomen[1], Niek Doelman[1]
[1]Eindhoven University of Technology, [2]Airbus Defence and Space Netherlands
Email: {m.sotelo.gonzalez1, g.witvoet, T.A.E.Oomen, n.j.doelman}@tue.nl
lex.meijer@airbus.com

1 Background

Free Space Optical Communications (FSOC) is a technology that has been under development for more than 20 years, especially in the Space sector, with the first inter-satellite link in 2001 achieved by ESA [1]. Currently, several companies, including SpaceX, are launching full satellite constellations that rely on inter-satellite laser communications since 2021. While Ground to Space links (G2S) have also been achieved [2], FSOC with mobile terminals, such as aircrafts, have not been established in current markets due to factors such as platform motion and vibration, field of regard specification, or SWaP constraints. The adoption of this technology could solve the bandwidth problem that is beginning to occur, reaching communication links of 200 Gbps [3].

2 Current control challenges

There are two main current challenges in FSOC systems: pointing, acquisition, and tracking (PAT); and atmospheric effects. Focusing on PAT, the initial pointing of both terminals must be accurate enough, in the order of 10^{-4} rad to reduce the uncertainty area that must be scanned. This is really hard constrained, in the case of an aircraft, to the positioning and orientation sensors it has, and if its trajectory is stable. For satellites, their trajectory is well established.

The tracking process is a challenge since the control architecture must deal with high frequency vibrations that depend on the terminal position in the aircraft. These high frequency contributions are not present in ground terminals that have more stable and isolated environment; neither in satellites since having less friction while moving, reduces induced vibrations due to a better conditioned environment. The tracking error specification is usually in the order of 10^{-6} rad, which is a significant challenge in terms of controller design.

The control should have a high noise rejection factor while maintaining a high bandwidth. High frequency components can be reduced by passive isolation systems, but that would imply larger SWaP terminal characteristics. This accuracy requirement is due to possible large link distance that can be up to 40.000 km if it is a GEO satellite and the targeted optical aperture is usually of 30 cm (typical size of CubeSats based satellites).

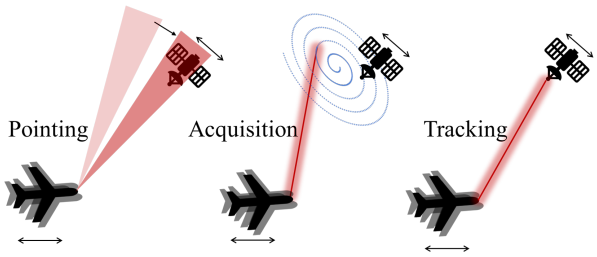

Figure 1: A2S FSO scheme where both platforms are moving and the aircraft is under high frequency vibrations. Both, aircraft and satellite are performing the same process in each stage

On top of tracking performance in terms of error, due to the nature of the platform, the terminal should target to have a full hemi-spherical field of regard. This usually involves a larger number of moving parts that make vibration isolation difficult. In relation to it, low drag and low volume design is important in order to install it in an aircraft fuselage.

Funded by the European Union

This work was funded by the European Union under the Horizon Europe programme (Grant Agreement No. 101 169 042)

References

[1] *A world first: Data transmission between European satellites using laser light*. Nov. 22, 2001.

[2] Ricardo Barrios et al. “Link budget assessment for GEO feeder links based on optical technology”. In: *International Journal of Satellite Communications and Networking* 39.2 (Mar. 2021), pp. 160–177.

[3] Curt M. Schieler et al. “On-orbit demonstration of 200-Gbps laser communication downlink from the TBIRD CubeSat”. In: *Free-Space Laser Communications XXXV*. Free-Space Laser Communications XXXV. San Francisco, United States: SPIE, Mar. 15, 2023.

Evolutionary Dynamics in Continuous-time Finite-state Mean Field Games

Leonardo Pedroso[1], Andrea Agazzi[2], W.P.M.H. (Maurice) Heemels[1], Mauro Salazar[1]
[1]Control Systems Technology section, Eindhoven University of Technology, Eindhoven, The Netherlands
[2]Institute for Mathematical Statistics and Actuarial Sciences, Alpeneggstrasse 22, 3012 Bern

1 Introduction

We study a dynamic game with a large population of players who choose actions from a finite set in continuous time. Each player has a state in a finite state space that evolves stochastically with their actions. A player's reward depends not only on their own state and action but also on the distribution of states and actions across the population, capturing effects such as congestion in traffic networks. While prior work in evolutionary game theory has primarily focused on static games without individual player state dynamics, we present the first comprehensive evolutionary analysis of such dynamic games. We propose an evolutionary model together with a mean field approximation of the finite-population game and establish strong approximation guarantees. We show that standard solution concepts for dynamic games lack an evolutionary interpretation, and we propose a new concept – the Mixed Stationary Nash Equilibrium (MSNE) – which admits one. We analyze the relationship between MSNE and the rest points of the mean field evolutionary model and study the evolutionary stability of MSNE.

2 MSNE and Evolutionary Equilibria

A solution concept grounded in an *evolutionary* model offers a more compelling notion of a game's outcome than one that is not, such as the NE, because it offers insight into how the outcome emerges under very limited assumptions about the players' knowledge. In *static games*, the rest points of the revision dynamics – the natural evolutionary solution concept – are NE for most meaningful families of revision protocols. In contrast, in *dynamic games*, state-of-the-art solution concepts cannot be given an evolutionary interpretation. The goal of this paper is to initiate a formal *evolutionary analysis of dynamic games*.

We formally define *behavioral* equilibria and show that they do not have an evolutionary interpretation. The key reason is that there is no room for heterogeneity in the players' behavior, which is essential for defining individual revision dynamics. To address this, we propose a *mixed* solution concept, called MSNE, that allows for such heterogeneity. To the best of our knowledge, such a solution concept has not been studied for the class of dynamic games considered in this paper.

The following two informal results of [1] establish a relation between a rest point of the evolutionary dynamics and the MSNE solution concept for an average payoff. Similar results are derived for a discounted payoff in [2].

Theorem 1 *Consider an imitative via comparison, excess payoff, or pairwise comparison revision protocol, then a MSNE is a rest point of the evolutionary dynamics.*

Theorem 2 *Consider an excess payoff or pairwise comparison revision protocol, then a rest point of the evolutionary dynamics is a MSNE.*

3 Evolutionary Stability of MSNE

In [3] we investigate the evolutionary stability of MSNE. We derive conditions on both the structure of the MSNE and the game's payoff map that ensure local and global stability under evolutionary dynamics. These results characterize when MSNE can robustly emerge and persist against strategic deviations, thereby providing insight into its long-term viability in large population dynamic games.

First, we show that non-MSNE rest points are unstable under evolutionary dynamics even if some (or all classes) use imitative via comparison revision protocols. Second, we show that if a MSNE is strict, i.e., each class places mass on a single policy with strictly higher payoff than all others, then the MSNE is locally asymptotically stable under meaningful revision protocols. Third, in a two time-scale regime where state dynamics are significantly faster than revision dynamics, we conclude that there are two classes of games (potential and stable games) for which all trajectories approach the whole set of MSNE. Fourth, in the two time-scale regime, we establish that a simple condition on the payoff structure of the game guarantees that all trajectories initialized sufficiently close to a subset of MSNE remain close to it.

References

[1] L. Pedroso, A. Agazzi, W. P. M. H. Heemels, and M. Salazar, "Evolutionary dynamics in continuous-time finite-state mean field games - Part I: Equilibria," 2025, arXiv preprint arXiv:2511.01452.

[2] ——, "Evolutionary analysis of continuous-time finite-state mean field games with discounted payoffs," 2025, arXiv preprint arXiv:2511.07655.

[3] ——, "Evolutionary dynamics in continuous-time finite-state mean field games - Part II: Stability," 2025, arXiv preprint arXiv:2511.03297.

Abstraction-Based Model Reduction of Uncertain Systems

Melvyn D. Wildeboer† m.d.wildeboer@rug.nl
Jacquelien M.A. Scherpen† j.m.a.scherpen@rug.nl
Bart Besselink‡ b.besselink@rug.nl
† Engineering and Technology Institute Groningen, Groningen, the Netherlands
‡ Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, Groningen, the Netherlands

1 Introduction

Motivated by the rapid increase in complexity of mathematical models of modern engineering systems, together with the inherent uncertainty that any mathematical model exhibits, model reduction of uncertain systems aims to accurately approximate the dynamics of high-order uncertain systems by models with a lower complexity. We propose an abstraction-based approach to model reduction of uncertain systems. Contrary to existing model reduction techniques for uncertain systems [1, 2], which focus on *behavioural approximation*, abstraction-based model reduction aims for *containment* of external behaviour, proving to be beneficial in specification verification for high-order systems.

2 Uncertain Systems

We consider uncertain systems as in Fig. 1, where the *nominal system* $G : L_2(\mathbb{R}_+, \mathbb{R}^{m+q}) \to L_2(\mathbb{R}_+, \mathbb{R}^{p+r})$ is a state-space operator, that is, the input-output operator of a stable linear time-invariant system. We define the *order* of a state-space operator G as the dimension of a minimal state-space representation. The variable $u \in L_2(\mathbb{R}_+, \mathbb{R}^m)$ denotes a control input, and $y \in L_2(\mathbb{R}_+, \mathbb{R}^p)$ is the corresponding measured output. The nominal system is perturbed by uncertain signals $d \in L_2(\mathbb{R}_+, \mathbb{R}^q)$, generated by an output $z \in L_2(\mathbb{R}_+, \mathbb{R}^r)$ and an operator Δ, the latter unknown but assumed to belong to the following class of linear operators

$$\mathbb{\Delta} := \{\, \Delta : L_2(\mathbb{R}_+, \mathbb{R}^r) \to L_2(\mathbb{R}_+, \mathbb{R}^q) \quad | \quad ||\Delta|| \leq 1 \,\}.$$

The set $\mathbb{\Delta}$ is referred to as the *uncertainty set*. We denote the uncertain system comprised of a nominal system, G, and an uncertainty set, $\mathbb{\Delta}$, by $(G, \mathbb{\Delta})$.

The interconnection in Fig. 1 is assumed to be described by

$$\begin{bmatrix} y \\ z \end{bmatrix} = \begin{bmatrix} G_{11} & G_{12} \\ G_{21} & 0 \end{bmatrix} \begin{bmatrix} u \\ d \end{bmatrix}, \qquad d = \Delta z, \tag{1}$$

where G_{11}, G_{12}, and G_{21} are state-space operators, compatibly partitioned with respect to the system's inputs and outputs. We regard the uncertain system $(G, \mathbb{\Delta})$ as a nominal mapping from u to y, given by G_{11}, while G_{12} and G_{21} prescribe how this mapping is perturbed by the uncertainty Δ. Therefore, we define the order of the uncertain system $(G, \mathbb{\Delta})$ as the order of the operator G_{11}.

From (1), we have that the mapping from u to y is given by

$$S(G, \Delta) := G_{11} + G_{21}\Delta G_{12}.$$

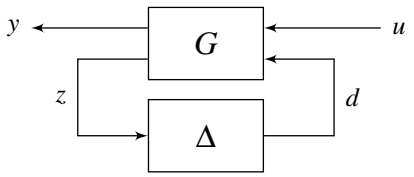

Figure 1: Uncertain system model

From this, we see that the external behaviour of the uncertain system $(G, \mathbb{\Delta})$, defined as all generable input-output trajectories of the system, is given by

$$\mathscr{B}(G, \mathbb{\Delta}) := \{\, (u, y) \mid \exists\, \Delta \in \mathbb{\Delta} \text{ such that } y = S(G, \Delta)u \,\}.$$

3 Problem Formulation

We propose an abstraction-based approach to performing model reduction for the class of uncertain systems as in Fig. 1. In this setting, we aim to construct a reduced-order uncertain system such that its external behaviour contains that of the high-order system. More explicitely, we consider the following problem.

Problem 1 *Given an uncertain system $(G, \mathbb{\Delta})$ of order n, find an uncertain system $(G_r, \mathbb{\Delta})$ of order $k < n$ such that*

$$\mathscr{B}(G, \mathbb{\Delta}) \subset \mathscr{B}(G_r, \mathbb{\Delta}). \tag{2}$$

4 Results

We show that under certain assumptions on the nominal systems, G and G_r, the external behaviour of both uncertain systems can be expressed as the solution set of an Integral Quadratic Constraint (IQC). Using well-established results on frequency-domain spaces and IQCs, we establish a verification technique for the behavioural containment in (2) in terms of a frequency-domain matrix inequality. Moreover, we develop a procedure for constructing a reduced-order uncertain system $(G_r, \mathbb{\Delta})$ that satisfies (2).

References

[1] C. L. Beck, J. Doyle, and K. Glover, "Model reduction of multidimensional and uncertain systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 10, pp. 1466-1477, 1996.

[2] L. Li and I. Petersen, "A gramian-based approach to model reduction for uncertain systems," *IEEE Transactions on Automatic Control*, vol. 55, no. 2, pp. 508-514, 2010.

Confidentiality of Linear Systems with Quadratic Output

Zeyad M. Manaa, Nathan van de Wouw, Michelle S. Chong

Dynamics and Control, Eindhoven University of Technology, The Netherlands. Emails: {z.manaa, n.v.d.wouw, m.s.t.chong}@tue.nl.

1 Introduction

Modern cyber-physical systems rely on networked sensors to monitor physical processes. An adversary with read and write access to sensor data can infer internal states and launch more dangerous, undetectable attacks that disrupt system performance while remaining undetected [1]. We adopt the adversary perspective, which is formulated as an estimation problem for a linear plant with a quadratic sensor output map, stabilized by a dynamic controller. We assume the closed-loop system has inherent stability properties. Yet the attack-free closed loop is unobservable at equilibrium due to the quadratic output map, which precludes observer designs based on local observability. We therefore ask whether an adversary can still estimate the plant and controller states. Indeed, this can be achieved by inducing observability by manipulating the sensor output while preserving closed-loop stability to avoid detection by existing abnormality detectors in the system.

2 Problem formulation

We consider a linear plant with scalar quadratic output $\Sigma_p:\ \dot{x}_p = A_p x_p + B_p u,\ y = x_p^\top Q_p x_p$, controlled by the dynamic output-feedback controller $\Sigma_c:\ \dot{x}_c = A_c x_c + B_c y,\ u = C_c x_c + D_c y$, where $x_p \in \mathbb{R}^{n_p}$, $x_c \in \mathbb{R}^{n_c}$, and $u, y \in \mathbb{R}$. Let $z := [x_p^\top\ x_c^\top]^\top \in \mathbb{R}^n$, the closed-loop dynamics are

$$\dot{z} = Az + Bh(z) =: f(z), \qquad h(z) = z^\top Q z, \tag{1}$$

with $Q = \begin{bmatrix} Q_p & 0 \\ 0 & 0 \end{bmatrix}$, $A = \begin{bmatrix} A_p & B_p C_c \\ 0 & A_c \end{bmatrix}$, $B = \begin{bmatrix} B_p D_c \\ B_c \end{bmatrix}$. We assume $z = 0$ is locally exponentially stable equilibrium, i.e., A is Hurwitz. The linearization of (1) is $F = \partial f/\partial z|_0 = A$ and $H = \partial h/\partial z|_0 = 0$, and is therefore unobservable at the origin. We assume the adversary knows (A, B, Q), can read and additively manipulate the sensor output such that $\tilde{y} = y + a$, with $a \in \mathbb{R}$, but has no access to the control input u. With the manipulated sensor measurements, the controller now sees $\tilde{y}$ instead of y and becomes $\dot{x}_c = A_c x_c + B_c \tilde{y},\ u = C_c x_c + D_c \tilde{y}$. In turn, the attacked version of (1) can be written as

$$\dot{z} = Az + B\big(h(z) + a\big) =: \bar{f}(z, a), \qquad h(z) = z^\top Q z. \tag{2}$$

Figure 1 illustrates the problem setup.

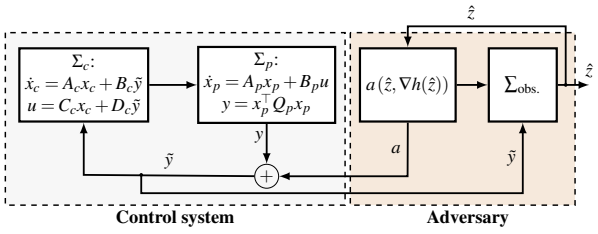

Figure 1: Problem setup. The block $\Sigma_{\text{obs.}}$ denotes the adversary's observer that will be designed.

In short, the adversary's objectives are:

Objective 1. *Design the injected signal a using the measured output $h(z)$ and the gradient $\nabla h(\hat{z}) = 2Q\hat{z}$ computed from the observer estimate $\hat{z}$ (from Objective 2), such that it renders system* (2) *locally observable at $z = 0$, i.e., the linearization of* (2) *is observable.*

Objective 2. *Design a state observer to get an estimate $\hat{z}$ of the state z of* (2)*, so that the origin is a locally exponentially stable equilibrium of the estimation error system $(z(t) - \hat{z}(t))$, with rate $\alpha > 0$.*

Objective 3. *Design the attack signal a such that the origin remains a locally exponentially stable equilibrium of the closed-loop system* (2)*.*

3 Approach

To achieve Objective 1, we design the attack as a projection of the estimated gradient using the observer state $\hat{z}$. Let $\pi \in \mathbb{R}^{n_p}$ be a projection direction and $\bar{\pi} := \big[\pi^\top,\ 0^\top\big]^\top \in \mathbb{R}^n$. Since $\nabla h(\hat{z}) = 2Q\hat{z}$, the attack is therefore $a(\hat{z}) := \bar{\pi}^\top \nabla h(\hat{z}) = \bar{H}(\pi)\hat{z}$ where $\bar{H}(\pi)[2\pi^\top Q_p \quad 0]$. Linearizing (2), at the origin yields the pair $(\bar{F}(\pi), \bar{H}(\pi))$ where $\bar{F}(\pi) = A + B\bar{H}(\pi)$. We design π such that $(\bar{F}(\pi), \bar{H}(\pi))$ is observable, thereby inducing local observability. We do so by applying the Hautus test to characterize the inadmissible set Π_{unobs} for which $(\bar{F}(\pi), \bar{H}(\pi))$ is unobservable, and then selecting $\pi \notin \Pi_{\text{unobs}}$.

Next, to achieve Objectives 2 and 3, we propose the observer

$$\Sigma_{\text{obs.}}: \dot{\hat{z}} = A\hat{z} + B(\hat{z}^\top Q\hat{z} + 2\bar{H}(\pi)\hat{z}) - L(\hat{z}^\top Q\hat{z} + 2\bar{H}(\pi)\hat{z} - \tilde{y}).$$

where L is the observer gain to be designed and π selected as discussed above. Recall that A is Hurwitz, and that the adversary's model as stated in section 2. In addition, assume $\sigma(A_p) \cap \sigma(A_c) = \emptyset$, and $B_p C_c \neq 0$ where $\sigma(\cdot)$ denotes the set of eigenvalues of a matrix. Under these assumptions, we design π and L such that $(\bar{F}(\pi), \bar{H}(\pi))$ is observable, and such parameters always exist under the recalled assumptions. The gain L then assigns the observer error eigenvalues to $\{s \in \mathbb{C} : \text{Re}(s) \leq -\alpha\}$, ensuring local exponential convergence with rate α.

4 Conclusion

We proposed a systematic sensor-attack design that induces local observability in closed-loop systems with linear dynamics and quadratic outputs while preserving stability. A Luenberger-type observer then guarantees local exponential convergence of plant and controller state estimates. Future work will address broader classes of nonlinearities.

References

[1] S. M. Dibaji, M. Pirani, D. B. Flamholz, A. M. Annaswamy, K. H. Johansson, and A. Chakrabortty, "A systems and control perspective of CPS security," *Annu. Rev. Control*, vol. 47, 2019.

Global Synchronization Preserving Model Reduction of Nonlinear Network Systems

Yangming Dou and Jacquelien M. A. Scherpen
University of Groningen
Email: {y.dou, j.m.a.scherpen}@rug.nl

Xiaodong Cheng
Wageningen University & Research
Email: xiaodong.cheng@wur.nl

Yu Kawano
Hiroshima University
Email: ykawano@hiroshima-u.ac.jp

1 Introduction

With the increasing scale of network systems, their analysis and control design become challenging, motivating the development of model reduction techniques [1]. Our former works [2] focus on reducing the topology of nonlinear network systems. In this work, we investigate how to reduce nonlinear subsystems in a homogeneous network system while preserving its synchronization. To this end, a matrix inequality is introduced to characterize the global exponential synchronization, and then a differential balanced truncation method [3] to reduce the dimension of each subsystem. We further derive an error bound between the original and reduced subsystem.

2 Differential Balancing for Nonlinear Subsystems

Consider the state space model of each subsystem

$$\mathbf{\Sigma}_i : \dot{x}_i = f(x_i) + Bu_i,\ y_i = Cx_i,\ i = 1,2,\cdots,N, \quad (1)$$

where $x_i(t) \in \mathbb{R}^n$, and $u_i(t) \in \mathbb{R}^m, y_i(t) \in \mathbb{R}^m$. The function $f(\cdot)$ is assumed to be $\mathscr{C}^1$, and B, C are constant matrices with proper dimensions. The subsystems on the network are interconnected according to $u_i = \sum_{j=1}^N w_{ij}(y_i - y_j), i = 1,2,\cdots,N$, where $w_{ij} \in \mathbb{R} \geq 0$ is the (i,j)-th entry of weighted adjacency matrix. A compact form of a nonlinear dynamical network can be written as

$$\mathbf{\Sigma} : \dot{x} = \mathscr{F}(x) - (L \otimes BC)x,\ y = (I_N \otimes C)x, \quad (2)$$

where $x := \left[x_1^\top, x_2^\top, \cdots, x_N^\top\right]^\top$, $y := \left[y_1^\top, y_2^\top, \cdots, y_N^\top\right]^\top$, $\mathscr{F}(x) := \left[f(x_1)^\top \quad f(x_2)^\top \quad \cdots \quad f(x_N)^\top\right]^\top$. The matrix $L \in \mathbb{R}^{N\times N}$ is the graph Laplacian of the underlying network.

If there exist a symmetric matrix positive definite $K \in \mathbb{R}^{n\times n}$ and $\varepsilon > 0$ such that

$$K\frac{\partial f(x_i)}{\partial x_i} + \frac{\partial f(x_i)}{\partial x_i}^\top K + KBB^\top K + \rho C^\top C \preceq -\varepsilon K, \quad (3)$$

where $\rho = \min_k \|(I + \mathbb{1}\mathbb{1}^T)L_{[k,k]}L_{[k,k]}^\top(I + \mathbb{1}\mathbb{1}^T)\|_2$, then the network system $\mathbf{\Sigma}$ in (2) is globally exponentially synchronized. Consider (3) and differential Lyapunov inequality

$$Q\frac{\partial f(x_i)}{\partial x_i} + \frac{\partial f(x_i)}{\partial x_i}^\top Q + C^\top C \preceq 0, \quad (4)$$

where K^{-1} and $Q \succ 0$ are generalized differential controllability and observability Gramians of system (1). There exist a transform matrix T such that $T^{-1}K^{-1}T^{-\top} = T^\top QT = \mathrm{diag}(\theta_1, \theta_2, \cdots, \theta_n) = \Theta$, where $\theta_1 \geq \theta_2 \geq \cdots \geq \theta_n$ are Hankel singular values. Following the procedure of differential balancing [3], with (3) and (4), we can obtain the reduced-order model after balanced truncation:

$$\hat{\mathbf{\Sigma}}_i : \dot{\hat{x}}_i = \hat{f}(\hat{x}_i) + \hat{B}\hat{u}_i,\ \hat{y}_i = \hat{C}\hat{x}_i,\ i = 1,2,\cdots,N, \quad (5)$$

where $\hat{x}_i(t) \in \mathbb{R}^r$ with $r \ll n$, and $\hat{u}_i(t), \hat{y}_i(t) \in \mathbb{R}^m$. The function $\hat{f}(\cdot)$ is assumed to be $\mathscr{C}^1$, and $\hat{B}$, $\hat{C}$ are constant matrices with proper dimensions. Furthermore, we can show that the network system with the reduced nonlinear subsystems remains globally exponentially synchronized.

Error Bound. Let $\hat{x}_i^j$, $j = r,\ldots,n$ denote the state of the j-dimensional reduced order model ($\hat{x}_i^n = \bar{x}_i$). Accordingly, define $\Theta_j := \mathrm{diag}\{\theta_1,\ldots,\theta_j\}$, $j = r+1,\ldots,n$ ($\Theta_n = \Theta$). Assume that $f(-x_i) = -f(x_i)$. The outputs of the systems (1) and (5) under the same input v_i satisfy $\|y_i - \hat{y}_i\|_{L_2} \leq 2\sum_{j=r+1}^n \theta_j \|v_i\|_{L_2} + c$, where c depends on the initial states:

$$c := \sum_{j=r+1}^n \left(\left| \hat{x}_i^j(0) + \begin{bmatrix} \hat{x}_i^{j-1}(0) \\ 0 \end{bmatrix} \right|_{\Theta_j} + \theta_j \left| \hat{x}_i^j(0) - \begin{bmatrix} \hat{x}_i^{j-1}(0) \\ 0 \end{bmatrix} \right|_{\Theta_j^{-1}} \right)$$

with $|x|_{\Theta_j} := \sqrt{x^\top \Theta_j x}$.

References

[1] X. Cheng and J. M. A. Scherpen, "Model reduction methods for complex network systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 425–453, 2021.

[2] Y. Dou, X. Cheng, and J. M. A. Scherpen, "Stability-preserving model reduction of networked Lur'e systems," in *2024 European Control Conference (ECC)*. IEEE, 2024, pp. 2059–2064.

[3] Y. Kawano and J. M. A. Scherpen, "Model reduction by differential balancing based on nonlinear Hankel operators," *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3293–3308, 2016.

Neural Linear Oscillator Networks and their Application to Soft Robots Model Learning and Control

Jingyue Liu Ebrahim Shahabi Cosimo Della Santina

Faculty of Mechanical Engineering, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands

Email: {J.Liu-14, E.ShahabiShalghouni, C.DellaSantina}@tudelft.nl

1 Abstract

Soft robots exhibit highly nonlinear, high-dimensional dynamics that are hard to model from pixels while remaining suitable for control. We propose Neural Linear Oscillator Networks (nLON), a vision-based latent dynamics model where a convolutional autoencoder encodes two consecutive images into a compact position–velocity latent state, and the latent evolution is constrained to a network of controlled linear mechanical oscillators. This structural prior yields a Wiener-type mechanical system, enabling stable long-horizon rollouts and direct LQR-based controller synthesis. We validate nLON on simulated and real soft tentacle systems, showing improved multi-step prediction quality over ConvLSTM and Koopman baselines and demonstrating robust reference tracking using model-based control with either internal-model propagation or pixel-level feedback.

2 Introduction

Modeling dynamics from high-dimensional sensory observations is central to robotics, and is particularly challenging for soft robots due to continuum deformation and strong nonlinearities. Latent-variable models compress images to low-dimensional states, but learning accurate and controllable latent dynamics remains difficult. Neural oscillator architectures [1] provide compact recurrent dynamics [2]; here we explicitly enforce a mechanical oscillator structure to obtain both interpretability and control compatibility.

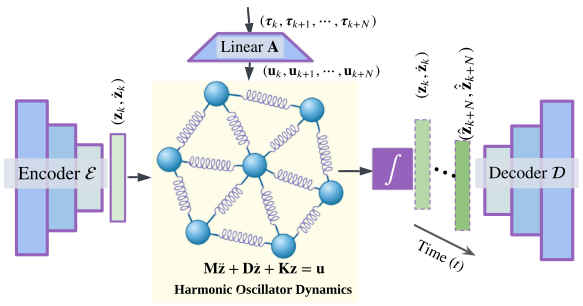

Figure 1: Overview of the proposed nLON modeling framework.

3 Method and Results

We propose Neural Linear Oscillator Networks (nLON) shown in Figure 1. Given two consecutive images $(\mathbf{I}_{k-1}, \mathbf{I}_k)$, an encoder outputs a latent state $(\mathbf{z}_k, \dot{\mathbf{z}}_k) \in \mathbb{R}^{2l}$. Latent dynamics are modeled as controlled coupled oscillators: $\mathbf{M}\ddot{\mathbf{z}} + \mathbf{D}\dot{\mathbf{z}} + \mathbf{K}\mathbf{z} = \mathbf{A}\boldsymbol{\tau}$. The learned matrices $\mathbf{M} \succ 0$, $\mathbf{D} \succeq 0$, and $\mathbf{K} \succ 0$ define an oscillator network, ensuring stable latent rollouts and enabling direct use of classical linear control tools.

Using the learned latent dynamics, we design LQR-based controllers shown in Figure 2 that operate entirely in the latent space. Both internal-model and pixel-feedback variants achieve robust tracking of desired shape trajectories across simulation and hardware experiments. We validate nLON on planar soft-tentacle robots in high-fidelity Cosserat-rod simulations [3] and on a hardware platform. The model achieves accurate multi-step predictions, consistently outperforming ConvLSTM and Koopman-based baselines in structural similarity and final-step reconstruction error, while using a compact three- to four-dimensional latent state.

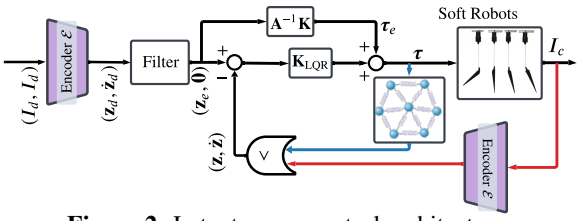

Figure 2: Latent-space control architecture.

These results show that embedding mechanical structure into learned latent dynamics provides a promising direction for vision-based modeling and control of soft and other high-dimensional robotic systems. More experimental details and visual results will be presented during the talk.

References

[1] T. Kapoor *et al.*, "Neural oscillators for generalizing parametric PDEs," in The Symbiosis of Deep Learning and Differential Equations III, 2023

[2] A. Ceni *et al.*, "Random oscillators network for time series processing." In International Conference on Artificial Intelligence and Statistics, pp. 4807-4815. PMLR, 2024.

[3] M. Gazzola *et al.*, "Forward and inverse problems in the mechanics of soft filaments." Royal Society open science 5, no. 6 (2018): 171628.

Bimodal Atomic Force Microscopy with Nonlinear Control for High-Speed Metrology

Jonas Hendrikx* Erik Steur* Robbert Voorhoeve† Henk Nijmeijer* Nathan van de Wouw*
Email: j.g.hendrikx@tue.nl

1 Introduction

As lateral transistor scaling reaches physical limits, the semiconductor industry is shifting toward 3D architectures such as FinFETs, GAAFETs, and stacked dies. Throughout the fabrication process of these complex structures, high-resolution, non-destructive metrology is essential to monitor critical dimensions and ensure process control. Atomic Force Microscopy (AFM) is a 3D imaging technique that achieves sub-nanometer accuracy, making it a key tool for next-generation devices. Due to its non-destructive operation and high throughput, amplitude-modulated AFM (AM-AFM) is widely employed in semiconductor metrology. In AM-AFM, the cantilever is driven near its flexural resonance, and a feedback loop aims to maintain the flexural oscillation at a predefined amplitude setpoint. As the tip scans across the surface, the sample topography acts as a disturbance that shifts the amplitude away from this setpoint. The controller then rejects this disturbance by moving the cantilever in the z-direction to restore the setpoint amplitude. Despite the advantages of AM-AFM, high-speed imaging of steep topographies remains challenging, as the flexural mode exhibits a slow transient response. This leads to tracking errors, ultimately limiting the achievable bandwidth and requiring a tradeoff between imaging accuracy and speed.

2 Bimodal AFM feedback

To overcome the bandwidth limitation of conventional AM-AFM, we propose a bimodal AFM approach that simultaneously excites the flexural and torsional resonance modes of the cantilever. The torsional mode, owing to its much higher resonance frequency, provides a faster dynamic response and enables earlier detection of steep topographic features. However, due to its limited vertical sensitivity, the torsional mode cannot serve as the sole feedback signal. In conventional AM-AFM, tip–sample separation is regulated by a PI controller acting on the flexural amplitude. Building on this structure, our approach exploits the fast torsional response by dynamically adjusting the integral gain of the flexural feedback using the torsional amplitude signal. To eliminate steady-state components, the torsional amplitude is first filtered by a standard second-order high-pass filter, and the integral gain is then formulated as

$$K_i(t) \;=\; K_{i,0} + K_{\text{tors}}\,\tilde{A}_{\text{tors}}(t). \tag{1}$$

In which $K_{i,0}$ is the nominal integral gain, K_{tors} is a tuning parameter, and $\tilde{A}_{\text{tors}}(s) = H_{\text{hp}}(s)A_{\text{tors}}(s)$ is the torsional amplitude after high-pass filtering. The resulting control law becomes

$$u_z(t) \;=\; K_p\,e(t) \;+\; K_i(t)\int_0^t e(\tau)\,\mathrm{d}\tau, \tag{2}$$

were $u_z(t)$ is the stage height, $e(t)$ is the flexural amplitude error, and K_p is the proportional gain. This adaptive integral action enables faster response during rapid height changes while preventing the overshoot that limits imaging accuracy in conventional PI control. We evaluate our bimodal control approach on a simulation model proposed by [1].

3 Results and Conclusions

We compare conventional flexural PI control and bimodal excitation with nonlinear adaptive gain PI control (Nonlinear MISO PI). The results in Fig. 1 show that the proposed nonlinear MISO PI controller achieves a reduction in overshoot compared to conventional flexural PI control. These results demonstrate the potential of bimodal AFM with nonlinear control to enhance AFM scanning performance.

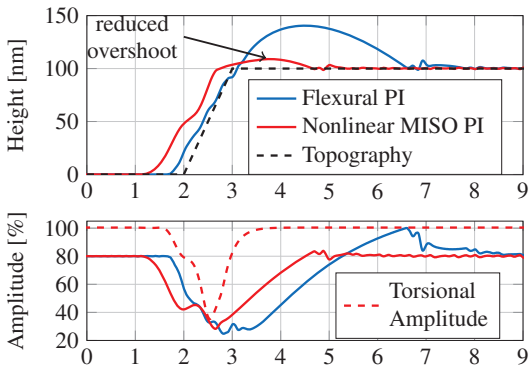

Figure 1: Simulated AFM scan at 10 µm/s comparing conventional flexural PI control (blue) with bimodal nonlinear MISO PI control (red). Flexural amplitudes are shown as solid lines, while the torsional amplitude in the bimodal case is shown as a dashed red line.

References

[1] Y. Song and B. Bhushan, "Simulation of dynamic modes of atomic force microscopy using a 3D finite element model," *Ultramicroscopy*, vol. 106, no. 8–9, pp. 847–873, 2006.

[2] J. G. Hendrikx, E. Steur, R. J. Voorhoeve, H. Nijmeijer, and N. van de Wouw, "Detecting topography changes with combined torsional and flexural resonance mode atomic force microscopy," in *Proc. SPIE Advanced Lithography and Patterning Conf.*, accepted for publication, 2026.

*Eindhoven University of Technology, Department of Mechanical Engineering, Dynamics & Control Group
†TNO, Optomechatronics Department, Delft
Nanotube project is funded by Holland High Tech — TKI HSTM via the PPP Innovation Scheme (PPP-I) for public-private partnerships.

Moment Matching by Recursive-Based Learning[1]

Francesco Corrini[1], Matteo Scandella[1], Alessio Moreschini[2]
[1]Dept. of Management, Information and Production Engineering University of Bergamo, Dalmine (BG), Italy
[2]Dept. of Electrical and Electronic Engineering, Imperial College London, SW72AZ London, U.K
Corresponding author e-mail: francesco.corrini@unibg.it

1 Introduction

Modeling dynamical systems is a fundamental task in the control field. The advancements of the industrial technologies demand for more complex models, which have to account for both high dimensionality and heavy nonlinearities of the systems. In practical applications, large models may represent a problem where the computational capacity is limited. Model reduction techniques have been proposed to approximate high-order systems with lower-dimensional models [1]. Moment matching methods rely on the *moments* of the system, which for LTI systems are linked to the coefficients of the Laurent series of the transfer function around a point. For nonlinear systems [2], moments are related to a steady-state output response that results from an invariant manifold, which has been proved to be a center manifold when the system is interconnected with an input generator with specific properties [2]. This manifold is characterized by a nonlinear mapping, which can be obtained by solving a PDE, which is known as the *invariance equation*. However, finding a solution to such a PDE may be impractical, since the complexity increases with model dimension and the high-order system may also be unknown. To address these challenges, data-driven methodologies have been developed, motivated by the increasing availability of input-output data from industrial processes. Recently, in [3], the moment is estimated by employing Tikhonov regularization on a reproducing kernel Hilbert space. While this method successfully achieves moment matching for high-order nonlinear systems, its complexity scales with the number of data points, which is necessary to increase the accuracy of the estimation. Furthermore, the accuracy decreases when data presents a transient. In this work, we present a recursive procedure that (i) updates the moment estimation using data collected online from the system and (ii) discards the oldest data point used for estimation. In this way, the number of data used for estimation is constant, thus the complexity does not increase. However, we prove that the accuracy of the estimation increases, since the effect of the transient vanishes asymptotically in time.

2 Preliminaries

In [3], the authors provide the kernel-based learning method to estimate the moment. In particular, the following theorem provides a probabilistic bound on the estimation error.

Theorem 1 ([3], Theorem 4) *Given the assumptions in [3], the reduced model obtained with the kernel-based method achieves ϑ-weak moment matching with probability $1-\varepsilon$, where $\varepsilon \in (0,1]$ and*

$$\vartheta := 2\rho|\tilde{\mu}^\star|_{\mathscr{L}_2} + \frac{\zeta\sqrt{T_n}}{n\sqrt{\rho-\delta_n}} + \frac{2\zeta\left(\frac{\Psi}{n}+\frac{\Sigma}{\sqrt{n}}\right)}{\sqrt{\rho-\delta_n}}\log\left(\frac{2}{\varepsilon}\right).$$

The meaning of Theorem 1 is that the estimation error can be decomposed into three terms: the first is introduced by the Tikhonov regularization and increases as the regularization parameter ρ increases; the second depends on the effect of the transient described by T_n; the third one depends on the uncertainty introduced by the noise.

3 Main contribution

The main contribution of this work is a recursive procedure that efficiently updates the moment estimation online. The proposed algorithm updates the estimation with the most recently collected datum while discarding the earliest collected datum to keep the size of the data set constant. Further, we update the regularization parameter ρ by minimizing the probabilistic bound provided by Theorem 1. Finally, the following theorem provides a new bound on the estimation error.

Theorem 2 *Suppose that the hypotheses of Theorem 1 hold and that the data are collected with constant sampling time. Then, for $t \to \infty$, the reduced model obtained with the recursive kernel-based method mentioned above achieves θ-weak moment matching with probability $1-\varepsilon$, where $\varepsilon \in (0,1]$ and*

$$\vartheta := 2\rho|\tilde{\mu}^\star|_{\mathscr{L}_2} + \frac{2\zeta\left(\frac{\Psi}{n}+\frac{\Sigma}{\sqrt{n}}\right)}{\sqrt{\rho-\delta_n}}\log\left(\frac{2}{\varepsilon}\right).$$

The proof follows from the fact that, discarding the earliest data points and substituting them with more recent data points collected online, the effect of the transient T_n vanishes, so that the second term in Theorem 1 goes to zero. In this way, it is possible to increase the accuracy of moment estimation while keeping the complexity of the estimation procedure constant.

References

[1] A. Antoulas, *Approximation of Large-Scale Dynamical Systems*, SIAM, 2005.
[2] A. Astolfi, "Model Reduction by Moment Matching for Linear and Nonlinear Systems," *IEEE Trans. Autom. Control*, vol. 55, no. 10, 2010.
[3] A. Moreschini, M. Scandella et al., "Moment Matching by Kernel-Based Learning," *IEEE Trans. Autom. Control*, (early access), 2025.

[1]This work was partially funded by the EU European Defence Fund Project KOIOS (EDF-2021-DIGIT-R-FL-KOIOS) and it was supported by the EPSRC grant EP/X033546/1. KOIOS is a project Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Defence Fund. Neither the European Union nor the granting authority can be held responsible for them.

Demographic Dependence of Vaccine Adoption under Opinion Persuasion

Alessandro Casu[1] Camilla Quaresmini[2] Robin Delabays[3] Lewis Mitchell[4] Philip E. Paré[5]
[1]Control Systems Group, Eindhoven University of Technology, NL [2]Politecnico di Milano, IT
[3]University of Applied Sciences of Western Switzerland, CH [4]Adelaide University, AU [5]Purdue University, US
Email: a.casu@tue.nl

1 Introduction

Vaccination is one of the most effective tools for reducing the burden of infectious diseases, yet hesitancy continues to limit uptake. We propose a networked model that couples epidemic spread, vaccination uptake with permanent immunity, and signed opinion dynamics about vaccines across interacting demographic subpopulations (1). Vaccination adoption is constrained by an opinion-dependent carrying capacity, while infection prevalence feeds back onto beliefs (2). Public health messaging enters as a control input whose perceived effect varies across demographics, enabling targeted communication. We characterize disease-free and endemic equilibria and provide a sufficient local stability condition for the healthy state. Simulations show that demographically aligned messaging increases vaccine uptake and reduces endemic infection compared to misaligned campaigns.

2 Model

We consider n interacting subpopulations. Each node i has infected share $x_i \in [0,1]$, vaccinated share $v_i \in [0,1]$, and opinion o_i, with $o_i > 0$ indicating pro-vaccine sentiment. Let $B = [\beta_{ij}] \geq 0$ denote the contact network and $A = [\alpha_{ij}]$ the signed opinion network.

Epidemic layer (SIS-V). Susceptible individuals become infected through contacts, recover at rate γ_i, and may vaccinate if capacity allows:

$$\dot{x}_i = (1 - x_i - v_i) \sum_j \beta_{ij} x_j - \gamma_i x_i, \tag{1}$$

$$\dot{v}_i = \rho \Big(1 - \frac{v_i}{\kappa_i(o_i)}\Big)(1 - x_i - v_i). \tag{2}$$

Vaccination slows as v_i approaches the opinion-dependent capacity $\kappa_i(o_i)$.

Opinion layer (signed influence and persuasion). Each group is characterized by a demographic vector $\boldsymbol{d}_i \in \mathbb{R}^k$, and influence weights are given by $\alpha_{ij} = \boldsymbol{d}_i^\top \boldsymbol{d}_j$. Opinions evolve through peer influence, infection awareness, and a vaccine promotion message $\boldsymbol{u} \in \mathbb{R}^k$:

$$\dot{o}_i = (x_i - o_i - \tfrac{1}{2}) + \sum_j \big(\alpha_{ij} o_j - |\alpha_{ij}| o_i\big) + \boldsymbol{u}^\top \boldsymbol{d}_i. \tag{3}$$

The persuasion term can shift opinions away from their nominal interval, although the resulting opinion dynamics remain bounded. Higher infection prevalence increases pro-vaccine sentiment, while persuasion (or distrust when the dot product is negative) strength depends on demographic alignment.

Opinion-to-uptake coupling. Vaccine capacity increases with opinion and is constrained so that changing opinions cannot reduce an already achieved vaccination level:

$$\kappa_i(o_i) = \max\left(v_i, \frac{1 + 2o_i}{1 - 2o_i}\right). \tag{4}$$

3 Equilibria and healthy-state stability

At equilibrium, vaccination satisfies $v_i^* = \min(1, \kappa_i^*)$, while infection levels solve the standard SIS fixed-point equation. The steady-state opinion matrix is

$$\boldsymbol{o}^* = (\tilde{A} + I)^{-1} \left(D^\top \boldsymbol{u} - \tfrac{1}{2}\mathbf{1} + \boldsymbol{x}^*\right), \tag{5}$$

where $D = [\boldsymbol{d}_1 \cdots \boldsymbol{d}_n]$ and $\tilde{A}$ is the signed Laplacian with $\tilde{A}_{ii} = \sum_j |\alpha_{ij}|$ and $\tilde{A}_{ij} = -\alpha_{ij}$ for $i \neq j$.

Healthy state. The disease-free equilibrium has $\boldsymbol{x}^* = \mathbf{0}$ and vaccination levels determined by $\boldsymbol{o}^*$. A sufficient condition for local stability of the healthy epidemic state is that

$$M = (I - \mathrm{diag}(\boldsymbol{\kappa}^*))B - \mathrm{diag}(\boldsymbol{\gamma})$$

has negative dominant eigenvalue.

4 Implications for intervention design

The demographic dependence of persuasion implies that message framing directly shapes vaccination capacity through opinion dynamics. Choosing $\boldsymbol{u}$ to align with population demographics increases uptake and suppresses infection, while anti-aligned messaging can sustain endemic disease in negatively opinionated groups. This motivates selecting $\boldsymbol{u}$ to minimize endemic burden, e.g., $\|\mathrm{diag}(N_1, \ldots, N_n)\boldsymbol{x}^*\|_1$ (with N_i being the number of individuals composing subpopulation i), subject to the coupled equilibrium conditions.

References

[1] B. She, J. Liu, S. Sundaram, and P. E. Paré, "On a Networked SIS Epidemic Model With Cooperative and Antagonistic Opinion Dynamics," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 3, pp. 1154–1165, 2022.

[2] H. C. H. Leung, M. E. Gibbs, and P. E. Paré, "The Impact of Vaccine Hesitancy on Epidemic Spreading," in *Proc. American Control Conference (ACC)*, 2022, pp. 3656–3661.

Optimal controllers in stochastic abstractions and sequential games

Adrien Janssens[1], Raphaël M. Jungers
UCLouvain, ICTEAM

1 Introduction

In safety-critical applications, such as autonomous vehicles, complex systems must be verified against rich specifications, which include notions of safety and liveness, e.g. never hit an obstacle and eventually reach a certain set. This has motivated the use of formal methods for the analysis and control of dynamical systems.

In this context, *Abstraction-based control* techniques involve replacing the original (concrete) control problem with an abstract, finite control problem that is easier to analyze and synthesize a controller for, while preserving behavioral properties. This abstract problem is typically obtained by partitioning the concrete state space, where each cell of this partition corresponds to one abstract state.

Abstraction-based control methods have recently been extended to the stochastic setting, following the same steps as in the deterministic setting but with the difference that their abstraction must capture the probabilistic behavior. The abstract system then comes in the form of a Markov Decision Process (MDP) [1] or, more recently, an Interval Markov Decision Process (IMDP) [2].

IMDPs had previously been introduced [3] as a generalization of standard MDPs. Informally, they allow the transition probabilities to lie in intervals, to model uncertainty about the system's dynamics. A key drawback of MDP abstractions is the introduction of a discretization error relative to the original stochastic system, which accumulates over time. Intuitively, each continuous state inside the cell leads to a different transition kernel, but we now consider a single representative kernel for the whole cell.

IMDP abstractions address this by transforming uncertainty on the exact position of the system inside a partition cell into uncertainty on the transition probabilities, trading error with additional conservatism.

2 IMDPs and sequential games

To formally define a stochastic process from an IMDP, one has to specify (see Figure 1):

- A policy (control law) π, if the system is controlled.
- An adversary ξ, whose role is to select the transition kernel.

[1]Corresponding author: adrien.janssens@uclouvain.be

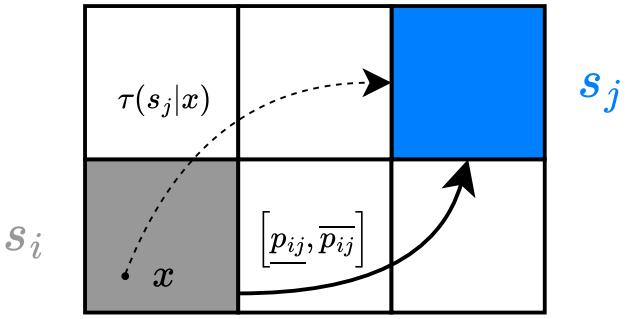

Figure 1: Illustration of an IMDP abstraction, with no actions. The state space is partitioned and the transition probabilities are intervals. An adversary selects the transition kernel that will be applied, which may not correspond to the exact position of the system (illustrated with a single transition $i \to j$).

This leads to a *sequential game*. In the context of abstraction-based control, the adversary does not exist: it is merely a model of the epistemic uncertainty introduced by the abstraction procedure. Because we want to be as robust as possible, the adversary selects the worst-case transition probability. In practice, we find the MDP compatible with the interval constraints which minimizes the satisfaction probability. The transition probabilities are thus fixed, whereas the effective transition kernel will change every time the system enters the cell.

We study the question *whether* and *when* a stationary adversary is sufficient. In particular, we ask whether the worst-case behavior can always be captured by a fixed choice of transition probabilities.

References

[1] A. Abate, J. P. Katoen, J. Lygeros, and M. Prandini, "Approximate model checking of stochastic hybrid systems," *European Journal of Control*, 2010.

[2] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis of discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, 2015.

[3] A. Nilim and L. El Ghaoui, "Robust control of markov decision processes with uncertain transition matrices," *Operations Research*, 2005.

Consolidating Probabilistic Optimal Control

A. Bhole, M. M. Filabadi, G. Crevecoeur, T. Lefebvre
Department of Electromechanical Systems and Metal Engineering, Ghent University, Ghent, Belgium
Core lab MIRO, Flanders Make, Belgium
Email: {ajinkya.bhole, mohammad.mahmoudifilabadi, guillaume.crevecoeur, tom.lefebvre}@ugent.be

Over the past two decades, growing interplay between control engineering and certain areas of machine learning has elevated the role of probabilistic policies in control design. Randomness, once confined to disturbances and uncertainty, is now deliberately employed to promote exploration, robustness, and generalization, with entropy regularization emerging as a principled means to balance exploration and exploitation. This paradigm shift also motivated the search for control formulations in which the optimal or posterior policy admits a expectation-based representation with respect to a prior policy, enabling tractable solutions.

The field of (stochastic) optimal control today, has thus diversified into a rich mosaic of distinct modeling paradigms, each employing its own mathematical tools and conceptual frameworks. *Probabilistic Control Design* frames decision-making as a density-matching problem using KL divergence; *Control as Inference* embeds optimality variables in a probabilistic graphical model; *Risk-Sensitive Control* accounts for the underlying cost distribution beyond the mean; *Distributionally Robust Control* uses ambiguity sets to hedge against model mismatch; *KL-Regularized Control* adds entropy penalties to improve tractability; while *Linearly Solvable Optimal Control* exploits linear Bellman equations under specific structural assumptions. While these approaches often address closely related questions, their differing formulations and solution techniques have led to a fragmented landscape in which connections are obscured and comparisons are nontrivial. We consolidate these diverse perspectives through the common lens of regularization, providing a consistent and self-contained framework that clarifies how modeling assumptions translate into different control formulations.

The unifying framework is anchored in a central control problem that goes beyond standard trajectory-level regularization, by explicitly optimizing over both the control policy and an auxiliary system transition kernel, each penalized by a divergence from a baseline, and weighted independently. This separation makes it possible to distinguish between uncertainty induced by stochastic control actions and that attributed to deviations in system dynamics, thereby providing a principled mechanism to encode and understand robustness and risk sensitivity within a single formulation.

The resulting problem admits a representation of the form

$$\min_{\underline{\pi}} \operatorname*{opt}_{\underline{\tau}}^{\lambda^S} \Big\{ \mathbb{E}_{p_{(\underline{\pi},\underline{\tau})}}[\underline{c}_T] + \frac{1}{\lambda^P} \mathbb{D}_{\underline{\rho}}^{\underline{\pi}} + \frac{1}{\lambda^S} \mathbb{D}_{\underline{\iota}}^{\underline{\tau}} \Big\}, \tag{1}$$

where $\underline{\pi} = (\pi_0, \ldots, \pi_{T-1})$ is the policy sequence, $\underline{\tau} = (\tau_0, \ldots, \tau_{T-1})$ is an auxiliary transition kernel sequence, $p_{(\underline{\pi},\underline{\tau})}$ is the resulting trajectory distribution, and $\underline{c}_T$ is the cumulative cost. The baseline policy $\underline{\rho}$ and baseline dynamics $\underline{\iota}$ provide reference behaviors. The operator $\operatorname{opt}^{\lambda^S}$ denotes minimization when $\lambda^S > 0$ (risk-seeking) and maximization when $\lambda^S < 0$ (risk-averse). The weights $\lambda^P > 0$ and $\lambda^S \neq 0$ independently control the strength of regularization toward the policy and transition baselines, respectively. Serving as an expressive umbrella, (1) provides a systematic way to represent, relate, and compare several classical stochastic control formulations by toggling two structural choices: (i) Whether the control policy is regularized against a baseline reference $\underline{\rho}$ (policy regularization on/off). (ii) Whether the transition kernel sequence is fixed to the true dynamics $\underline{\iota}$ or left free as an optimization variable (transition optimization on/off). Through these toggles we recover:

- **Stochastic Optimal Control (SOC)**: No policy regularization by setting $\underline{\rho} = \underline{\pi}$ and constraining transitions $\underline{\tau} = \underline{\iota}$.
- **Risk-Sensitive Optimal Control (RSOC)**: No policy regularization, while leaving $\underline{\tau}$ free.
- **Soft-Policy SOC (SP-SOC)**: Policy $\underline{\pi}$ regularized w.r.t. baseline $\underline{\rho}$ and fixing $\underline{\tau} = \underline{\iota}$.
- **Soft-Policy RSOC (SP-RSOC)**: Allowing both $\underline{\pi}$ and $\underline{\tau}$ to be regularized and optimized.

Using the Majorization–Minimization (MM) framework it can be shown that the regularized problems SP-SOC and SP-RSOC majorize the original SOC and RSOC objectives, respectively, providing tractable surrogates for iterative algorithm design. When the baseline dynamics are deterministic, the auxiliary transition kernel loses its expressive freedom, and we have: SOC $\equiv$ RSOC, and SP-SOC $\equiv$ SP-RSOC.

For SP-RSOC, the special case of synchronized regularization, where the policy and transition regularization weights coincide ($\lambda^P = \lambda^S = \lambda > 0$), enjoys powerful properties: a linear Bellman equation, a path integral solution, and compositionality that enables modular control design. These properties, previously known only in restricted settings, are thus extended to a broader class.

These insights open up several promising research avenues. Our future work will explore time-varying regularization weights to adapt risk attitudes throughout the horizon, shift from soft to hard constraints, and extend the framework to partially observable settings. Practical algorithmic development could also build on these connections to design efficient solvers for a wide range of optimal control applications.

References

[1] A. Bhole, M. M. Filabadi, G. Crevecoeur, T. Lefebvre, "Unifying Entropy Regularization in Optimal Control: From and Back to Classical Objectives via Iterated Soft Policies and Path Integral Solutions," arXiv:2512.06109, 2025.

Accurate Remnant Positioning Through Combined Hysteresis and Creep Modeling

J. R. Keulen[1,*], K. Classens[2], M. Acuautla[1], B. Jayawardhana[1]
[1]University of Groningen, The Netherlands
[2]ASML, The Netherlands
*Email: j.r.keulen@rug.nl

Abstract

'Set-and-forget' or shape memory piezoelectric actuators enable high-precision, long term positioning without continuous power supply or feedback control. After the actuator is driven to a desired position it maintains this position, which is defined as the remnant position. This makes the set-and-forget actuator ideal for high-precision applications, such as space research, advanced optical systems, and lithography, where ultra-stable, drift-free positioning is required over extended periods. An additional advantage is that a single set of electronics can drive multiple actuators sequentially. The working principle of the set-and-forget piezo actuator relies on the hysteretic memory of the strain–electric field curve, also known as the butterfly loop. This loop is typically symmetric in conventional piezoelectric ceramics such as PZT (lead zirconate titanate). Strategic defects in the material introduced by niobium doping yields asymmetry in the butterfly loop [1]. This enables a wider range of reachable remnant positions.

To exploit this hysteretic memory, iterative feedback algorithms have been developed that recursively adjust the input amplitude to drive the actuator to a desired remnant position [2]. However, once this position has been reached, the actuator still exhibits creep, a slow, time-dependent deformation that changes the remnant position. Accurately predicting or compensating for this drift requires a dynamic model of the underlying relaxation processes inside the material. Following [3], the creep behavior is represented as a combination of fast and slow visco-elastic dynamics using one elastic element K_0 and n Kelvin–Voigt (KV) elements, each consisting of a linear spring K_i in parallel with a damper D_i, as illustrated in Figure 1. These parameters are identified from experimental measurements of a PNZT pellet. While the KV model captures the time-dependent relaxation, the static hysteresis that defines the static remnant position is described using the Preisach model [4]. Combining both components yields a model where the Preisach operator captures the instantaneous hysteresis, while the KV model describes the slow creep, as shown for a 1800 V excitation in Figure 2. This combined modeling framework enables explicit compensation of creep, thereby improving accuracy of the remnant positioning.

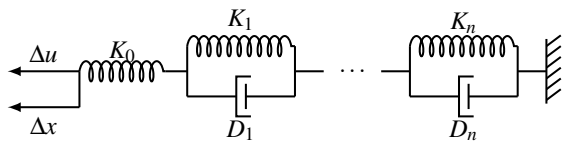

Figure 1: Visco-elastic creep model based on Kelvin Voigt elements. This can be expressed the Laplace domain by $C(s) = \frac{1}{K_0} + \sum_{i=1}^{n} \frac{1}{D_i s + K_i}$, where $C : \Delta u \mapsto \Delta x$, where Δu is the voltage and Δx the displacement.

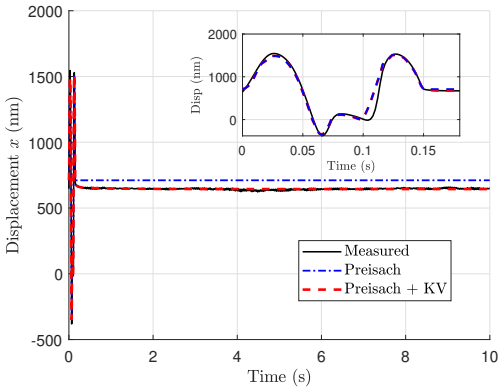

Figure 2: Measured PNZT actuator response plotted with the Preisach hysteresis model and the combined Preisach–KV model. Incorporating the KV creep elements reduces the relative modeling error from 10.27% to 2.42%.

References

[1] E. van der Veer, B. Noheda, and M. Acuautla, "Piezoelectric properties of pzt by an ethylene glycol-based chemical solution synthesis," *Journal of Sol-Gel Science and Technology*, vol. 100, no. 3, pp. 517–525, 2021.

[2] J. R. Keulen, M. Ugalde-Reygadas, N. K. Hamelynck, J. Roorda, M. Acuautla, and B. Jayawardhana, "Development and validation of remnant control for pnzt hysteretic piezo actuators." In preparation, 2025.

[3] K. Seki, J.-Y. Hirata, M. Ruderman, and M. Iwasaki, "Creep modeling with time-dependent damping parameters in piezoelectric actuators," in *2019 IEEE International Conference on Mechatronics (ICM)*, vol. 1, pp. 135–140, IEEE, 2019.

[4] M. A. Vasquez-Beltran, B. Jayawardhana, and R. Peletier, "On the characterization of butterfly and multiloop hysteresis behavior," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3494–3506, 2021.

Optimal setpoint adaptation under regime change for control of mechatronic systems

Foeke Vanbecelaere, Miel Van der Wilt, and Taranjitsingh Singh
MotionS, Flanders Make VZW
3920 Lommel, Belgium
foeke.vanbecelaere@flandersmake.be

1 Motivation

Mechatronic motion systems struggle when operating under varying load conditions or regimes, particularly when relying on fixed, precomputed motion profiles or setpoints. Adapting control parameters or adding sensors to handle regime changes is usually impractical due to cost, complexity, access to controllers, or reliability concerns.

A more practical solution is to adapt motion profiles using optimal control, allowing the system to respond to observed performance rather than direct load measurements. This is particularly challenging for pneumatic systems due to nonlinear dynamics, actuator limits, dead zones, and sensitivity to fluctuations in supply pressure and leaks. The objective of this work is to validate a dual-control approach for real-time adaptation of setpoints during regime change.

2 Approach

The considered pneumatic system is modeled as a pressure-dependent transfer function, preceded by actuator dead-band and saturation. The system performs repetitive motions with predefined timing, where the motion profiles are parameterized using Chebyshev polynomials [1] to enable numerically efficient optimization while respecting motion constraints at the rest positions. For both forward (FW) and backward (BW) motion, an optimal control problem is solved using CasADi. The RMS value of the control effort u_{rms} is minimized over a range of supply pressure p_{s} values. A hard constraint is imposed on the position error e_{p} at the rest positions. The motion profile's coefficients are combined using principal component analysis, resulting in a scalar value z. Its relation with the varying pressure p_{s} and control effort u_{rms} is mapped in figure 1 and is implemented on-line. Note that the optimum does not always correspond with the minimum u_{rms} because of the position constraint at rest.

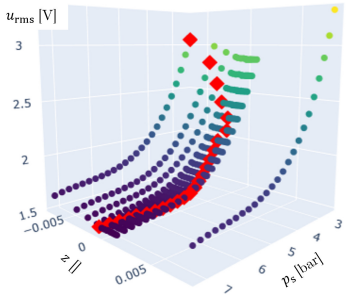

Figure 1: Control effort u_{rms} in function of motion profile z and pressure p_{s}, with optimum (♦) indicated.

During on-line operation, a supervisory layer keeps track of the control effort and adapts the motion profile. Changes in u_{rms} are interpreted as changes in p_{s}. Using the stored map, the optimal profile z for the next FW or BW movement is generated by interpolation at the estimated pressure value.

3 Results

A drop in supply pressure is simulated. Figure 2 shows that the dual control strategy remains optimal under a changing regime. It is 13% more efficient than a nominal-pressure profile, maintains accurate positioning to ensure productivity, and adds no computational burden since optimization is performed off-line.

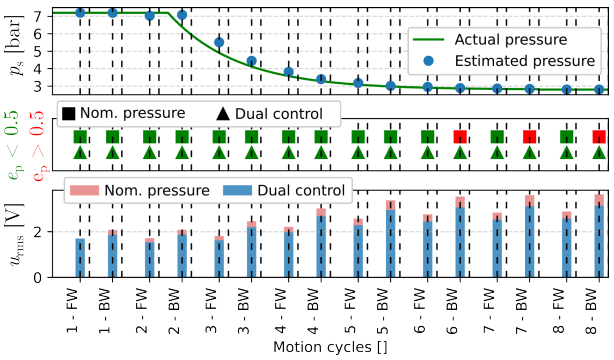

Figure 2: Performance over time. Top to bottom: tracking the supply pressure, position constraint, and control effort.

References

[1] F. Vanbecelaere, N. Van Oosterwyck, S. Derammelaere, A. Cuyt, M. Monte, and K. Stockman, "On-line motion profile optimization for reciprocating mechanisms," *Mechanism and Machine Theory*, vol. 173, p. 104833, 2022.

Acknowledgment This research was supported by Flanders Make, the strategic research center for the manufacturing industry, in the framework of the ADCEE_SBO project.

Funnel control for TCP congestion control with state-dependent input delays

Yanxin Li
Bernoulli Institute, University of Groningen
yanxin.li@rug.nl

Stephan Trenn
Bernoulli Institute, University of Groningen
s.trenn@rug.nl

1 Introduction

Transmission Control Protocol (TCP) congestion control is a key mechanism to ensure the stability and fairness of communication networks. However, the time delays in packet transmission and acknowledgment pose fundamental challenges to stability and performance. For this reason, many TCP network congestion control algorithms have been proposed, such as the RED algorithm. Since these algorithms are highly sensitive to static parameters, they cannot effectively adapt to changing network environments. Motivated by this limitation, Misra et al. [1] developed the following fluid-flow based model of TCP dynamics:

$$\dot{W}(t) = \frac{1}{R(t)} - \frac{W(t)}{2}\frac{W(t-R(t))}{R(t-R(t))}p(t-R(t)),$$
$$\dot{q}(t) = \begin{cases} -C(t) + \dfrac{N(t)}{R(t)}W(t), & q > 0, \\ \max\left\{0, -C(t) + \dfrac{N(t)}{R(t)}W(t)\right\}, & q = 0, \end{cases} \quad (1)$$

where $W(t)$ is the TCP window size at time $t \geq 0$, $q(t)$ is the queue length of the router, $R(t) = T_p(t) + \frac{q(t)}{C(t)}$ is the round-trip time, $C(t)$ is the available link capacity, $N(t)$ is the number of TCP sessions, and $p(t) \in [0,1]$ is the probability of packet loss.

Based on model (1), various control methods have been developed to improve the response speed of network systems, such as PID controllers and sliding mode controllers. These controllers often require precise knowledge of system parameters (e.g., round-trip time, link capacity, gain factors), which are typically unknown or time-varying in real networks. Consequently, maintaining stable and predictable queue dynamics under uncertain and time-delayed network environments remains a challenging and unresolved problem.

Funnel control [2] is an output reference tracking adaptive control approach that guarantees prescribed transient and steady-state performance without requiring an accurate system model. It can rapidly adapt to dynamic network conditions, thereby significantly enhancing the robustness and stability of congestion control. It thus represents a novel and effective approach for controlling network congestion.

2 Funnel control

As a starting point, we ignore the complex window dynamics and assume that the window size changing can be directly controlled via an input u, i.e. the model (1) simplifies to

$$\dot{W}(t) = u(t),$$
$$\dot{q}(t) = -C(t) + \frac{N(t)}{R(t)}W(t). \quad (2)$$

It is easily seen, that this model has relative degree two, hence a classical relative-degree two funnel control can be used to control system (2). The corresponding Matlab simulation results are shown in Figure 1. Clearly, the funnel control works well, and the actual queue length tracks the desired queue length $q_d = 20$.

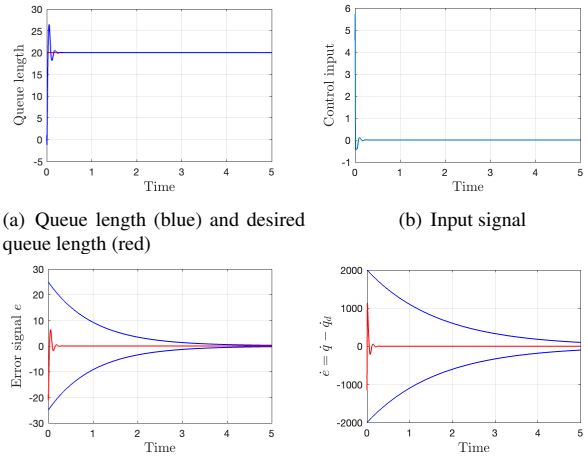

(a) Queue length (blue) and desired queue length (red)
(b) Input signal
(c) Error signal $e = q - q_d$ (red) and funnel boundary (blue)
(d) $\dot{e} = \dot{q} - \dot{q}_d$ (blue) with derivative funnel (blue)

Figure 1: Simulation results of funnel controller for (2).

Preliminary simulations of the classical funnel controller to the original model (1) indicate that the error *cannot* be kept within the funnel without additional assumptions. The key challenge is the presence of time delays in the state and the input, which is further complicated by the fact, that the delay is not constant, but depends on the state.

References

[1] V. Misra, W. Gong, and T. Donald, "Fluid-based analysis of network of AQM routers supporting TCP flows with an application to RED," ACM SIGCOMM Computer Communication Review, vol. 30, no. 4, pp. 151-160, 2000.

[2] Thomas Berger, Achim Ilchmann, and Eugene P. Ryan, "Funnel control — A survey," Annual Reviews in Control, vol. 60, pp. 101024.1-36, 2025.

Probabilistic Safety under Arbitrary Disturbance Distributions using Piecewise-Affine Control Barrier Functions

Matisse Teuwen, Mathijs Schuurmans, Panagiotis Patrinos
KU Leuven, Department of Electrical Engineering (ESAT-STADIUS),
Kasteelpark Arenberg 10, 3001 Leuven, Belgium.[1]
Email: matisse.teuwen@esat.kuleuven.be

1 Introduction

Driven by advances in machine learning and high-fidelity simulation, learning-based control methods have gained popularity, but often lack formal safety guarantees in uncertain environments. Safety filters based on Control Barrier Functions (CBFs) have emerged as a useful tool to enforce these safety constraints. Probabilistic CBFs (pCBFs) account for uncertainty by bounding the probability of safety violations, but existing approaches typically rely on restrictive distributional assumptions or conservative concentration inequalities.

This work aims to construct a probabilistic safety filter for arbitrary disturbance distributions that exploits structural properties to reduce conservatism while preserving modeling flexibility. To this end, we leverage Piecewise-Affine CBFs, allowing us to capture complex safety sets while only requiring quantiles of the disturbance distribution. Supplementary material can be found at
https://mathijssch.github.io/ecc26-supplementary/.

2 Problem statement

We consider a discrete-time stochastic system of the form:

$$x_{k+1} = f(x_k, u_k) + g(x_k, \xi_k),\ k \in \mathbb{N} \tag{1}$$

where $x_k \in \mathbb{R}^n$ is the state, $u_k \in \mathbb{R}^m$ is the control input, and $\xi_k \in \mathbb{R}^p$ is a random disturbance with arbitrary distribution. Safety is defined using the N-step exit probability i.e., the probability that the system remains within a safe set $S \subseteq \mathbb{R}^n$ over a finite horizon N. The objective is to design a control policy that guarantees constraint satisfaction with a user-specified probability.

3 Piecewise-Affine pCBFs

We focus on Piecewise-Affine CBFs of the form:

$$h(x) = \min_{i\in[d]} \max_{j\in[p_i]} c_{ij}^\top x + b_{ij} \tag{2}$$

for which the associated safe set can be expressed as the complement of a finite union of polyhedral sets. This representation yields a highly expressive model class capable of capturing complex, nonconvex safety constraints. We derive sufficient conditions ensuring that h is a pCBF with a guaranteed upper bound on the N-step exit probability. The key insight is to exploit quantiles of affine transformations of the disturbance leading to less conservative bounds. Based on this formulation, we design a safety filter that involves solving a small mixed-integer program at every time step. We propose a simple heuristic that allows us to solve these problems within a few milliseconds.

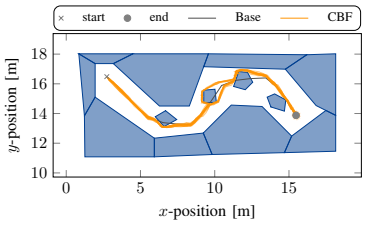

Figure 1: Path-planning experiment in a cluttered environment.

In practical scenarios, the disturbance distribution is often unknown and must be estimated from data. We propose a data-driven approach to estimate the required quantiles using order statistics and provide probabilistic safety guarantees with high confidence.

4 Numerical Experiments

We validate the proposed approach on stochastic navigation and path planning tasks. Compared to existing methods [1, 2], our approach gives tighter bounds on the safety violation probability and larger feasible regions under higher disturbance levels. Furthermore, we empirically demonstrate that our approach is both less conservative and less computationally demanding than other data-driven methods [3]. We further illustrate the flexibility of the method by considering complex safe sets with multiple obstacles and by evaluating performance under different disturbance distributions.

[1] This research has received funding from the European Union under the Horizon Europe Grant Agreement AIGGREGATE, no. 101202457. The work of M. Schuurmans was supported by an FWO postdoctoral fellowship under project number 1248626N.

References

[1] R. Cosner, P. Culbertson, A. Taylor, and A. Ames, "Robust safety under stochastic uncertainty with discrete-time control barrier functions," in *Robotics: Science and Systems XIX*, 2023.

[2] S. Fushimi, K. Hoshino, and Y. Nishimura, "Safety-critical control for discrete-time stochastic systems with flexible safe bounds using affine and quadratic control barrier functions." arXiv:2510.01501, 2025.

[3] P. Mestres, B. Werner, R. K. Cosner, and A. D. Ames, "Probabilistic control barrier functions: Safety in probability for discrete-time stochastic systems." arXiv:2510.01501, 2025.

Constrained Performance Boosting Control for Nonlinear Systems via ADMM

Gianluca Giacomelli[1], Danilo Saccani[2], Siep Weiland[1], Giancarlo Ferrari-Trecate[2], Valentina Breschi[1]
[1]Control Systems Group, Eindhoven University of Technology, 5600MB Eindhoven, The Netherlands
[2]Institute of Mechanical Engineering, EPFL, Lausanne, CH-1015, Switzerland
Email: g.giacomelli@tue.nl

1 Introduction

The increasing demand for high-performing controllers calls for the design of control architectures that push system performance to its limit while ensuring closed-loop stability and guaranteeing constraint specifications.
Within this challenging landscape, we propose ADMM-PB, a learning procedure that extends the Performance Boosting (PB) approach presented in [1] using the Alternating Direction Method of Multipliers (ADMM) [2] to design a state-feedback policy boosting the performance of a nonlinear system while achieving closed-loop stability by design and promoting constraint satisfaction.

2 Setting & goal

Consider an ℓ_p-stable nonlinear system, i.e., the system state's norm is bounded for input and disturbance with bounded norm, with a set of user-defined input and state constraints and a boosting objective. By following [1], we adopt an Internal Model Control (IMC) scheme featuring a learnable $\mathscr{L}_p$ operator, allowing us to decouple the boosting objective from the stability requirements, as stability is preserved independently from its parameters. In practice, this operator is parametrized by a stable NN, trained to minimize a boosting loss over a control horizon following the closed-loop system dynamics. However, including the hard input and state constraints prevents adopting the Gradient Descent (GD) training procedure of [1] due to computational intractability [3]. Prompted by boosting the system performance under hard constraint specifications, our goals are to (i) preserve stability, (ii) maintain the same control structure proposed by [1], and (iii) satisfy a set of user-defined input and state constraints.

3 ADMM for Constrained Performance Boosting

To retain the main features (and advantages) of the seminal PB approach proposed in [1], while not altering the structure of the control law, we propose an ADMM-based scheme for constrained Performance Boosting (ADMM-PB), following three iterative steps. The first step requires solving the same unconstrained training problem of [1], yet with a loss augmented with a soft penalty on the closed-loop input and state constraint violations, thereby allowing us to find a solution with GD. This soft penalty is iteratively updated, as it is formed by elements computed in the two following steps, namely, the safe inputs and states along with the Lagrangian multipliers. The safe inputs and states are determined within the second step by projecting the inputs and states returned by the unconstrained training onto the input and state constraint sets, respectively. In the last step, the Lagrangian multipliers are obtained by recursively adding the difference between the inputs and states of the unconstrained training and the projected ones. The ADMM-PB routine is concluded upon satisfying a pre-defined termination criterion, providing the NN parameters for the PB scheme of [1].

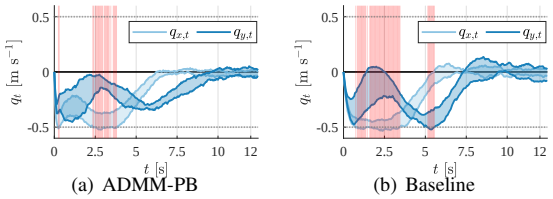

(a) ADMM-PB (b) Baseline

Figure 1: ADMM-PB *vs* CBF-based baseline velocities.

4 Constrained control of a point-mass robot

To evaluate the performance of ADMM-PB, we consider a pre-stabilized point mass robot having the origin of the 2D plane as its equilibrium. As desired boosting goals, we set faster convergence to the equilibrium while avoiding an obstacle, under constraints on the robot's Euclidean velocities $q_{x,t}$ and $q_{y,t}$, respectively. In this setting, we compare the proposed strategy with the literature baseline used in [1], which resorts to augmenting the loss with soft Control Barrier Functions (CBFs) components penalizing constraint violations. As showcased by Figure 1, where we compare the speed range of the robot, i.e., the minimum and maximum speed achieved across every testing scenarios, the baseline approach pushes performance at the cost of a higher constraint violation, while the speed trajectories resulting from the use of ADMM-PB tend to be more cautious, reducing constraints' violation at the price of more conservative performance.

References

[1] L. Furieri, C.L. Galimberti, and G. Ferrari-Trecate, "Learning to Boost the Performance of Stable Nonlinear Systems," in IEEE Open Journal of Control Systems, 2024.
[2] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," in Foundations and Trends in Machine learning, 2011.
[3] P. Márquez-Neila, M. Salzmann, and P. Fua, "Imposing hard constraints on deep networks: Promises and limitations," in arXiv preprint arXiv:1706.02025, 2017.

Compositional design for time-varying and nonlinear coordination

Jonas Hansson
ICTEAM Institute, UCLouvain, Louvain-la-Neuve, Belgium
Email: jonas.hansson@uclouvain.be

Introduction

The coordination of networked dynamical systems is a central topic in control, with consensus protocols providing a canonical solution for first-order agents. While linear first-order consensus is well understood, the design of *high-order* coordination protocols remains challenging under locality constraints. Classical approaches typically consider networks of nth-order integrators $x^{(n)} = u$, and apply the graph Laplacian to multiple derivatives of the state. A representative example is the protocol proposed in [1],

$$x^{(n)} = u = -r_{n-1}Lx^{(n-1)} - \cdots - r_1 L\dot{x} - r_0 Lx, \qquad (1)$$

where the gains r_k must be carefully tuned. However, such designs may suffer from *scale fragility*: stability and performance can depend critically on the network size and spectral properties of the Laplacian [2]. This motivates alternative architectures that enable scalable high-order coordination using only local and relative information.

Problem Formulation

We consider a network of N identical n^{th}-order integrators and aim to synchronize agent states and the first $n-1$ derivatives using local, relative feedback. Our approach is *compositional*: instead of applying Laplacian feedback directly to multiple derivatives as in (1), we design the closed loop as a cascade of first-order coordination protocols. Specifically, we study the *compositional consensus* dynamics

$$(\frac{\mathrm{d}}{\mathrm{d}t} + \mathscr{L}_n) \circ \cdots \circ (\frac{\mathrm{d}}{\mathrm{d}t} + \mathscr{L}_1)(x) = 0, \qquad (2)$$

where each operator $\mathscr{L}_k$ is Laplacian-like and may be nonlinear and time-varying. Each module enforces relative feedback, i.e., invariance under translations of the form $x \mapsto x + \mathbf{1}a$, and can be designed independently.

The proposed formulation generalizes existing high-order consensus architectures. In the linear time-invariant case, (2) reduces to the serial consensus framework introduced in [3], which was shown to overcome scale fragility affecting classical Laplacian-based designs. In contrast to (1), the compositional structure avoids coordinated tuning across derivative orders and naturally supports modular synthesis and analysis.

This work was partially funded by Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation under the postdoctoral grant KAW 2024.0437.

Main Results

Theorem 1 (Compositional high-order consensus)**.** *Consider* (2)*. Assume that each coordination operator $\mathscr{L}_k$ only uses relative feedback and that the first-order systems $\dot{z} = -\mathscr{L}_k(z,t)$ are asymptotically consensus-seeking with standard regularity and well-posedness assumptions (precise conditions are given in [4]). Then, the compositional consensus system admits a unique solution and achieves n^{th}-order consensus, i.e., all agents synchronize their states and their first $n-1$ derivatives.*

The assumptions in Theorem 1 are deliberately stated at a structural level. They allow each operator $\mathscr{L}_k$ to be nonlinear, time-varying, and heterogeneous across layers, while requiring only relative feedback and standard regularity and stability assumptions. This enables the treatment of coordination protocols with saturations, switching interaction graphs, and time-dependent gains within a unified analytical framework. Figure 1 illustrates the proposed framework for a second-order system using the compositional protocol.

References

[1] W. Ren, K. Moore, and Y. Q. Chen, "High-order consensus algorithms in cooperative vehicle systems," in *Proc. IEEE Int. Conf. Netw. Sens. Control*, 2006, pp. 457–462.

[2] E. Tegling, B. Bamieh, and H. Sandberg, "Scale fragilities in localized consensus dynamics," *Automatica*, vol. 153, p. 111046, 2023.

[3] J. Hansson and E. Tegling, "Closed-loop design for scalable performance of vehicular formations," *IEEE Trans. Control Netw. Syst.*, pp. 1–10, 2025.

[4] ——, "Compositional design for time-varying and nonlinear coordination," *Preprint: arXiv:2504.07226*, 2025.

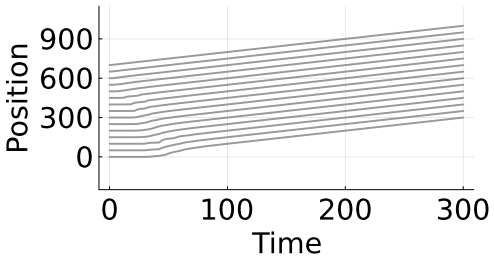

Figure 1: Second-order compositional consensus on a line graph with time-varying coordination weights. Despite intermittent disconnections and initial standstill, the cascade of first-order consensus systems recovers desired formation, consistent with Theorem 1.

`freq-statespace`: A Python package for identification of (non)linear state-space models from periodic input-output data

Merijn Floren[1,2,*] and Jan Swevers[1,2]

[1]MECO Research Team, Department of Mechanical Engineering, KU Leuven, Leuven, Belgium
[2]Flanders Make@KU Leuven, Leuven, Belgium
[*]Email: `merijn.floren@kuleuven.be`

1 Introduction

This work presents `freq-statespace`: a flexible JAX-based [1] package for frequency-domain state-space identification. Its primary focus is on the nonlinear linear fractional representation (NL-LFR) model, which combines a linear time-invariant (LTI) system with a static feedback nonlinearity (see Fig. 1). This internal feedback structure enables the modeling of a broad class of complex dynamics, with classical block-oriented models such as Wiener, Hammerstein, and Wiener–Hammerstein appearing as special cases. Support for standard linear state-space identification is also included. The package is open-source and available on GitHub[1] and via the Python Package Index.

2 Software overview

The package works with (multiple periods and realisations of) multivariate input–output data sequences $u(n) \in \mathbb{R}^{n_u}$ and $y(n) \in \mathbb{R}^{n_y}$ for $n = 0, \ldots, N-1$, assuming periodic excitation and an integer number of steady-state output periods. The specific NL-LFR state-space structure is defined as:

$$x(n+1) = Ax(n) + B_u u(n) + B_w w(n), \tag{1a}$$
$$y(n) = C_y x(n) + D_{yu} u(n) + D_{yw} w(n), \tag{1b}$$
$$z(n) = C_z x(n) + D_{zu} u(n), \tag{1c}$$
$$w(n) = f(z(n)), \tag{1d}$$

consisting of appropriately-sized linear state-space matrices and a static nonlinear function approximator $f(\cdot)$.

A typical step-wise identification procedure is as follows:

1. *Best Linear Approximation (BLA) parametrization.* Initializes the matrices A, B_u, C_y and D_{yu} using the frequency-domain subspace method [2], and refines these estimates through iterative optimization. If only linear state-space models are of interest, the identification process can be terminated at this point.
2. *NL-LFR initialization.* Applies the frequency-domain inference and learning method [3] to efficiently initialize the remaining model parameters while keeping the BLA parameters fixed. This step requires that $f(\cdot)$ is linear in the parameters, i.e., $f(\cdot) = \beta^\top \phi(\cdot)$, with β the parameter vector and $\phi(\cdot)$ the nonlinear feature mapping (e.g., polynomials).

[1]`www.github.com/merijnfloren/freq-statespace`

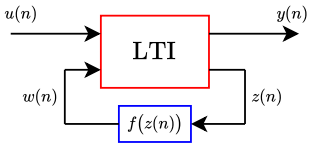

Figure 1: The NL-LFR model structure.

3. *NL-LFR optimization.* Performs iterative refinement of all model parameters using time-domain simulations. This is the most computationally demanding step, mainly due to the sequential nature of the forward simulations. Fortunately, the previous steps should have provided an initialization that is already close to a good local minimum.

Alternatively, the inference and learning step can be skipped in favor of direct nonlinear optimization, which imposes no restrictions on the structure of $f(\cdot)$, i.e., it does not require a model that is linear in its parameters.

The implementation leverages JAX [1] for automatic differentiation, JIT compilation, and GPU acceleration, resulting in highly efficient and scalable computations. Combined with a frequency-domain formulation that naturally exposes parallelism, this enables fast and robust identification even for high-dimensional problems. The framework provides seamless access to both Optimistix [4] solvers (e.g., Levenberg–Marquardt, BFGS) and Optax [5] optimizers (e.g., Adam, SGD), and has been validated on multiple benchmark systems, including a large-scale, high-dimensional multi-input multi-output system [6].

References

[1] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018. software available from `http://github.com/google/jax`.

[2] T. McKelvey and M. Gibanica, "FSID-A frequency weighted MIMO frequency domain identification and rational matrix approximation method for Python, Julia and Matlab," *IFAC-PapersOnLine*, vol. 54, no. 7, pp. 403–408, 2021.

[3] M. Floren, J.-P. Noël, and J. Swevers, "Inference and learning of nonlinear LFR state-space models," *IEEE Control Systems Letters*, vol. 9, pp. 1345–1350, 2025.

[4] J. Rader, T. Lyons, and P. Kidger, "Optimistix: modular optimisation in JAX and Equinox," *arXiv:2402.09983*, 2024.

[5] M. Hessel, D. Budden, F. Viola, M. Rosca, E. Sezener, and T. Hennigan, "Optax: composable gradient transformation and optimisation, in JAX!," 2020.

[6] M. Floren, L. Peri, J. De Maeyer, W. De Munter, D. Vandepitte, and J.-P. Noël, "Data-driven state-space identification and nonlinearity assessment of the CubeSpec Fine Steering Mirror," *ISMA-USD2024*, 2024.

Acknowledgment. This work was done within the framework of Flanders Make's IRVA project ASSISStaNT.

Adaptive-Gain Control of Rational Decision Makers

Global Convergence & Comparative Performance Study[1]

R. Gavin and M. Cao
University of Groningen
{r.c.gavin, m.cao}@rug.nl

1 Abstract

In a large population, peoples' interactions collectively result in population-level behavioral dynamics. These dynamics are sensitive to the way each individual chooses to interact with others [2]. Access to information, the actions of one's peers, and logical reasoning can all factor into an individual's decision making.

Given the variety of ways people can come to decisions, how can a government or other authority figure best steer their constituency towards desired behavior? Given evolutionary game theory's (EGT's) wide range of tools for modeling and analyzing emergent behavioral dynamics [2–4], researchers have studied the control of EGT models to gain insight into this problem.

Many past studies develop controllers for EGT models that ensure desirable long-term behaviors (e.g., [5, 6]), and most of these studies assume that people imitate others. While true in certain situations, this assumption does not hold in general. For example, people often make decisions based on the information and technological means at their disposal. This begs the question, how does controlling other decision makers differ from controlling imitative ones?

To address this question, we investigate the control of *rational* decision makers. Specifically, we apply the adaptive-gain controller from [6] to a population modeled by the *logit dynamics* [2,7,8] and investigate how the control rational decision makers differs from the control of imitative decision makers.

We find that while the same control scheme can guarantee convergence to the desired equilibrium, there are different requirements upon the controller's design parameters for rational and imitative decision makers. Furthermore, we find that, compared to imitative decision makers,

1. The controller is easier to design for rational decision makers; and

2. The controller generally works harder to steer rational decision makers to the desired equilibrium.

[1] This abstract and the accompanying presentation is based on [1].

References

[1] R. Gavin, K. Paarporn, M. Ye, L. Zino, and M. Cao, "Adaptive-gain control for equilibrium selection in the logit dynamics [Manuscript awaiting publication]," in *64th IEEE Conf. Decis. Control*, 2025.

[2] W. H. Sandholm, *Population games and evolutionary dynamics*. MIT Press, 2010.

[3] T. L. Vincent and J. S. Brown, *Evolutionary game theory, natural selection, and Darwinian dynamics*. Cambridge University Press, 2005.

[4] J. M. Smith, *Evolution and the Theory of Games*. Cambridge University Press, 1982.

[5] J. Riehl, P. Ramazi, and M. Cao, "Incentive-based control of asynchronous best-response dynamics on binary decision networks," *IEEE Trans. Control Netw. Syst.*, vol. 6, no. 2, pp. 727–736, 2018.

[6] L. Zino, M. Ye, A. Rizzo, and G. C. Calafiore, "On adaptive-gain control of replicator dynamics in population games," in *62nd IEEE Conf. Decis. Control*, 2023, pp. 485–490.

[7] R. Gavin, M. Cao, and K. Paarporn, "An analysis of logit learning with the r-lambert function," in *63rd IEEE Conf. Decis. Control*, 2024, pp. 6774–6779.

[8] L. Cianfanelli and G. Como, "On the stability of the logit dynamics in population games," *IEEE Trans. Autom. Control*, 2025.

Stability of fast extremum seeking scheme for Wiener systems using frequency response data

Juan Javier Palacios Roman,
Maurice Heemels & Thijs van Keulen.
Control Systems Technology,
Eindhoven University of Technology.
PO Box 513, 5600 MB Eindhoven,
The Netherlands.
Email: {j.j.palacios.roman,
m.heemels, t.a.c.v.keulen}@tue.nl.

Matthijs van Berkel.
Energy Systems and Control,
DIFFER.
PO Box 6336, 5600 HH Eindhoven,
The Netherlands.
Email: M.vanBerkel@differ.nl.

1 Introduction

Extremum seeking control (ESC) is an online, adaptive, model-free optimization method for optimizing the steady-state input-output behaviour of general nonlinear systems. By perturbing the input of a dynamic system and measuring the performance output, the local derivative of the steady-state input-output mapping can be estimated and the performance output can iteratively be maximized. The method relies on a time-scale separation between the perturbation and the system, i.e., the angular frequency of the perturbation ω must be small with respect to the dynamics of the system, which limits the rate of convergence to the optimum. Fast ESC [1] circumvents this time-scale separation, i.e., ω can be chosen freely, and thereby achieves faster convergence to the optimum. It does this by imposing a Wiener system, i.e., the system can be described by

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= h(Cx(t) + Du(t)), \end{aligned} \tag{1}$$

with $x(t) \in \mathbb{R}^n$ the state, $u(t), y(t) \in \mathbb{R}$ the input and performance output, respectively, at time $t \geq 0$, matrices A, B, C, D of appropriate dimensions and nonlinear function h, and by incorporating an estimate $\hat{H}(s)$ of the transfer function $H(s) = C(sI - A)^{-1}B + D$, where $s \in \mathbb{C}$ is the Laplace variable and $I \in \mathbb{R}^{n \times n}$ the identity matrix. The algorithm is visualized in Figure 1, where a and c are the design parameters, $T \coloneqq 2\pi/\omega$ and

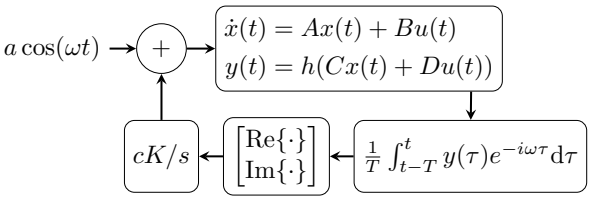

Figure 1: Visualization of the fast ESC scheme in [1].

$$K \coloneqq \frac{2}{a} \cdot \frac{\hat{H}(0)}{\left|\hat{H}(i\omega)\right|^2} \cdot \begin{bmatrix}\mathrm{Re}\{\hat{H}(i\omega)\}\\ \mathrm{Im}\{\hat{H}(i\omega)\}\end{bmatrix}^{\top}. \tag{2}$$

The method can be interpreted as combining the classic model-free ESC algorithm with frequency response data, which are commonly used (and thus commonly available) in the industry, to achieve faster convergence.

2 Contributions

In [2], we present the following contributions: we prove uniform semi-global practical asymptotic stability of the algorithm in [1]; we provide explicit upper bounds on the design parameters a and c; we relax the required accuracy of the estimate $\hat{H}(s)$ of $H(s)$ compared to [1]; and we demonstrate the effectiveness and robustness of the method in a numerical use case, see Figure 2. Ongoing research includes the extension to multi-input Wiener systems and general nonlinear systems.

References

[1] T. van Keulen, R. van der Weijst, and T. Oomen, "Fast extremum seeking using multisine dither and online complex curve fitting," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 5362–5367, 2020.

[2] J. J. Palacios Roman, M. van Berkel, M. Heemels, and T. van Keulen, "Stability analysis of fast extremum seeking control for Wiener systems using online complex curve fitting," *IEEE Control Systems Letters*, vol. 9, pp. 2993–2998, 2025.

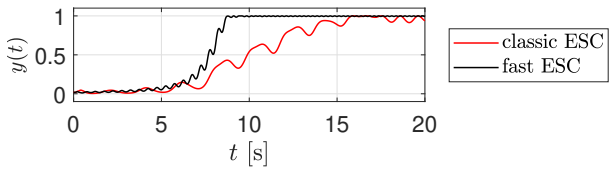

Figure 2: Example of enhanced convergence rate of $y(t)$ when using fast ESC over classic ESC.

Adaptive Control for Resilient Wireless Optical Communication Links

Arturo Morales-Bugueno[1,2,*,†], Rodrigo A. González[1], Tom Oomen[1], and Niek Doelman[1,2]

[1]Control Systems Technology, Eindhoven University of Technology, Eindhoven, The Netherlands
[2]Optomechatronics Department, TNO, Delft, The Netherlands
[*]Email: a.r.h.morales.bugueno@tue.nl

1 Introduction

Free-space optical (FSO) communication refers to the transmission of information using optical waves that travel through an unguided medium. Unlike fiber optics, its guided equivalent, it can transmit data through the atmosphere or free space. Regarding its radio-frequency (RF) counterpart, the advantages are worth mentioning: (1) license-free spectrum, (2) higher data rates, (3) stronger security, and (4) lower power consumption [1]. However, these advantages come at the expense of having to deal with pointing errors and atmospheric distortion, among other challenges. In this work, we explain the impact of these issues and discuss how a control system approach could address them.

2 Channel statistical model

Consider an FSO link with a transmitter and a receiver. We can express the received signal y at the detector as

$$y = hx + n, \tag{1}$$

where x is the transmitted power, $n \sim \mathcal{N}(0, \sigma_n^2)$ is white noise, and h is the channel gain that models the random attenuation of the propagation medium, formulated as

$$h = h_\ell h_a h_p, \tag{2}$$

where h_ℓ represents the path loss, which is a deterministic albeit time-varying function of the weather conditions and the link distance, h_a denotes the atmospheric turbulence loss, which is random with a known probability density function (pdf) that may change over time, and h_p refers to the pointing error loss, which depends on the fixed displacement between the beam and the detector centers (boresight) and the random movements of the beam center on the detector plane (jitter) [2].

Notice that h is a random variable whose pdf, $f_h(h)$, is time-varying ($\sim$ ms) due to the fact that their factors (i.e., h_ℓ, h_a, and h_p) may vary over time. However, as the typical bit interval in FSO links is on the order of ns, $f_h(h)$ can be considered constant over many transmitted bits. An example of $f_h(h)$ under weak turbulence, with zero and non-zero boresight, is shown in Fig. 1. We can notice that boresight negatively impacts h by shifting its pdf toward lower values.

As $h \ll 1$, is random, and time-varying, there will be values for which x will not be successfully transmitted, resulting in bit errors and a reduction in the average channel capacity.

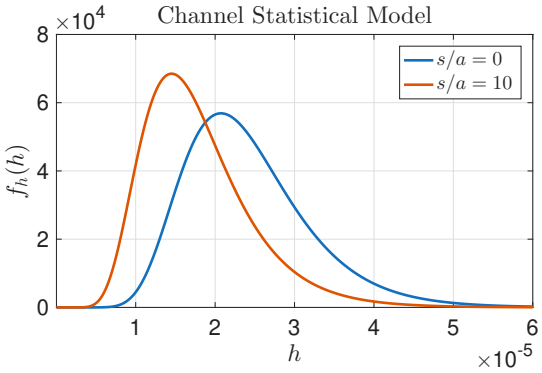

Figure 1: Probability density function of the channel gain h with and without boresight error s under weak turbulence, detector radius $a = 10\,\text{cm}$, beam width at detector $w_z \approx 1\,\text{m}$, and jitter standard deviation $\sigma_s \approx 10\,\text{cm}$.

3 Control approach

Actions that can be taken to overcome the random attenuation of the channel include adapting the transmitted power and improving the tracking system to reduce boresight and jitter. Adaptive methods have been proposed in the literature, but not within a control framework. Therefore, in this work, we investigate how the problem can be expressed using the formalism of the control theory and how this theory could address these adaptive challenges. Our approach has two stages: (i) estimate the channel state h_t and (ii) design an adaptive controller $C(s)$ that adjusts its parameters based on this estimate, while accounting for the non-negligible feedback delay. The output of $C(s)$ may be the transmitted power, which we may minimize to save energy.

References

[1] M. A. Khalighi and M. Uysal, "Survey on free space optical communication: A communication theory perspective," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 4, pp. 2231–2258, 2014.

[2] F. Yang, J. Cheng, and T. A. Tsiftsis, "Free-space optical communication with nonzero boresight pointing errors," *IEEE Transactions on Communications*, vol. 62, no. 2, pp. 713–725, Feb. 2014.

[†]This project has received funding from the EU's Horizon Europe research and innovation programme under the MSCA grant agreement No 101.169.042.

Experiment design for continuous-time systems using generalized filtering

Jiwei Wang
University of Groningen
9747 AG, Groningen
The Netherlands
jiwei.wang@rug.nl

Simone Baldi
Southeast University
211189, Nanjing
China
simonebaldi@seu.edu.cn

Henk J. van Waarde
University of Groningen
9747 AG, Groningen
The Netherlands
h.j.van.waarde@rug.nl

1 Introduction

This abstract investigates the problem of experiment design for continuous-time systems under piecewise constant input signals. To obviate the need for measuring time derivatives of (data) trajectories, we introduce a generalized filtering framework. Our main result is to establish conditions on the input and the filter functions under which the filtered data are informative for system identification, i.e., they satisfy a certain rank condition.

2 Problem Formulation

Consider the continuous-time linear time-invariant system

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0, \tag{1}$$

where $x(t) \in \mathbb{R}^n$ is the system state, $u(t) \in \mathbb{R}^m$ is the control input, and $A \in \mathbb{R}^{n\times n}, B \in \mathbb{R}^{n\times m}$ are the system matrices, assumed to be unknown, with the pair (A,B) being controllable. Let $\mathsf{T} > 0$ and $\mathscr{AC}([0,\mathsf{T}),\mathbb{R}^n)$ be the space of absolutely continuous functions from $[0,\mathsf{T})$ to $\mathbb{R}^n$, and let $\mathscr{PC}([0,\mathsf{T}),\mathbb{R}^m)$ be the space of piecewise continuous functions from $[0,\mathsf{T})$ to $\mathbb{R}^m$. We define the behavior of (1) as

$$\mathfrak{B}_{\mathsf{T}} := \{(x,u) \in \mathscr{AC}([0,\mathsf{T}),\mathbb{R}^n) \times \mathscr{PC}([0,\mathsf{T}),\mathbb{R}^m) \mid x \text{ is a Carathéodory solution of (1) on } [0,\mathsf{T})\}.$$

For $(u,x) \in \mathfrak{B}_{\mathsf{T}}$, we denote the collection of *filtered data* by

$$x^{\mathrm{df}}_{[1,M]} := [x^{\mathrm{df}}_1 \ \cdots \ x^{\mathrm{df}}_M], x^{\mathrm{f}}_{[1,M]} := [x^{\mathrm{f}}_1 \ \cdots \ x^{\mathrm{f}}_M], u^{\mathrm{f}}_{[1,M]} := [u^{\mathrm{f}}_1 \ \cdots \ u^{\mathrm{f}}_M],$$

where

$$x^{\mathrm{df}}_\ell = \int_0^{\mathsf{T}} g_\ell(\tau)\dot{x}(\tau)d\tau, x^{\mathrm{f}}_\ell = \int_0^{\mathsf{T}} g_\ell(\tau)x(\tau)d\tau, u^{\mathrm{f}}_\ell = \int_0^{\mathsf{T}} g_\ell(\tau)u(\tau)d\tau$$

for $\ell \in \{1,2,\dots,M\}$. The filter functions $g_\ell : [0,\mathsf{T}) \to \mathbb{R}$ are continuously differentiable on each interval $[t_{j-1},t_j)$ for $j = 1,2,\dots,q$, where $0 = t_0 < t_1 < \cdots < t_q = \mathsf{T}$ and $g_\ell(t_j^-) := \lim_{t\uparrow t_j} g_\ell(t)$. We consider the following problem.

Problem 1 *Consider the trajectory $(u,x) \in \mathfrak{B}_{\mathsf{T}}$ and the filter functions g_ℓ for $\ell = 1,2,\dots,M$. Provide conditions on this trajectory and filter functions such that*

$$\operatorname{rank}\left(\begin{bmatrix} x^{\mathrm{f}}_{[1,M]} \\ u^{\mathrm{f}}_{[1,M]} \end{bmatrix}\right) = n+m. \tag{2}$$

3 Main Results

Let $T > 0$ and $N \in \mathbb{N}$ and define $\mathsf{T} := NT$. We consider a piecewise constant input $u(t+kT) = \mu_k$, where $t \in [0,T)$ and $\mu_k \in \mathbb{R}^m$ for $k = 0,1,\dots,N-1$. For the trajectory $(u,x) \in \mathfrak{B}_{NT}$, we consider *sampled data* at the time instants kT, captured in the matrices

$$\chi_{[0,N-1]} := \begin{bmatrix} \chi_0 & \cdots & \chi_{N-1} \end{bmatrix}, \mu_{[0,N-1]} := \begin{bmatrix} \mu_0 & \cdots & \mu_{N-1} \end{bmatrix}.$$

where $\chi_k = x(kT)$.

We then make the following assumption for g_ℓ.

Assumption 1 *For all $\tau \in [0,T)$ and $j \in \{0,1,\dots,N-1\}$, the filter function g_ℓ satisfies*

$$g_\ell(\tau + jT) = g(\tau)f_\ell(jT),$$

where $g : [0,T) \to [0,\infty)$ is a continuously differentiable function such that $\int_0^T g(\tau)d\tau > 0$, and $f_\ell : \mathbb{R} \to \mathbb{R}$ for $\ell \in \{1,\dots,M\}$.

Finally, we have the following result.

Proposition 1 *Let Assumption 1 hold. If $N = n+m$,*

$$\operatorname{rank}\left(\begin{bmatrix} f_1(0) & \cdots & f_M(0) \\ \vdots & \ddots & \vdots \\ f_1((N-1)T) & \cdots & f_M((N-1)T) \end{bmatrix}\right) = n+m,$$

and

$$\operatorname{rank}\left(\begin{bmatrix} \chi_{[0,N-1]} \\ \mu_{[0,N-1]} \end{bmatrix}\right) = n+m, \tag{3}$$

then (2) holds.

Proposition 1 states that under an appropriate design of f_ℓ for $\ell \in \{1,\dots,M\}$, the rank condition (2) on the filtered data directly follows from the rank condition (3) on the sampled data. Hence, it suffices to formulate an experiment design for the input samples: more details can be found in [1].

References

[1] J. Wang, S. Baldi, and H. J. van Waarde, "Experiment design for continuous-time systems using generalized filtering," *arXiv preprint arXiv:2511.09386*, 2025.

Data-driven oscillatory model of memory consolidation: Perspective for control by optogenetics

Mario Valencia Torres, Joshua Amato, and Matin Jafarian
Delft University of Technology
Email: {mvalenciatorre, j.amato, m.jafarian @tudelft.nl}

1 Introduction

Memory consolidation relies on coordinated interactions between the hippocampus and the neocortex. Newly acquired information is rapidly encoded in the hippocampus (HPC) and subsequently reorganized into distributed neocortical representations through repeated replay of hippocampal activity, a process that occurs predominantly during sleep. Experimental evidence indicates that this transfer of information is tightly associated with specific oscillatory bands, most notably hippocampal sharp-wave ripples (SWRs) and gamma rhythms that are temporally coordinated with slow cortical oscillations [1]. A schematic representation of the bidirectional coupling between the HPC and the retrosplenial cortex (SRC) is shown in Figure 1.

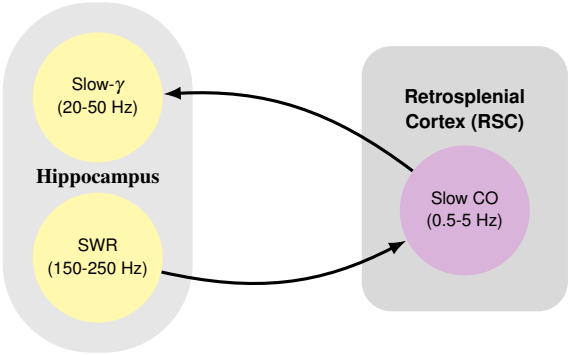

Figure 1: Schematic of hippocampo-cortical connectivity during memory consolidation in the NREM sleep phase.

In this work, we adopt a data-driven modeling approach grounded in experimental recordings. Our objective is twofold: first, to capture the dominant oscillatory interactions underlying hippocampo-cortical communication, and second, to establish a modeling framework that can support future closed-loop modulation of memory consolidation.

2 Data-driven modeling

While biophysical models, e.g., Wilson-Cowan or Jansen-Rit have often used to model brain dynamics, their complexity often limits their suitability for real-time estimation and control. To address this limitation, we employ the HAVOK-ARX framework [2], which provides a low-dimensional linear representation of nonlinear dynamics through delay embeddings and Koopman operator theory.

This approach identifies a small number of dominant modes whose evolution can be approximated by a forced linear time-invariant system.The resulting forced linear system can be written as

$$x_{k+1} = Ax_k + N_r x_r + Bu, \tag{1}$$

where Ax_k captures the intrinsic linearized dynamics, $N_r x_r$ represents the latent forcing terms associated with nonlinear effects in the delay-embedded Koopman space, and Bu denotes external inputs acting on the system.

3 Perspective for control

Optogenetics enables temporally precise and cell-type-specific modulation of neural activity through light-sensitive ion channels [3]. Compared to electrical stimulation, it offers improved specificity and reduced disruption of surrounding circuits, making it well suited for targeting memory-related dynamics.

We propose a closed-loop strategy in which optogenetic stimulation modulates hippocampal activity linked to SWR generation, with the goal of shaping gamma-band dynamics. By incorporating optogenetic inputs as control terms in the data-driven model, this framework provides a principled path toward real-time closed-loop regulation of memory consolidation during sleep.

4 Acknowledgment

We would like to thank DBI2 consortium, F. Battaglia, A. Schouten, and W. Serdijn for helpful discussions.

References

[1] R. Pedrosa et al., "Hippocampal gamma and sharp wave/ripples mediate bidirectional interactions with cortical networks during sleep," *PNAS*, vol. 119, no. 44, 2022.

[2] J. C. S. Amato, "On Koopman Delay Embeddings for Nonlinear Oscillators: A Data-Driven Modeling Approach of Memory Consolidation in the Mouse Brain," MSc Thesis, Delft University of Technology, 2025.

[3] E. S. Boyden et al., "Millisecond-timescale, genetically targeted optical control of neural activity," *Nature Neuroscience*, vol. 8, no. 9, pp. 1263–1268, 2005.

Data-driven mechanisms of cortical memory networks in cognition

Matin Jafarian
Delft University of Technology
Email: m.jafarian@tudelft.nl

1 Introduction

The importance of understanding human memory functioning is evident from its central role in cognitive processes and its influence on the development of artificial memory networks. Simply put, memory is acquired information that is stored, processed, recalled, or forgotten. Its dynamics range from the connection of two cells [1] to networks of integrated cells and populations across spatial and temporal scales, including short-term and long-term processes [2–5].

The neuronal dynamics of memory have been understood via detailed bio-physical, or phenomenological, abstract [2] or purely data-driven approaches. While detailed biophysical models are mainly relevant to the cell level [1] and networks of cells, they are less useful in cortical representations. They simply have too much parameters to be identified, rapidly scale in complexity and being computationally demanding.

Cortical memory dynamics of human brain are relevant in studying the process of memory consolidation as well as the interplay of memory with higher order cognitive tasks, such as attention regulation, decision making, and the interplay with emotions. To reveal these mechanisms, mainly data analysis techniques have been employed, and partly phenomenological models. In addition to mechanistic understanding, regulation and control is a new focus in brain science owing to the advancement of measuring and intervention technology.

In this talk, we present some examples of our recent work on data-driven mechanistic understanding of memory, discuss the challenges, contributions, and perspectives.

2 Identification of memory dynamics

We focus on data-driven modeling, informing the models with no or with only some high-level biological hypotheses. We consider two experiments using fMRI and MEG recordings of several brain regions whose states are stacked in vector $\mathbf{x} \in \mathrm{R}^n$. In the fMRI study, we consider an associate memory encoding task in which the interplay of emotions and memory was examined. We consider a geneal nonlinear dynamics with inputs, in the form,

$$\dot{\mathbf{x}} = f(\mathbf{x}, u, \theta),$$

where the function f, its relation with the exogenous input u, and its corresponding parameters are to be identified. Using the dynamic casual modeling framework, we show that a stochastic bilinear model best describes the data [4].

In the MEG study, we consider a task of working memory in which the interplay of attention and memory was examined. We propose the following oscillatory network model [5]

$$\dot{\mathbf{x}} = f(\mathbf{x}, \theta_f) + \lambda g(\mathbf{x}),$$

where the function f is assumed known with unknown parameters, function g which represents the attention mechanism, $\lambda \in \{0, 1\}$, is to be identified. For the latter, We use universal differential equations (UDEs) [5].

3 Results and Discussion

Our focus will be devoted to the identified dynamics, their level of complexity compared with biophysical and abstract models. We discuss the results, their utilization in neuroscience and medical research. We also discuss how analytical tools from systems and control theory could enhance revealing, analyzing and regulating these mechanisms.

References

[1] Würzbauer, V., Kerstin L., and M. Jafarian. *Astrocytic gliotransmission as a pathway for stable stimulation of postsynaptic spiking: Implications for working memory.* IFAC-PapersOnLine 56.2 (2023): 488-493.

[2] M. Jafarian, D., Huerta, G., Villani, A., Lansner, and K.H., Johansson. *Cluster Synchronization as a Mechanism of Free Recall in Working Memory Networks.* IEEE Open Journal of Control Systems, 2023.

[3] S. Bakels, M. van de Ruit, and M. Jafarian. *Accurate linear state-space model of EEG-Based cortical activity during a passive motor task with exogenous input: Perspective for control.* arXiv, 2025.

[4] W. Dziarnowska, M. Orhun, Y. Zhu, N. Kohn, G. Fernandez, and M. Jafarian. *Associative emotional memory encoding: Insights from network stability analysis of an fMRI-driven bilinear dynamics.* bioRXiv, 2026.

[5] S.M. Ohkawa, S. Haegens, and M. Jafarian. *A novel data-driven approach for identifying causal interplay of working memory and attention using MEG data.* 2026.

Spatio-temporal Excitation with Negative Threshold

Thomas Burger, Rodolphe Sepulchre
STADIUS, ESAT, KU Leuven, Leuven 3001, Belgium
thomas.burger@kuleuven.be, rodolphe.sepulchre@kuleuven.be

1 Introduction

Spiking events, of interest to neuromorphic engineers and neuroscientists, are the result of excitable systems crossing a threshold. The spiking behaviors exhibited by such systems span both space and time, as activity can be restricted to one neuron or an entire ensemble of neurons, and different types of spikes can have durations ranging from one millisecond to hundreds of milliseconds.

In this talk, we investigate the behaviors helped by equipping a neuronal population with a *negative* threshold, that is, a threshold below the resting potential, which can be reached by an inhibitory (negative) applied current, traditionally regarded as being a 'brake' on activity. To investigate this, we use a model that allows for the seamless generalization of spiking systems from a single neuron to large neuronal populations, and from fast temporal scales to ultraslow ones. The key features that enables this are the memristive elements that we know from biophysical neuron models, explicitly separated in multiple spatial and temporal scales. The memristive elements in biophysical models, which are combined with parallel capacitors and banks of voltage sources, is what gives neurons their complex nonlinear dynamics. Therefore, traditionally, they are regarded as a source of complexity. Simplified neuronal models do away with this complexity by replacing the memristive elements with nonlinear amplifiers [1].

Instead, we highlight that the memristive nature of the elements helps in designing and simulating neuromorphic systems across scales. We illustrate the flexibility of this modeling framework by constructing and contrasting neuronal populations with and without a negative threshold.

2 Spatio-temporal Memristive Elements

We use a canonical excitable neuron model, where the currents are memristive elements. The memductances are modelled as first-order lags of the membrane potential, operating in finite and distinct time scales. In other words, they are modeled as an input-output fading memory operator of the voltage history. The spatial interactions of neurons are modeled analogously, with memductances which are modeled as input-output spatio-temporal operators on the space of voltage trajectories. The temporal dependence is modeled as a time-invariant fading memory operator, whilst the spatial dependence is modeled as a space-invariant myopic operator. These spatio-temporal memristive elements model the synaptic currents between neurons, analogously to how biophysical ion channels are modeled as temporal convolution operators.

To contrast between populations with and without an inhibitory threshold, we model two neuronal populations. The first population has memristive elements which below the resting potential are purely Ohmic, and therefore can only facilitate a return to the resting voltage. The second population is equipped with memductances which have dynamics such that negative spatial connections can actually facilitate excitability, as opposed to merely functioning as a brake on activity.

3 Neuron Populations with Negative Threshold

The features described above mean that the model displays all interesting behaviors shown by excitable circuits: The separate timescales for the memristive currents in the individual neurons mean that the neuron displays the same array of excitable behaviors as biophysical neurons do. The spatial couplings mean that the circuit is capable of the same interactions that orchestrate spiking computations in the brain.

We investigate the importance of negative thresholds in a neuronal population setting, and demonstrate spatio-temporal excitability with both a negative and a positive thresholds. We also illustrate the usefulness of a negative threshold in two population-level phenomena which are dominated by inhibition: neuronal oscillations and winner-take-all circuits.

References

[1] Fitzhugh, R. Thresholds and Plateaus in the Hodgkin-Huxley Nerve Equations. J Gen Physiol 43, 867–896 (1960).

Acknowledgment This work was supported by European Research Council under the Advanced ERC Grant Agreement "SpikyControl" 101054323.

Context-dependent manifold learning in dynamical systems: A neuromodulated constrained autoencoder approach

Jérôme Adriaens, Guillaume Drion, Pierre Sacré
Dept. of Electrical Engineering and Computer Science, University of Liège, Belgium
jadriaens@uliege.be, gdrion@uliege.be, p.sacre@uliege.be

1 Introduction

Dimensionality reduction (DR) is a key tool for analyzing high-dimensional data from complex dynamical systems. While linear methods such as principal component analysis (PCA) and nonlinear approaches like autoencoders effectively capture structure in fixed settings [1], they typically assume a unique data-generating process. In practice, however, the geometry of the data manifold often depends on external conditions such as physical parameters or environmental context. Simply augmenting the input with contextual variables increases dimensionality and may obscure the underlying structure. We propose a context-dependent DR approach based on the constrained autoencoder (cAE) [2], which enforces smooth manifold structure through geometric constraints. By integrating neuromodulation mechanisms conditioned on a context vector, the model adapts its projection while preserving geometric consistency [3]. We demonstrate the effectiveness of the approach on a 16-degree-of-freedom pendulum system with different context parameters.

2 Neuromodulated Constrained Autoencoder

We propose a context-adaptive constrained autoencoder (cAE) that combines geometric manifold learning with neuromodulation. The cAE framework enforces that the encoder ρ and decoder φ compose into an idempotent projection, *i.e.*, $\rho \circ \varphi = \mathrm{Id}$, so that $P = \varphi \circ \rho$ defines a smooth projection and the learned representation $\hat{M} = \mathrm{Range}(P)$ forms a smooth embedded submanifold of the data space [2]. The encoder and decoder are constructed from paired layers with biorthogonal weights $\vec{\Psi}_l, \vec{\Phi}_l$ and inverse activation functions (σ_-, σ_+):

$$\rho^{(l)}(\vec{x}^{(l)}) = \sigma_- \left(\vec{\Psi}_l^T (\vec{x}^{(l)} - \vec{b}_l) \right),$$
$$\varphi^{(l)}(\vec{z}^{(l-1)}) = \vec{\Phi}_l \sigma_+ \left(\vec{z}^{(l-1)} \right) + \vec{b}_l,$$

with the constraint $\vec{\Psi}_l^T \vec{\Phi}_l = \vec{I}$. We enforce this constraint via Riemannian optimization on the biorthogonal manifold. To enable adaptation to external conditions, we introduce neuromodulation [3]. Neuromodulation is a fundamental mechanism that uses neuromodulators to tune neuronal properties and circuit function dynamically. It is essential for generating flexible brain states and complex behaviors [4]. Fixed activation functions $\sigma_\pm$ are replaced by parameterized functions $\sigma_\pm(\cdot; \vec{\alpha}^{(l)})$, where the parameters depend on a context vector $\vec{c}$. Specifically, the context is mapped to a global modulation signal $\vec{s}$ via a fully connected layer, and each layer parameters are obtained as $\vec{\alpha}^{(l)} = g(\vec{W}_l^T \vec{s})$, where g is a saturating function to keep $\vec{\alpha}^{(l)}$ bounded. This allows the nonlinear transformations to adapt smoothly to the context while preserving the geometric structure of the cAE.

3 Experiments & Results

We evaluate three architectures on two 16-DoF pendulum tasks. The *standard pendulum* randomly samples the lengths of the first four links in $[0.35, 0.65]$m while keeping the remaining 12 DoF fixed. The *context-dependent pendulum* introduces a coupling between link lengths and joint angles, yielding a context-dependent setting. We compare: (i) a constrained autoencoder (*cAE*), (ii) a modulated cAE with irrelevant context (*IcAE*, $\vec{c} = \vec{0}$), and (iii) a modulated cAE with relevant context (*McAE*). Performance is assessed via RMSE on position and velocity reconstruction on unseen test trajectories. In the standard task, all models perform similarly, as expected for context-independent dynamics. In contrast, in the context-dependent task, McAE outperforms both cAE and IcAE, reducing reconstruction errors by a factor of 2–3. These results demonstrate that neuromodulation effectively exploits relevant contextual information for context-dependent dimensionality reduction.

Acknowledgment This work was supported by the Belgian Government through the Federal Public Service Policy and Support.

References

[1] Y. Wang, H. Yao, and S. Zhao, "Auto-encoder based dimensionality reduction," *Neurocomputing*, RoLoD: Robust Local Descriptors for Computer Vision 2014, vol. 184, pp. 232–242, 2016.

[2] S. E. Otto, G. R. Macchio, and C. W. Rowley, *Learning Nonlinear Projections for Reduced-Order Modeling of Dynamical Systems using Constrained Autoencoders*, 2023.

[3] N. Vecoven et al., "Introducing neuromodulation in deep neural networks to learn adaptive behaviours," *PLOS ONE*, vol. 15, no. 1, e0227922, 2020.

[4] C. I. Bargmann and E. Marder, "From the connectome to brain function," *Nature Methods*, vol. 10, no. 6, pp. 483–490, 2013.

Inter-personal variability in muscle synergies for gait control

Rodrigo Moreno-Morton
Allée de la Découverte 10, Liege, Belgium
lrmoreno@uliege.be

Alessio Franci
Allée de la Découverte 10, Liege, Belgium
afranci@uliege.be

Human gait control allows people to walk in diverse and quickly-changing environments that require adaptation to changes in external (*i.e.* slope, evenness) and internal (*i.e.* movement range, lesions) conditions. This is facilitated by a feedback loop between the central nervous system (CNS), which sends out motor commands, and the muscles in the periphery, which execute said commands and send sensory information back to the CNS. This sensory information is then integrated to modify future motor commands. Understanding how this sensory-motor feedback loop leads to fast and flexibly adaptive motor behavior would enable us to design better prosthetic and robotic controllers.

Motor and sensory signals live in a high-dimensional space, roughly determined by the amount of muscles in the body. Trying to implement a controller in those dimensions would require a very high computational cost, and would make its implementation impractical. However, it has been reported repeatedly that the brain works in lower dimensions to simplify the control problem [1, 2]. We call these lower-dimensional signals *motor* and *sensory synergies*. While motor synergies can be easily obtained through dimension-reduction techniques using non-invasive muscle activation data, sensory synergies are mostly a hypothesis due to experimental limitations [2].

In this project, we want to calculate muscle synergies in different walking conditions, and use the changes in muscle activation to estimate the dynamics of sensory signals, calculate their synergies, and close the control loop with neuro-inspired decision-making models [3]. A first result is the calculation of motor synergies with classical tools like principal component analysis (PCA) or non-negative matrix factorization (NMF). These synergies can then be compared across activities and subjects to determine their similarity, *i.e.* how similar the solutions to the motor problems are. As expected, synergies are similar across tasks, showing that a person uses similar strategies to walk in similar conditions, instead of learning entirely different ways to walk. However, comparing across subjects we can see that muscle synergies also differ significantly [4], indicating that different people reach different solutions for locomotion, due to either different exercise habits, body capacity, or chance.

Understanding how motor synergies vary not only across tasks but also across subjects allows us to better understand the control of gait, but also renders classical techniques like PCA and NMF impractical. In order to understand this variation fully, we need to separate the effects of different variables before doing the analysis of these "marginalized" traces. The analysis with demixed PCA (dPCA) [5] allows us to do just that. In figure 1, we show the activity of two muscles during one step for three subjects (left column), and the isolation of different values that allow us to see the variance between an "average step" (third column) and the step for each individual subject (fourth column). Dimension reduction can then be performed on each demixed signal to understand how muscle activation and synergies vary across different variables.

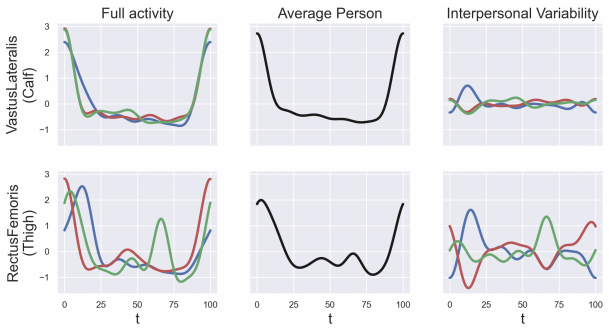

Figure 1: Marginalization of the data according to time and subject. Each row represents a muscle. Each column represents different contributions to the overall activity (first column). Most importantly, the last column shows the individual variation with respect to an average step.

References

(1) L. H. Ting and J. L. McKay, *Current Opinion in Neurobiology*, 2007, **17**, 622–628.

(2) W. J. Magrath, C. S. Qiu, P. J. Hanwright, S. H. Tuffaha and N. Khavanin, *Plastic and Reconstructive Surgery - Global Open*, 2022, **10**, e4438.

(3) N. E. Leonard, A. Bizyaeva and A. Franci, *Annual Review of Control, Robotics, and Autonomous Systems*, 2024, **7**, annurev-control-090523–100059.

(4) J. Tsuchiya, K. Momose, H. Saito, K. Watanabe and T. Yamaguchi, *Frontiers in Neuroscience*, 2024, **Volume 18 - 2024**, DOI: 10.3389/fnins.2024.1485066.

(5) D. Kobak, W. Brendel, C. Constantinidis, C. E. Feierstein, A. Kepecs, Z. F. Mainen, X.-L. Qi, R. Romo, N. Uchida and C. K. Machens, *eLife*, 2016, **5**, e10989.

Reinforcement Learning for Diverse Locomotion of the Human Musculoskeletal System

Kelvin Toonen
Delft Center for Systems and Control
Delft University of Technology
k.a.b.toonen@tudelft.nl

Shengling Shi
Delft Center for Systems and Control
Delft University of Technology
shengling.shi@tudelft.nl

1 Introduction

Investigating how the human body works is an influential research topic, due to its many possible applications in healthcare. Among these is the area of human locomotion, where details about the dynamics of the body can help understand one of the most basic activities in everyday life. To avoid invasive procedures, studying these inner workings is done through musculoskeletal modeling, which has led to advances in exoskeleton design, injury risk, disease analysis and synthetic gait generation.

We use deep reinforcement learning based on motion imitation to train a controller that can give the necessary muscle activations to reach the desired next state, similarly to inverse dynamics. This has been done successfully in other research, but these approaches rely on multiple reinforcement learning controllers that specialize in different movements during locomotion. In our work, we will focus on designing the architecture to allow for a single controller to achieve good performance on all aspects of standard locomotion.

2 Musculoskeletal Models

Human musculoskeletal models have a number of muscles acting on certain joints of the skeleton. The dynamics of this system are usually described as:

$$M(q)\ddot{q} + c(q,\dot{q}) = \sum_m J_m^\top f_m(l_m, \dot{l_m}, a_m) + J_c^\top f_c + \tau_{ext}, \quad (1)$$

where q are the generalized coordinates, $M(q)$ is the mass matrix and $c(q,\dot{q})$ is the combination of Coriolis and gravitational force. f_m and f_c are muscle and constraint forces, respectively, J_m and J_c are Jacobians that map forces to generalized coordinates and τ_{ext} are the external forces.

For the muscle forces f_m, the Hill-type muscle formulation is often used, such that for each muscle we have

$$f_m(l_m, \dot{l_m}, a_m) = a \cdot f_l(l_m) \cdot f_v(\dot{l_m}) + f_p(l_m), \quad (2)$$

where l_m is the muscle length, a_m is the muscle activation, f_{l_m} and f_v are force length and velocity functions, respectively, and f_p is the passive muscle force that represents the elasticity of the muscle. In many cases, however, the length of the muscles is assumed to be constant, simplifying it to only need the variable a_m.

3 Proposed Method

For the optimal control of systems similar to (1), deep reinforcement learning has already been compared with optimization-based approaches and has been shown to be faster than approaches with a similar level op precision. Analyzing muscle behavior, given a recorded motion, is typically done by optimizing a musculoskeletal model to this reference. However, this can also be achieved similarly through deep reinforcement learning.

We will use deep reinforcement learning based on motion imitation. The goal of deep reinforcement learning is to collect experience with a controller to optimize the return R

$$\max_\theta \mathbb{E}_\pi(R) = \max_\theta \mathbb{E}_\pi(\sum_{t=0}^{\infty} \gamma^t r_t), \quad (3)$$

where θ are the network parameters describing the control policy, γ is a discount factor and r_t is the reward gained at time step t. The reward function guides the controller to perform behavior that we want to see. In the case of motion imitation, we have that

$$r_t = \sum_j (||q_t^j - \hat{q}_t^j||^2 + ||\dot{q}_t^j - \hat{\dot{q}}_t^j||^2), \quad (4)$$

where $\hat{q}$ and $\hat{\dot{q}}$ are the reference generalized positions and velocities, respectively, and j indexes the joints. The benefit of using motion imitation is that this is a clear way to describe natural human motions.

The novelty of our research will be in designing the training pipeline of the controller, such that a single controller can perform different aspects needed to perform natural human motion. We will use Lattice [1] to explore efficiently in an overactuated space and a hierarchical network, inspired by [2], to incorporate the difference between high-level movements and low-level muscle activations into the model.

References

[1] Chiappa, Alberto Silvio, et al. "Latent exploration for reinforcement learning." Advances in Neural Information Processing Systems 36 (2023): 56508-56530.

[2] Lee, Seunghwan, et al. "Scalable muscle-actuated human simulation and control." ACM Transactions On Graphics (TOG) 38.4 (2019): 1-13.

Teacher–Student Reinforcement Learning for Reeds–Shepp Vehicles: Kinematically Feasible Motion Planning

Ibrahim Ibrahim, Wilm Decré, and Jan Swevers
MECO Research Team, Department of Mechanical Engineering, KU Leuven
Leuven, 3000, Belgium
Flanders Make@KU Leuven, Belgium
ibrahim.ibrahim@kuleuven.be

1 Introduction

Numerous works have demonstrated reinforcement learning (RL) and imitation learning (IL) in robot navigation, motion planning, and driving [1]. Typical methods use RL to make high-level decisions, like selecting among a discrete set of actions, or to generate intermediate waypoints, but often face challenges with instability in training and producing dynamically/kinematically infeasible or jerky paths [2]. While such methods enable systems to adapt behaviors beyond hand-engineered heuristics, they still face open challenges, namely out-of-distribution generalizability, sample inefficiency, limited interpretability, and safety concerns — which leads to certification challenges in safety-critical domains. In this work, we design and implement a physics-informed hierarchical mixture-of-experts teacher–student IL & RL framework for kinematically feasible motion planning of Reeds–Shepp vehicles in 2D Cartesian grids.

2 Methodology & Results

We leverage distance-optimal Reeds–Shepp (RS) planning [3] as a geometric prior within a learned planner that combines classical geometry with modern deep learning. A fast free-space RS solver [3] serves as the low-level layer, and a clutter-aware RS planner provides expert trajectories. Training follows a three-stage teacher–student paradigm (privileged teacher (access to expert plan) → distillation → RL fine-tuning), which is designed to efficiently distill expert knowledge into a lightweight, deployable student policy. Grounding decisions in distance-optimal RS geometry yields inherently kinematically feasible actions at selection time (no post-hoc projections), i.e., feasibility-by-construction.

Preliminary experiments in randomized cluttered environments achieve $\sim 95\%$ training rollout success and 100% validation success (success = goal reached with collision-free, kinematically feasible motion), as shown in Figure 1, using compact simplified models enabling real-time inference (on the order of a few million parameters). We also observe improved efficiency and safety (Figure 2), indicating faster and safer navigation in randomly generated environments and motivating further optimization.

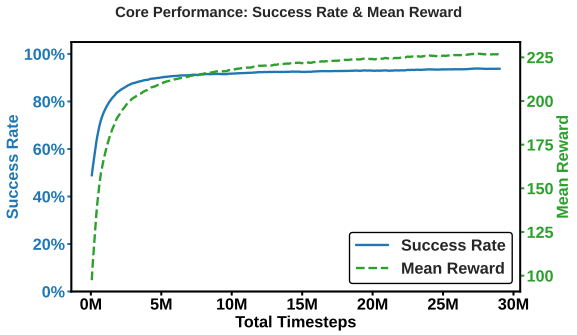

Figure 1: Success means that the policy managed to connect the start and goal configurations p_s and p_f using a chain of intermediate configurations in a safe, collision-free manner, while adhering to behavior shaped by the reward function.

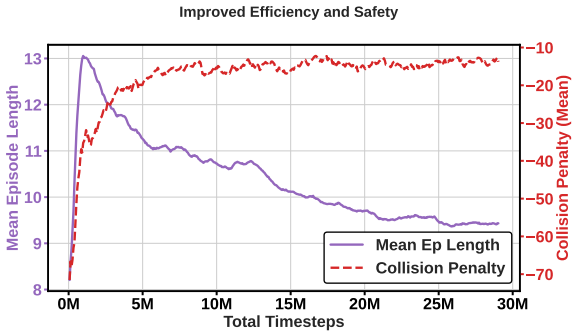

Figure 2: Efficiency (number of steps taken before termination), and safety (penalty for collisions) metrics, shaped by our defined reward function, improving over time.

References

[1] B. Ravi Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani and P. Pérez, "Deep Reinforcement Learning for Autonomous Driving: A Survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2022, doi: 10.1109/TITS.2021.3054625

[2] M. Reda, A. Onsy, A. Y. Haikal and A. Ghanbari, "Path planning algorithms in the autonomous driving system: A comprehensive review," *Robotics and Autonomous Systems*, vol. 174, p. 104630, 2024, doi: 10.1016/j.robot.2024.104630

[3] I. Ibrahim, W. Decré and J. Swevers, "Accelerated Reeds–Shepp and Under-specified Reeds–Shepp Algorithms for Mobile Robot Path Planning," *IEEE Transactions on Robotics*, vol. 41, pp. 2691–2708, 2025, doi: 10.1109/TRO.2025.3554406

Acknowledgement This work has been carried out within the framework of Flanders Make's SBO project ARENA (Agile & Reliable Navigation). Flanders Make is the Flemish strategic research centre for the manufacturing industry.

AI-enhanced wastewater treatment process automation

Hadisi, S.[1,2,*], Wagterveld, R.M.[2,3], Mels, A.[1], Hans, J.D.[1], Schott, C.[2], and Keesman, K.J.[1,2]
[1]Mathematical and Statistical Methods (Biometris), Wageningen University and Research,
PO Box 16, 6700 AA, Wageningen, The Netherlands
[2]Wetsus, European Centre of Excellence for Sustainable Water Technology,
Oostergoweg 9, 8911 MA, Leeuwarden, The Netherlands
[3]Aqa.earth, Tweede Sluisweg 35, 8413 NN Oudehorne, The Netherlands
*Email: soheil.hadisi@wur.nl

1 Introduction

Wastewater treatment plants (WWTPs) must meet effluent quality targets under strongly fluctuating influent conditions (e.g., rain-driven dilution and seasonal variability) while also reducing energy consumption. Recent studies highlight the growing role of digitalisation, hybrid modelling, and control strategies to improve process monitoring and operational efficiency, but practical implementation is often limited by data quality issues and the complexity of real plants [1]. This project investigates sensor-driven hybrid modelling for improved state estimation and predictive decision support in WWTP operation

2 Case study and data acquisition

An industrial wastewater treatment plant will serve as the first full-scale case study for deploying online sensing and collecting high-frequency operational data under real conditions. The plant treats wastewater generated by a municipal waste recycling facility and discharges to the municipal sewer system for downstream treatment. Online sensor measurements and SCADA/operational data will be complemented with laboratory analyses to support model development, calibration, and performance evaluation.

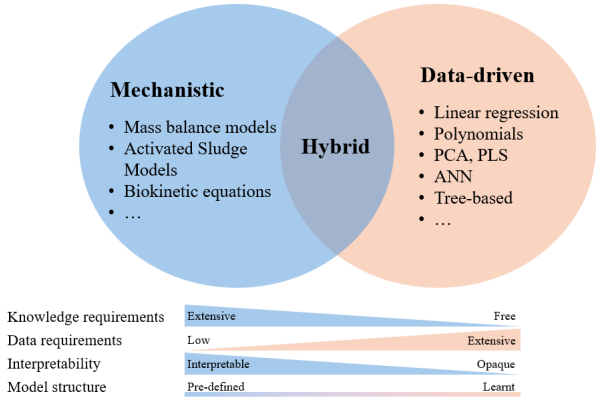

Figure 1: Venn diagram of hybrid modelling, showing requirements for expert knowledge and data, interpretability, and model structure.

3 Model and control framework

We will develop a digital-twin-oriented workflow that integrates real-time sensor streams with hybrid modelling, combining mechanistic and data-driven components. The mechanistic part provides interpretability and robustness, while the data-driven part captures plant-specific dynamics and time-varying disturbances [2]. The resulting hybrid model will enable near real-time state estimation and forecasting to support closed-loop decision support and automated control actions, such as energy-efficient aeration setpoint adaptation. **(Figure 1)**.

4 Expected outcomes and transferability

The project will identify practical requirements related to sensing strategy, data quality, model architecture, and control integration. The goal is to improve resilience to fluctuations in water quality and quantity while enhancing operational efficiency. Insights from the industrial case study will be used to develop a transferable methodology that can later be evaluated in municipal WWTP settings, with an emphasis on robustness across varying influent conditions and operational constraints.

References

[1] Wang, A.-J., Li, H., He, Z., Tao, Y., Wang, H., Yang, M., Savic, D., Daigger, G.T., & Ren, N. (2024). *Digital Twins for Wastewater Treatment: A Technical Review*. Engineering, 36, 21–35. https://doi.org/10.1016/j.eng.2024.04.012

[2] Schneider, M.Y., Quaghebeur, W., Borzooei, S., Froemelt, A., Li, F., Saagi, R., Wade, M.J., Zhu, J.-J., & Torfs, E. (2022). *Hybrid modelling of water resource recovery facilities: status and opportunities*. Water Science & Technology, 85(9), 2503–2524. https://doi.org/10.2166/wst.2022.115

Acknowledgement This work was performed in the cooperation framework of Wetsus, European Centre of Excellence for Sustainable Water Technology (www.wetsus.nl). Wetsus is co-funded by the European Union, the Province of Fryslân and the Dutch Government: Ministry of Economic Affairs (SBO & PPS-l/TKI Water Technology), Ministry of Education, Culture and Science (SBO) and Ministry of Infrastructure and Water Management (National Growth Fund - UPPWATER). The authors thank the participants of the research theme "Monitoring & Quality" for the fruitful discussions and their financial support.

Overlapping Covariance Intersection: Fusion with Partial Structural Knowledge of Correlation from Multiple Sources

Leonardo Pedroso[1,2], Pedro Batista[2], W.P.M.H. (Maurice) Heemels[1]
[1]Control Systems Technology section, Eindhoven University of Technology, Eindhoven, The Netherlands
[2]Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Portugal

1 Introduction

Emerging large-scale engineering systems rely on distributed fusion for situational awareness, where agents combine noisy local sensor measurements with exchanged information to obtain fused estimates. However, at the sheer scale of these systems, tracking cross-correlations becomes infeasible, preventing the use of optimal filters. Covariance intersection (CI) tools have been devised for the past quarter century to avoid double-counting [1]. These techniques aim at fusing estimates from multiple sources whose correlation is (partially or totally) unknown. The fusion procedure is designed to minimize the worst-case uncertainty of the fused estimate under the available information. Existing CI extensions exploit limited correlation knowledge but cannot incorporate structural knowledge of correlation from multiple sources, which naturally arises in distributed fusion problems. We introduce Overlapping Covariance Intersection (OCI), a generalized CI framework that accommodates this novel information structure.

2 Outline of Results

The distributed fusion problem under the OCI information structure has not been previously addressed in the literature. It is graphically compared to the state-of-the-art in Fig. 1. It is encountered in the motivating example and it is exemplary of many cooperative localization problems for the next generation of ultra large-scale engineering systems [2], e.g., the problem in Fig. 2. In these settings, each agent has access to noisy data from multiple sensors and from communication with other agents that must be fused. It is infeasible to keep track of the covariance between all measurements, but it is feasible that partial structural knowledge about the joint estimation error covariance matrix is tracked by the agents in a distributed fusion framework.

In [3] the following conclusions were drawn. First, this problem can be expressed as a generalized CI problem. Second, it is possible to establish very mild necessary and sufficient conditions on the available information for the feasibility of the OCI problem. Third, we can restrict the problem to a parameterized family of bounds for the joint estimation error covariance matrix (which is a generalization of the family of bounds used in the basic CI setting). A solution to the restricted OCI problem is given by a semidefinite program (SDP), which is computationally tractable and suitable for real-time implementation.

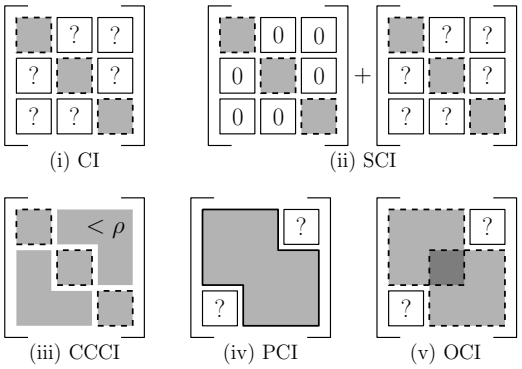

Figure 1: Comparison of the estimation error covariance matrix with three estimates for distinct partial information structures. Dashed contours represent knowledge of bounds and solid contours represent exact knowledge.

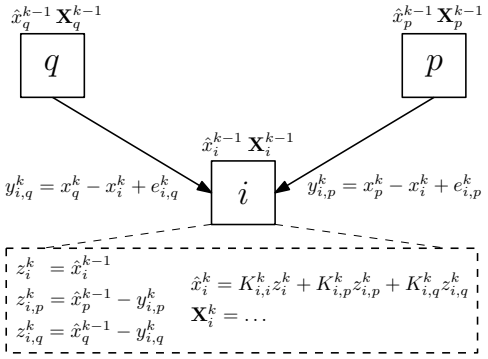

Figure 2: Scheme of the cooperative localization toy problem.

References

[1] R. Forsling, B. Noack, and G. Hendeby, "A quarter century of covariance intersection: Correlations still unknown?" *IEEE Control Systems Magazine*, vol. 44, no. 2, pp. 81–105, 2024.

[2] L. Pedroso, P. Batista, and W. P. M. H. Heemels, "Distributed design of ultra large-scale control systems: Progress, challenges, and prospects," *Annual Reviews in Control*, vol. 59, p. 100987, 2025.

[3] ——, "Overlapping covariance intersection: Fusion with partial structural knowledge of correlation from multiple sources," 2026, submitted.

Biologically Inspired Sparse Event-Based Communication and Decoding of Spatiotemporal Stimuli

Ben Efron, Guillaume Drion, Pierre Sacré, Alessio Franci
Dept. of Electrical Engineering and Computer Science, University of Liège, Belgium
Email: {befron, gdrion, p.sacre, afranci}@uliege.be

Sensory systems encode high-dimensional, continuous inputs under constraints of bandwidth, energy, and delay, posing a coding challenge: how to transmit information efficiently while preserving spatiotemporal structure. We leverage principles of sensory organization—labeled lines, receptive fields, and dual temporal pathways—to design a bio-inspired, event-based scheme that reconstructs spatiotemporal stimuli from sparse spikes. Using the tactile system as a model [1], we simulate Slowly Adapting (SA) and Rapidly Adapting (RA) neurons that encode sustained and transient components of a dynamic pressure stimulus. Both populations sample a two-dimensional receptor grid through Gaussian-weighted, overlapping receptive fields, forming a labeled-line architecture in which each neuron carries a unique, known spatial signature. SA neurons have smaller, sparser receptive fields with less overlap, while RA neurons have larger, denser receptive fields with higher overlap. Together with their distinct temporal dynamics, SA and RA neurons enable robust, efficient encoding across diverse stimuli.

Our proposed pipeline has three stages: (1) spatial compression via fixed, sparse innervation, (2) spiking for bandwidth efficiency and attention, and (3) attention-modulated Kalman filter for optimal localized reconstruction.

This event-based representation provides key advantages: bandwidth scales with active neurons rather than total sensor count, and sparsity enables reconstructing only active locations and times.

Encoding: Our framework extends the tactile encoding model of Parvizi-Fard et al. [2]. Each simulated sensory neuron samples a two-dimensional receptor grid through a Gaussian-weighted receptive field. Randomized weights produce heterogeneous, overlapping receptive fields that preserve spatial coverage and diversity. SA and RA populations apply complementary temporal filters for sustained and transient components. Filtered inputs pass through Izhikevich neuron models to generate sparse spikes. Although receptive fields are well established in neuroscience, their application to sensor communication has not been widely studied. Barranca et al. [3] explored receptive fields to reconstruct stationary images using compressed sensing. Their approach was limited to stationary images, did not exploit population dynamics, and relied on computationally expensive reconstruction.

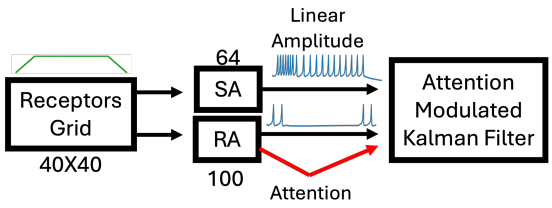

Figure 1: Three-stage pipeline: sparse innervation, event-based spiking, and RA-modulated Kalman reconstruction.

Decoding: We extend this work by leveraging neuron dynamics for high compression, accurate reconstruction of dynamical stimuli, and real-time decoding with low computational cost. We use spatial and temporal priors to stabilize reconstruction from sparse, overlapping receptive fields. SA neurons continuously measure pressure amplitude (firing rate $\propto$ intensity), enabling reconstruction of slowly varying stimuli. RA neurons detect and encode rate of change, providing fine-grained localization and triggering event-based attention. By modulating reconstruction locally where RA neurons fire, we achieve higher spatial resolution and reconstruct only active regions. RA activation dynamically adjusts the Kalman filter: measurement noise (R) is reduced for RA and increased for SA locally, prioritizing change detection, while process noise (Q) is elevated to permit rapid state changes. Combining SA and RA populations minimizes spike count while enabling localized processing across stimulus regimes. By restricting reconstruction to regions where RA neurons fire, computational complexity scales with active area rather than total grid size, enabling generalization to arbitrarily large sensor arrays. We propose an event-based SA/RA attention scheme over structured receptive fields that enables spatial compression and targeted reconstruction from sparse spikes. This strategy can be generalized to other sensor arrays and modalities by modifying priors and hidden states without altering the core receptive-field and dynamics framework.

References

[1] B. P. Delhaye et al., "Neural Basis of Touch and Proprioception in Primate Cortex," *Compr. Physiol.*, vol. 8, no. 4, pp. 1575–1602, 2018.

[2] A. Parvizi-Fard et al., "Functional spiking neuronal model for tactile perception," *J. Neurophysiol.*, vol. 126, pp. 1385–1397, 2021.

[3] J. C. Barranca et al., "Sparsity and compressed coding in sensory systems," *PLoS Comput. Biol.*, vol. 10, no. 8, e1003793, 2014.

Acknowledgment This work was supported by the Belgian Government through the Federal Public Service Policy and Support, and by the Wallonia–Brussels International (WBI) Postdoctoral Excellence Grant.

Design of Reduced Distributed Observers for Nonlinear Systems in Prime Form

Yaodong Li, Nathan van de Wouw, Erik Steur
Department of Mechanical Engineering
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven
The Netherlands
Email: (l.y.li,n.v.d.wouw,e.steur)@tue.nl

1 Introduction

In distributed state estimation, every sensor in a sensor network has the objective to reconstruct the state of a dynamical system. Each sensor should produce this estimate using only local measurements and neighbor-to-neighbor communication. In many applications such as smart grid and autonomous vehicles, no single sensor has sufficient information to observe the full state of the system, which makes information exchange essential. Compared to centralized observers, distributed observers offer improved scalability, robustness to failures, and reduced communication and computational burdens. Distributed observers are mostly developed for linear systems. However, for nonlinear systems, distributed observer design is particularly challenging due to nonlinear coupling and locally unobservable sub-state.

Existing nonlinear distributed observers typically rely on consensus dynamics to reconstruct unobservable states, leading to full-order observer designs and convergence rates that depend on the network topology. In this work, we considers nonlinear systems that admit a block-triangular observable canonical form (prime form) and develops a distributed observer design where unobservable states are reconstructed through static algebraic averaging of neighbor information. Distinguishing features of the proposed design are summarized as follows: (i) The design of each local observer is standard and simple, relying solely on locally available information. (ii) The overall order of the distributed observer is substantially reduced compared to existing designs, which are of full order. (iii) Convergence rate of the estimation (global performance) is easily tunable by setting the local observer gains.

2 Problem Formulation

Consider the nonlinear system whose state vector $x \in \mathbb{R}^n$ is partitioned into N subsets, i.e., $x = \mathrm{col}(x_1,\ldots,x_N)$ with $x_i \in \mathbb{R}^{n_i}$ and $\sum_i^N n_i = n$. The dynamics of each subsystem, with state x_i, take the following prime form [1]:

$$\dot{x}_i = A_i x_i + B_i \phi_i(x_1,\ldots,x_i), \quad y_i = H_i x_i, \quad i = 1,\ldots,N, \quad (1)$$

where the nonlinearities ϕ_i are assumed to be *globally Lipschitz*. The matrices A_i, B_i and H_i take the form:

$$A_i = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}, \quad B_i = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}, \quad H_i = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}^\top.$$

Each node has access only to its local measurement $y_i \in \mathbb{R}$. The goal is to design a distributed observer such that each local observer produces a full state estimate $\hat{x}_i \in \mathbb{R}^n$ which exponentially converges to the state x with a prescribed rate $\alpha > 0$, i.e., for some $C > 0$,

$$||\hat{x}_i(t) - x(t)|| \leq Ce^{-\alpha t} \text{ for all } i = 1,2,..,N, \quad t \geq 0. \quad (2)$$

3 Main Result

We propose the following distributed observer:

$$\dot{z}_i = A_i z_i + B_i \phi_i(\hat{x}_{i,1},\ldots,\hat{x}_{i,i-1},z_i) + \mathrm{diag}(L_i,\ldots,L_i^{n_i}) K_i (H_i z_i - y_i), \quad (3a)$$

$$\hat{x}_i = \mathrm{col}(\hat{x}_{i,1},\ldots,\hat{x}_{i,N}), \quad \text{with} \quad \hat{x}_{i,i} = z_i, \quad (3b)$$

$$\hat{x}_{i,j} = \frac{1}{N_i} \sum_{k \in \mathcal{N}_i} \hat{x}_{k,j}, \quad j \neq i, \quad (3c)$$

Equation (3a) is a standard high-gain observer for estimating the locally observable state x_i, but due to the structure of the plant (1), individual nodes generally lack full observability and cannot achieve the estimation objective independently. However, the prime form structure of (1) guarantees joint observability, implying that full-state estimation can be achieved through information exchange (3c) over a strongly connected communication graph. Consider system (1) and a strongly connected communication graph. If the observer gains K_i are chosen such that $A_i + K_i H_i$ are Hurwitz, $P_i = P_i^\top \succ 0$ solves $(A_i + K_i H_i)^\top P_i + P_i (A_i + K_i H_i) = -I_{n_i}$, and the high-gain parameters satisfy $L_i > \max(1, 2\alpha\lambda_i + 2\lambda_i c_i)$, where λ_i is the largest eigenvalue of P_i and c_i is the Lipschitz constant of ϕ_i, then objective (2) is satisfied.

References

[1] Wu, Y., Isidori, A., Lu, R. (2021). On the design of distributed observers for nonlinear systems. IEEE Transactions on Automatic Control, 67(7), 3229-3242.

Part 3

List of Participants

Mels, A.

Monica Acuautla
University of Groningen

m.i.acuautla.meneses@rug.nl

Jérôme Adriaens
University of Liege

jadriaens@uliege.be

Andrea Agazzi
University of Bern

andrea.agazzi@unibe.ch

Dirk Alferink
Eindhoven University of Technology

d.w.t.alferink@tue.nl

Mohsen Alirezaei
Siemens Digital Industries Software

mohsen.alirezaei@siemens.com

Lucas Alves
Universidade Federal de Minas Gerais

lucasvra@ufmg.br

Joshua Amato
Delft University of Technology

amatojoshua@live.nl

Duarte Antunes
Eindhoven University of Technology

d.antunes@tue.nl

Arviandy Aribowo
Eindhoven University of Technology

A.G.Aribowo@tue.nl

Zemerart Asani
Vrije Universiteit Brussel

zemerart.asani@vub.be

Jonathan Baake
Flanders Make

Gustave Bainier
University de Liège

gustave.bainier@uliege.be

Simone Baldi
Southeast University

simonebaldi@seu.edu.cn

Tren Baltussen
Eindhoven University of Technology

t.m.j.t.baltussen@tue.nl

Adrien Banse
UCLouvain

adrien.banse@uclouvain.be

Joao Carlos Basilio
Federal University of Rio de Janeiro

basilio@dee.ufrj.br

Valentin de Bassompierre
UCLouvain

valentin.debassompierre@uclouvain.be

Pedro Batista
University of Lisbon

pbatista@isr.tecnico.ulisboa.pt

Marzhan Baubekova
Eindhoven University of Technology

m.baubekova@tue.nl

Amir Bayat
UCLouvain

Amir.Bayat@UCLouvain.be

David Becker
DIFFER

d.l.c.becker@differ.nl

Martijn de Beij
DIFFER

m.p.debeij@differ.nl

Guillaume Berger
UCLouvain

guillaume.berger@uclouvain.be

Matthijs van Berkel
DIFFER

M.vanBerkel@differ.nl

Riccardo Bertollo
UCLouvain

riccardo.bertollo@uclouvain.be

Bart Besselink
University of Groningen

b.besselink@rug.nl

Ajinkya Bhole
Ghent University

ajinkya.bhole@ugent.be

Gianluca Bianchin
UCLouvain

gianluca.bianchin@uclouvain.be

François-Grégoire Bierwart
University of Namur

francois-gregoire.bierwart@unamur.be

Jef Billet
KU Leuven

jef.billet@kuleuven.be

Senne Bogaerts
KU Leuven

senne.bogaerts@kuleuven.be

Léon Borremans
Vrije Universiteit Brussel

leon.borremans@vub.be

Mathias Bos
KU Leuven

mathias.bos@kuleuven.be

Boulaid Boulkroune
Flanders Make

boulaid.boulkroune@flandersmake.be

Martijn Boussé
Maastricht University

m.bousse@maastrichtuniversity.nl

Taha Bouzid
Eindhoven University of Technology

t.bouzid@tue.nl

Julien Brandoit
University of Liège

jbrandoit@uliege.be

Léopold Van Brandt
UCLouvain

leopold.vanbrandt@uclouvain.be

Valentina Breschi
Eindhoven University of Technology

v.breschi@tue.nl

Herman Bruyninckx
KU Leuven

herman.bruyninckx@kuleuven.be

Robin Bulsink
University of Groningen

r.m.bulsink@rug.nl

Thomas Burger
KU Leuven

thomas.burger@kuleuven.be

Xinjiang Cai
University of Groningen

x.cai@rug.nl

Louis Callens
KU Leuven

louis.callens@kuleuven.be

Kanat Camlibel
University of Groningen

m.k.camlibel@rug.nl

Ming Cao
University of Groningen

m.cao@rug.nl

Stefano De Carli
University of Bergamo

stefano.decarli@unibg.it

Raffaella Carloni
University of Groningen

r.carloni@rug.nl

Alessandro Casu
Eindhoven University of Technology

a.casu@tue.nl

Alessandro Cecconi
KU Leuven

alessandro.cecconi@kuleuven.be

Lennard Ceelen
DIFFER

l.ceelen@differ.nl

Paula Chanfreut
Eindhoven University of Technology

p.chanfreut.palacio@tue.nl

Xiaodong Cheng
Wageningen University and Research

xiaodong.cheng@wur.nl

Michelle Chong
Eindhoven University of Technology

m.s.t.chong@tue.nl

Alessandro Cicalese
University of Groningen

a.cicalese@rug.nl

Radu Cimpean
University of Groningen

r.cimpean@rug.nl

Koen Classens
ASML

koen.classens@asml.com

Adeline Colson
UCLouvain

adeline.colson@uclouvain.be

Patrick Conjarts
Eindhoven University of Technology

p.m.d.conjarts@student.tue.nl

Bryan Convens
Vrije Universiteit Brussel

Bryan.Convens@vub.be

Francesco Corrini
University of Bergamo

francesco.corrini@unibg.it

Guillaume Crevecoeur
Ghent University

guillaume.crevecoeur@ugent.be

Michele Cucuzzella
University of Groningen

m.cucuzzella@rug.nl

Amritam Das
Eindhoven University of Technology

am.das@tue.nl

Janset Dasdemir
University of Twente

j.dasdemir@utwente.nl

Ruth David
Eindhoven University of Technology

r.r.k.david@tue.nl

Wilm Decré
KU Leuven

wilm.decre@kuleuven.be

Giannis Delimpaltadakis
Eindhoven University of Technology

i.delimpaltadakis@tue.nl

Jean-Charles Delvenne
UCLouvain

jean-charles.delvenne@uclouvain.be

Bruno Depraetere
Flanders Make

bruno.depraetere@flandersmake.be

Margot Devillers
UCLouvain

margot.devillers@uclouvain.be

Laurent Dewasme
Mons University

laurent.dewasme@umons.ac.be

Vijayanand Digge
UCLouvain

vijayanand.digge@uclouvain.be

Dimos Dimarogonas
KTH Royal Institute of Technology

dimos@kth.se

Dries Dirckx
KU Leuven

dries.dirckx@kuleuven.be

Nic Dirkx
ASML

nic.dirkx@asml.com

Victor Dolk
Eindhoven University of Technology

v.s.dolk@tue.nl

Sander Doodeman
Eindhoven University of Technology

s.doodeman@tue.nl

Yangming Dou
University of Groningen

y.dou@rug.nl

Enrico Dozzi
Eindhoven University of Technology

e.dozzi@tue.nl

Philippe Dreesen
Maastricht University

philippe.dreesen@maastrichtuniversity.nl

Roel Drenth
Eindhoven University of Technology

r.drenth@tue.nl

Guillaume Drion
University of Liège

gdrion@uliege.be

Yuvan Dwaraga
Eindhoven University of Technology

Felix Ebert
TU Munich

felix.ebert@tum.de

Ben Efron
The University of Liège

befron@uliege.be

Sebastiaan van den Eijnden
Eindhoven University of Technology

s.j.a.m.v.d.eijnden@tue.nl

Katrien Eloot
Flanders Hydraulics

katrien.eloot@mow.vlaanderen.be

Magne Johannes Erlandsen
Eindhoven University of Technology

m.j.erlandsen@tue.nl

Damien Ernst
University of Liège

dernst@uliege.be

Erik van den Eshof
Eindhoven University of Technology

r.c.p.v.d.eshof@tue.nl

Matthias Van Eysendeyk
KU Leuven

matthias.vaneysendeyk@kuleuven.be

Chaoyue Fei
Vrije Universiteit Brussel

chaoyue.fei@vub.be

Diego Fernandez-Zapico
Eindhoven University of Technology

d.fernandez.zapico@tue.nl

Rob Fey
Eindhoven University of Technology

r.h.b.fey@tue.nl

Paulo Figueiredo
DIFFER

p.a.marinhofigueiredo@differ.nl

Mohammad Mahmoudi Filabadi
Ghent University

mohammad.mahmoudifilabadi@ugent.be

Merijn Floren
KU Leuven

merijn.floren@kuleuven.be

Alvaro Florez
KU Leuven

alvaro.florez@kuleuven.be

Alessio Franci
University de Liège

afranci@uliege.be

Franz Fuchshumer
AGCO GmbH (Fendt)

franz.fuchshumer@agcocorp.com

Arthur Fyon
University of Liège

afyon@uliege.be

Kangping Gao
Wageningen University and Research

kangping.gao@wur.nl

Shaya Garjani
University of Groningen

s.garjani@rug.nl

Rory Gavin
University of Groningen

r.c.gavin@rug.nl

Jarne van Gemert
Eindhoven University of Technology

j.j.h.v.gemert@tue.nl

Bart Gerlagh
University of Twente

b.gerlagh@utwente.nl

Gianluca Giacomelli
Eindhoven University of Technology

g.giacomelli@tue.nl

Andrea Giusti
Flanders Make

andrea.giusti@flandersmake.be

Raphaël Goetz
Eindhoven University of Technology

r.p.p.f.goetz@tue.nl

Rodrigo Gonzalez
Eindhoven University of Technology

r.a.gonzalez@tue.nl

Mauricio Sotelo Gonzalez
Eindhoven University of Technology

m.sotelo.gonzalez1@tue.nl

Alejandro Gonzalez-Garcia
KU Leuven

alex.gonzalezgarcia@kuleuven.be

Martijn Goorden
Eindhoven University of Technology

m.a.goorden@tue.nl

Timo de Groot
Eindhoven University of Technology

t.d.groot2@tue.nl

Bendeguz Györök
Institute for Computer Science and Control

gyorokbende@sztaki.hu

Abdelrahman Habash
Wageningen University and Research

abdelrahman.habash@wur.nl

Abdullah Habboush
Eindhoven University of Technology

a.habboush@tue.nl

Taco Hacquebord
Philips

taco.Hacquebord@philips.com

Yulyan Hadi
University of Groningen

y.w.hadi@rug.nl

Soheil Hadisi
Wageningen University and Research,

soheil.hadisi@wur.nl

Sofie Haesaert
Eindhoven University of Technology

s.haesaert@tue.nl

Wouter Hakvoort
University of Twente

w.b.j.hakvoort@utwente.nl

Jonas Hansson
UCLouvain

jonas.hansson@uclouvain.be

Amirreza Haqshenas
KU Leuven

amirreza.haqshenasmojaveri@kuleuven.be

Mehran Raisi Hatmabadi
Vrije Universiteit Brussel

Mehran.Raisi.Hatmabadi@vub.be

Mohaddese Hazini
University of Twente

m.hazini@utwente.nl

Maurice Heemels
Eindhoven University of Technology

m.heemels@tue.nl

Marcel Heertjes
Eindhoven University of Technology

M.F.Heertjes@tue.nl

Marieke Heidema
University of Groningen

h.m.heidema@rug.nl

Julien Hendrickx
UCLouvain

julien.hendrickx@uclouvain.be

Jonas Hendrikx
Eindhoven University of Technology

j.g.hendrikx@tue.nl

Domagoj Herceg
Eindhoven University of Technology

d.herceg@tue.nl

Sebastian Hermann
Technical University of Munich

hermannsebastian@protonmail.com

Jan Hoekstra
Eindhoven University of Technology

j.h.hoekstra@tue.nl

Theo Hofman
Eindhoven University of Technology

t.hofman@tue.nl

Ali Hosseini
Delft University of Technology

s.a.hosseini@tudelft.nl

Hassan HosseinNia
Delft University of Technology

s.h.hosseinniakani@tudelft.nl

Binbin Hu
University of Groningen

b.hu@rug.nl

Huayuan Huang
University of Groningen

huayuan.huang@rug.nl

Thierry Hubert
Vrije Universiteit Brussel

Thierry.Rene.Hubert@vub.be

Mischa Huisman
Eindhoven University of Technology

m.r.huisman@tue.nl

Jilles van Hulst
Eindhoven University of Technology

j.s.v.hulst@tue.nl

Maarten van der Hulst
Eindhoven University of Technology

M.v.d.hulst@tue.nl

Lucian Cristian Iacob
Flanders Make

cristian.iacob@flandersmake.be

Lucian Cristian Iacob,
Flanders Make

cristian.iacob@flandersmake.be

Ibrahim Ibrahim
KU Leuven

ibrahim.ibrahim@kuleuven.be

Santiago Iregui
KU Leuven

santiago.iregui@kuleuven.be

Francesco Lo Iudice
University of Naples Federico II

francesco.loiudice@unina.it

Laurens Jacobs
Nikwist

laurens.jacobs@nikwist.be

Laurent Jacques
UCLouvain

laurent.jacques@uclouvain.be

Matin Jafarian
Delft University of Technology

m.jafarian@tudelft.nl

Marin Jalsic
Eindhoven University of Technology

m.j.jalsic@tue.nl

Loes Jansen
DIFFER

L.L.T.C.Jansen@differ.nl

Adrien Janssens
UCLouvain

adrien.janssens@uclouvain.be

Bayu Jayawardhana
University of Groningen

b.jayawardhana@rug.nl

Kevin Jehasse
University of Liège

Kevin.Jehasse@uliege.be

Ashutosh Jindal
University of Groningen

a.jindal@rug.nl

Thomas de Jong
Eindhoven University of Technology

t.o.d.jong@tue.nl

Jan de Jong
University of Twente

j.j.dejong@utwente.nl

Marissa Jonker
University of Twente

l.m.jonker@utwente.nl

Raphaël M. Jungers
UCLouvain

Raphael.Jungers@UCLouvain.be

Raphaël Jungers
UCLouvain

raphael.jungers@uclouvain.be

Fatemeh Zahed Kachaee
University of Liège

fzahed@uliege.be

Jorn van Kampen
Eindhoven University of Technology

j.h.e.v.kampen@tue.nl

Alexander Katriniok
Eindhoven University of Technology

a.katriniok@tue.nl

Yu Kawano
Hiroshima University

ykawano@hiroshima-u.ac.jp

Hussain Kazmi
KU Leuven

hussainsyed.kazmi@kuleuven.be

Jurrien Keulen
University of Groningen

j.r.keulen@rug.nl

Achim Kienle
Otto von Guericke Univ. / Max Planck Institute

achim.kienle@ovgu.de

Edward Kikken
Flanders Make

edward.kikken@flandersmake.be

Michel Kinnaert
University Libre de Bruxelles

michel.kinnaert@ulb.be

Maté Kiss
Eindhoven University of Technology

m.kiss@tue.nl

Lizan Kivits
DIFFER

e.m.m.kivits@differ.nl

Ward Klip
Eindhoven University of Technology

w.m.klip@tue.nl

Tamas Kovacs
Eötvös University

tamas.kovacs@ttk.elte.hu

Julius Krebbekx
Eindhoven University of Technology

j.p.j.krebbekx@tue.nl

Dinesh Krishnamoorthy
Norwegian University of Science and Technology

dinesh.krishnamoorthy@ntnu.no

Kevin Tom Kurian
Eindhoven University of Technology

k.tom.kurian@tue.nl

Martin Lames
Technical University of Munich

martin.lames@tum.de

Mircea Laza
Eindhoven University of Technology

m.lazar@tue.nl

Mircea Lazar
Eindhoven University of Technology

M.Lazar@tue.nl

Erjen Lefeber
Eindhoven University of Technology

A.A.J.Lefeber@tue.nl

Benoît Legat
UCLouvain

benoit.legat@uclouvain.be

Sophie Lequeu
UCLouvain

sophie.lequeu@uclouvain.be

Yanxin Li
University of Groningen

yanxin.li@rug.nl

Yaodong Li
Eindhoven University of Technology

l.y.li@tue.nl

Xinghua Liu
University of Groningen

xinghua.liu@rug.nl

Siyuan Liu
Eindhoven University of Technology

s.liu5@tue.nl

Gaoyuan Liu
Vrije Universiteit Brussel

Jingyue Liu
Delft University of Technology

J.Liu-14@tudelft.nl

Christian Fernandez Lorden
University of Liège

christian.fernandezlorden@uliege.be

Surya Malladi
University of Groningen

v.s.p.malladi@rug.nl

Zeyad Manaa
Eindhoven University of Technology

z.manaa@tue.nl

Chiara Manfletti
TU Munich

chiara.manfletti@tum.de

Mariana Guimaraes Marques
Eindhoven University of Technology

m.guimaraes.marques@tue.nl

Carlos Martinez
Maastricht University

carlos.martinezgarcia@maastrichtuniversity.nl

Norbert Marwan
Potsdam Institute for Climate Impact Research (PIK)

marwan@pik-potsdam.de

Louise Massager
University Libre de Bruxelles

louise.massager@ulb.be

Sébastien Mattenet
UCLouvain

sebastien.mattenet@uclouvain.be

Fakhira Maulana
Wageningen University and Research

fakhira.maulana@wur.nl

Alexandre Mauroy
University of Namur

alexandre.mauroy@unamur.be

Mirko Mazzoleni
University of Bergamo

mirko.mazzoleni@unibg.it

Gijs Van Meerbeeck
Eindhoven University of Technology

g.v.meerbeeck@student.tue.nl

Bas Meere
Eindhoven University of Technology

b.g.l.meere@tue.nl

Amir Mehrnoosh
UCLouvain

amir.mehrnoosh@uclouvain.be

Tomas Meijer
Lund University

tomas.meijer@lu.se

Lex Meijer
Airbus Defence and Space Netherlands

lex.meijer@airbus.com

Loris Mendolia
University de Liège

lmendolia@uliege.be

Kelly Merckaert
Vrije Universiteit Brussel

Kelly.Merckaert@vub.be

Hendrik Meth
Stuttgart Media University

meth@hdm-stuttgart.de

Daan Meulendijks
Eindhoven University of Technology

d.p.meulendijks@tue.nl

Erwan Meunier
UCLouvain

erwan.meunier@uclouvain.be

Wim Michiels
KU Leuven

wim.michiels@kuleuven.be

Marijn Minkenberg
Eindhoven University of Technology

m.minkenberg@tue.nl

Tom Minten
Eindhoven University of Technology

t.p.minten@tue.nl

Georgios Mitsos
Eindhoven University of Technology

g.mitsos@tue.nl

Mehrdad Mohammadi
Eindhoven University of Technology

m.mohammadi1@tue.nl

Judicaël Mohet
University of Namur

judicael.mohet@unamur.be

René van de Molengraft
Eindhoven University of Technology

m.j.g.v.d.Molengraft@tue.nl

Nima Monshizadeh
University of Groningen

n.monshizadeh@rug.nl

Maxime Monsieur
Flanders Make

maxime.monsieur@flandersmake.be

Yann de Mont-Marin
Inria Paris

yann.de-mont-marin@inria.fr

Louis Montaut
Inria Paris

louis.montaut@inria.fr

Adarsh Subrahamanian Moosath
Eindhoven University of Technology

a.subrahamanian.moosath@tue.nl

Arturo Morales
Eindhoven University of Technology

a.r.h.morales.bugueno@tue.nl

Rodrigo Moreno-Morton
University of Liege

lrmoreno@uliege.be

Alessio Moreschini
Imperial College London

a.moreschini@imperial.ac.uk

Branimir Mrak
Flanders Make

branimir.mrak@flandersmake.be

Junhwi Mun
Eindhoven University of Technology

j.mun@tue.nl

Arthur Mélot
UCLouvain

arthur.melot@uclouvain.be

Maedeh Izadi Najafabadi
Eindhoven University of Technology

m.izadi.najafabadi1@tue.nl

Aditya Natu
Delft University of Technology

a.m.natu@tudelft.nl

Léa Ninite
UCLouvain

lea.ninite@uclouvain.be

Hafsa Nouhi
Vrije Universiteit Brussel

Ali El Hadi Noura
Mons University

ali.noura@ucb.com

Javier Olucha
Eindhoven University of Technology

e.j.olucha.delgado@tue.nl

Mohayad Omer
Vrije Universiteit Brussel

Mohayad.Abdelmonim.Mahmoud.Omer@vub.be

Tom Oomen
Eindhoven University of Technology

T.A.E.Oomen@tue.nl

Christopher Orrico
Eindhoven University of Technology

c.a.orrico@tue.nl

Gé van Otterdijk
Eindhoven University of Technology

g.j.e.v.otterdijk@tue.nl

Pournami Padmalatha
KU Leuven

pournami.padmalatha@kuleuven.be

Faezeh Pak
Eindhoven University of Technology

f.pak@tue.nl

Panagiotis Patrinos
KU Leuven

panos.patrinos@esat.kuleuven.be

Patricia Pauli
Eindhoven University of Technology

p.d.pauli@tue.nl

Leonardo Pedroso
Eindhoven University of Technology

l.pedroso@tue.nl

Patricia Pena
Universidade Federal de Minas Gerais

ppena@ufmg.br

Uto Perra
Eindhoven University of Technology

u.perra@tue.nl

Greet Van de Perre
Vrije Universiteit Brussel

Greet.Van.de.Perre@vub.be

Elena Petri
Eindhoven University of Technology

e.petri@tue.nl

Juan Peypouquet
University of Groningen

j.g.peypouquet@rug.nl

Alessandro Piazza
Eindhoven University of Technology

a.p.piazza1@tue.nl

Violaine Piengeon
University of Namur

violaine.piengeon@unamur.be

Carlos Segura Pina
University of Mons

carlosroberto.segurapina@umons.ac.be

Andrés Pino-Santana
Mons University

AndresIgnacio.PINOSANTANA@umons.ac.be

Armin Pirastehzad
University of Groningen

a.pirastehzad@rug.nl

Jeroen Ploeg
Siemens Industry Software Netherlands B.V., Helmond

jeroen.ploeg@siemens.com

Filip Polakowski-Karol
Maastricht University

f.polakowski-karol@student.maastrichtuniversity.nl

Romain Postoyan
University de Lorraine, CNRS, CRAN, F-54000 Nancy, France

Davide Previtali
University of Bergamo

davide.previtali@unibg.it

Wagterveld, R.M

Ali Rahdarian
University Libre de Bruxelles

ali.rahdarian@ulb.be

Carlos Murguia Rendon
Eindhoven University of Technology

C.G.Murguia@tue.nl

Michel Reniers
Eindhoven University of Technology

Michaela Repasova
University of Groningen

m.repasova@rug.nl

Federico Ulloa Rios
KU Leuven

federico.ulloarios@kuleuven.be

Carlos Jose Gonzalez Rojas
Eindhoven University of Technology

c.j.gonzalez.rojas@tue.nl

Juan Javier Palacios Roman
Eindhoven University of Technology

j.j.palacios.roman@tue.nl

Thijs Romberg
Eindhoven University of Technology

t.l.romberg@tue.nl

Wesley Roozing
University of Twente

w.roozing@utwente.nl

Danilo Saccani
EPFL

danilo.saccani@epfl.ch

Pierre Sacré
University of Liège

p.sacre@uliege.be

Marco Saerans
UCLouvain

marco.saerens@uclouvain.be

Mauro Salazar
Eindhoven University of Technology

m.r.u.salazar@tue.nl

Mauro Salazar,
Eindhoven University of Technology

m.r.u.salazar@tue.nl

Tom van der Sande
Eindhoven University of Technology

t.p.j.v.d.sande@tue.nl

Edward Sandra
KU Leuven

edward.sandra@kuleuven.be

Cosimo Della Santina
Delft University of Technology

C.DellaSantina@tudelft.nl

Sonia De Santis
KU Leuven

sonia.desantis@kuleuven.be

Daniel Sarsur
Eindhoven University of Technology

d.sarsur.camara@tue.nl

Ajay Suresha Sathya
Inria Paris

ajay.sathya@inria.fr

Matteo Scandella
University of Bergamo

matteo.scandella@unibg.it

Patrick Scheffe
KU Leuven

patrick.scheffe@kuleuven.be

Koen Scheres
KU Leuven

koen.scheres@kuleuven.be

Jacquelien Scherpen
University of Groningen

j.m.a.scherpen@rug.nl

Mathias Schietecat
KU Leuven

mathias.schietecat@kuleuven.be

Maarten Schoukens
Eindhoven University of Technology

m.schoukens@tue.nl

Mathijs Schuurmans
KU Leuven

mathijs.schuurmans@esat.kuleuven.be

Jordy Senden
Eindhoven University of Technology

j.p.f.senden@tue.nl

Rodolphe Sepulchre
KU Leuven

rodolphe.sepulchre@kuleuven.be

Ebrahim ShahabiShalghouni
Delft University of Technology

E.ShahabiShalghouni@tudelft.nl

Amir Shahhosseini
KU Leuven

amir.shahhosseini@kuleuven.be

Erfan Shakhesi
Eindhoven University of Technology

e.shakhesi@tue.nl

Fahim Shakib
Eindhoven University of Technology

m.f.shakib@tue.nl

Amir Shakouri
University of Groningen

a.shakouri@rug.nl

Brayan Shali
KU Leuven

brayan.shali@kuleuven.be

Shengling Shi
Delft University of Technology

shengling.shi@tudelft.nl

Taranjitsingh Singh
Flanders Make

taranjitsingh.singh@flandersmake.be

Peter Slaets
KU Leuven

peter.slaets@kuleuven.be

Denis Steckelmacher
Vrije Universiteit Brussel

Denis.Steckelmacher@vub.be

Alef Sterk
University of Groningen

a.e.sterk@rug.nl

Erik Steur

Lucas Stordeur
University of Liège

Lucas.Stordeur@uliege.be

Ankit Surti
Flanders Make

ankit.surti@flandersmake.be

Jan Swevers
KU Leuven

jan.swevers@kuleuven.be

Matisse Teuwen
KU Leuven

matisse.teuwen@esat.kuleuven.be

Rick Thijssen
Eindhoven University of Technology

r.e.w.l.thijssen@tue.nl

Diego Tielens
KU Leuven

diego.tielens@kuleuven.be

Bavo Tistaert
KU Leuven

bavo.tistaert@kuleuven.be

Guanchun Tong
KU Leuven

guanchun.tong@kuleuven.be

Samarth Toolhally
Eindhoven University of Technology

s.toolhally@tue.nl

Kelvin Toonen
Delft University of Technology

k.a.b.toonen@tudelft.nl

Mario Valencia Torres
Delft University of Technology

M.ValenciaTorres@student.tudelft.nl

Roland Toth
Eindhoven University of Technology

r.toth@tue.nl

Stephan Trenn
University of Groningen

s.trenn@rug.nl

Edoardo Vacchini
University of Groningen

e.vacchini@rug.nl

Foeke Vanbecelaere
Flanders Make

foeke.vanbecelaere@flandersmake.be

Julien Vanderheyden
University de Liège

julien.vanderheyden@uliege.be

Martina Vanelli
UCLouvain

martina.vanelli@uclouvain.be

Hari Prasad Varadarajan
DIFFER

h.varadarajan@differ.nl

Finn Vehlhaber
Eindhoven University of Technology

f.n.vehlhaber@tue.nl

Kees Verbaan
Eindhoven University of Technology

c.a.m.verbaan@tue.nl

Kevin Verhaegh
Eindhoven University of Technology

k.h.a.verhaegh@tue.nl

Hans Vermeulen
Eindhoven University of Technology

j.p.m.b.vermeulen@tue.nl

Mats Versmissen
Eindhoven University of Technology

m.versmissen@tue.nl

Carolina Vissers
Eindhoven University of Technology

c.m.e.vissers@tue.nl

Jasper Visterin
KU Leuven

jasper.visterin@kuleuven.be

Anastasios Vlachos
University of Groningen

a.vlachos@rug.nl

Maarten Vlaswinkel
Eindhoven University of Technology

m.g.vlaswinkel@tue.nl

Robbert Voorhoeve
TNO

Brazil de Vries
University of Groningen

b.de.vries.22@student.rug.nl

Wytze de Vries
Eindhoven University of Technology

w.a.b.d.vries@student.tue.nl

Alexis Vuille
UCLouvain

alexis.vuille@uclouvain.be

Henk van Waarde
University of Groningen

h.j.van.waarde@rug.nl

Martijn Wagterveld
Wetsus, European Centre of Excellence for Sustainable Water Technology

Martijn.Wagterveld@wetsus.nl

Laurens van der Wal
Eindhoven University of Technology

l.e.v.d.wal@tue.nl

Zepeng Wang
University of Groningen

zepeng.wang@rug.nl

Ruohan Wang
Eindhoven University of Technology

r.wang2@tue.nl

Cuiling Wang
KU Leuven

cuiling.wang@kuleuven.be

Baolong Wang
Tsinghua University

wangbl@tsinghua.edu.cn

Jiwei Wang
University of Groningen

jiwei.wang@rug.nl

Sam Weckx
Flanders Make

sam.weckx@flandersmake.be

Siep Weilan
Eindhoven University of Technology

s.weiland@tue.nl

Siep Weiland
Eindhoven University of Technology

s.weiland@tue.nl

Elionai Weldeghebreal
Eindhoven University of Technology

e.weldeghebreal@tue.nl

Teun Wijfjes
Eindhoven University of Technology

t.wijfjes@student.tue.nl

Marc van der Wijst
ASML

Melvyn Wildeboer
University of Groningen

m.d.wildeboer@rug.nl

Frank Willems
Eindhoven University of Technology

F.P.T.Willems@tue.nl

Jeroen Willems
Flanders Make

bruno.depraetere@flandersmake.be

Michiel Wind
Eindhoven University of Technology

m.wind@tue.nl

Max Winkel
DIFFER

winkel@differ.nl

Joseph Winkin
University of Namur

joseph.winkin@unamur.be

Eline Wisse
Eindhoven University of Technology

e.k.wisse@tue.nl

Gert Witvoet
Eindhoven University of Technology

G.Witvoet@tue.nl

Nathan van de Wouw
Eindhoven University of Technology

Alain Vande Wouwer
Mons University

alain.vandewouwer@umons.ac.be

Robin Wroblowski
KU Leuven

robin.wroblowski@kuleuven.be

Dongjun Wu
Lund University

dongjun.wu@control.lth.se

Lufeng Xu
University of Groningen

lufeng.xu@rug.nl

Chenyang Yuan
Toyota Research Institute

ycy@mit.edu

Mohammad Javad Zallaghi
Vrije Universiteit Brussel

Mohammad.Javad.Zallaghi@vub.be

Shuhao Zhang
KU Leuven

shuhao.zhang@kuleuven.be

Yan-Yun Zhang
KU Leuven

yanyun.zhang@kuleuven.be

Zhe Zhang
Eindhoven University of Technology

z.zhang4@tue.nl

Menno van Zutphen
Eindhoven University of Technology

m.j.t.c.v.zutphen@tue.nl

Leyla Özkan
Eindhoven University of Technology

L.Ozkan@tue.nl

Welcome

The Organizing Committee has the pleasure of welcoming you to the *45th Benelux Meeting on Systems and Control*, Center Parcs De Vossemeren, Elzen 145, 3920 Lommel, Belgium.

Aim

The aim of the Benelux Meeting is to promote research activities and to enhance cooperation between researchers in Systems and Control. This is the forty-fifth in a series of annual conferences that are held alternately in Belgium and The Netherlands.

Scientific Program Overview

1. One plenary lecture by *Florian Dörfler* (Eidgenössische Technische Hochschule Zürich, Switzerland) on **Data-Enabled Predictive Control (DeePC)**.
2. One plenary lecture by *Alessandro Abate* (University of Oxford, UK) on **Bits and bobs: sound synthesis with logic and data**.
3. Two plenary lectures (Mini Course) by *Maryam Kamgarpour* (École Polytechnique Fédérale de Lausanne, Switzerland) on **Decentralized Learning of Equilibria in Multiagent Games**.
4. Contributed short lectures. See the list of sessions for the titles and authors of these lectures.

Directions for speakers

For a contributed lecture, the available time is 20 minutes. Please leave a few minutes for discussion and room changes, and adhere to the indicated schedule. In each room projection screen are available, as well as HDMI cables. When using a projector, you have to provide a notebook yourself.

Registration

The registration desk is located in the check-in building at the entrance of the park. Upon arrival, participants will receive an envelope containing their bracelet (room key). Conference badge and certificate of participation will be available in the lobby of the business center; please pick up an envelope with your name on it.

Cottages and hotel rooms will be available from 15:00 on Tuesday.

Late registrations can be made when space is still available, please contact Dora Ourives Sereno at `dora.ourives@uclouvain.be` and `committee2026@beneluxmeeting.be`.

Arrangement	Early	Late	On-site
Accommodation (single)	600	700	725
Accommodation (shared)	450	550	575
only meals (no dinners)	300	375	400
one day (no dinner)	200	275	300

The registration fee includes:

- Admission to all sessions.
- Coffee and tea during the breaks.
- In the case of accommodation arrangement: lunch and dinner on Tuesday; breakfast, lunch, dinner on Wednesday; breakfast and lunch on Thursday.

- In the case of only meals (no dinner) arrangement: lunch on Tuesday, Wednesday, Thursday.
- In the case of one day (no dinner) arrangement: lunch on Tuesday, or Wednesday, or Thursday.
- Free use of a wireless Internet connection is available via KNP and there is no specific code.

PS: - Please note that drinks during dinner are included and that drinks after dinner are not included.

- Participants with allergies to gluten, cashews, or hazelnuts may contact the restaurant staff to request an adapted meal.

- You are all invited to the Benelux Networking Drink on Wednesday at 18.30 at the Fagus room.

Organization

The meeting has been organized by

- Gianluca Bianchin (UCLouvain)
- Julien Hendrickx (UCLouvain)
- Raphaël Jungers (UCLouvain)

with the administrative support of John Aoga (UCLouvain), Dora Ourives Sereno (UCLouvain) and Julie Vandekerckhove (UCLouvain), and the technical/website support of Erjen Lefeber (Eindhoven University of Technology).

The Advisory Committee of the 45th Benelux Meeting consists of

- Gianluca Bianchin (UCLouvain)
- Ming Cao (University of Groningen)
- Rodolphe Sepulchre (Katholieke Universiteit Leuven)
- Bart De Schutter (Delft University of Technology)
- Dennis Dochain (UCLouvain)
- Jorge Goncalves (University of Luxemburg)
- Maurice Heemels (Eindhoven University of Technology)
- Julien Hendrickx (UCLouvain)
- Bayu Jayawardhana (University of Groningen)
- Raphaël Jungers (UCLouvain)
- Karel Keesman (Wageningen University)
- Michel Kinnaert (Université libre de Bruxelles)
- Henk Nijmeijer (Eindhoven University of Technology)
- Tom Oomen (Eindhoven University of Technology)
- Hans Stigter (Wageningen University)
- Siep Weiland (Eindhoven University of Technology)
- Joseph Winkin (Université de Namur)
- Nathan van de Wouw (Eindhoven University of Technology)
- Hans J. Zwart (University of Twente)

Sponsor

The meeting is supported by UCLouvain.

Conference location

The lecture rooms of the Centre Parcs are situated in the business center. Consult the map at the end of this booklet to locate rooms. Breakfast will be served between 7:30 and 8:30 AM. Room keys (bracelet) can be picked up at the reception as of Tuesday. Rooms need to be free by 10:00 am on Thursday and the bracelet need to be returned before your departure. Parking is free on-site, subject to availability.

The address of the Center Parcs De Vossemeren is:

Elzen 145
3920 Lommel
Belgium
Tel. +32 (0)11 54 82 00

Website: Center Parcs De Vossemeren

Best junior presentation award

Continuing a tradition that started in 1996, the 45th Benelux Meeting will close with the announcement of the winner of the Best Junior Presentation Award. This award is given for the best presentation, given by a junior researcher, and it consists of a trophy that may be kept for one year and a certificate. The award is specifically given for quality of presentation rather than quality of research, which is judged in a different way.

At the meeting, the chairs of sessions will ask three volunteers in the audience to fill out an evaluation form. After the session, the evaluation forms will be collected by the Prize Commissioners who will then compute a ranking. The winner will be announced on Thursday, March 26, in room Fagus, 13:15–14:15. The evaluation forms of each presentation will be returned to the junior researcher who gave the presentation.

The Prize Commissioners are Dominik Bongartz (KULeuven, Belgium), Guillaume Drion (ULiège, Belgium), Mircea Lazar (Eindhoven University of Technology, The Netherlands).

The organizing committee counts on the cooperation of the participants to make this contest a success.

Website

An electronic version of the Book of Abstracts can be downloaded from the Benelux Meeting web site.

The mobile version of the program is available here.

Meeting

The Benelux 2027 organization team is scheduled on Wednseday, March 25, business center Lobby, 15:30–16:30.

The DISC Board meeting is scheduled on Wednseday, March 25, Betula room, 18:00–20:00.

The DISC Staff meeting is scheduled on Wednseday, March 25, Salix room, 20:00–21:00.

The meeting IFAC WC2029 is scheduled on Wednseday, March 25, Salix room, 21:00–22:00.

DISC certificates, DISC Best Thesis Award and SOCN Best Thesis Award

The ceremony for the distribution of the DISC certificates, the DISC Best Thesis Award and for the SOCN Best Thesis Award will be held on Thursday, March 26, room Fagus, 13:15–14:15.

The DISC certificates will be given by Bayu Jayawardhana, scientific director of DISC.

The jury of the DISC Best Thesis Award is formed by Leyla Özkan (Eindhoven University of Technology, The Netherlands), Ming Cao (University of Groningen, The Netherlands) and Tamas Keviczky (Delft University of Technology).

The jury of the SOCN Best Thesis Award is formed by Alexandre Mauroy (Unamur), Masoud Ahookhosh (Uantwerpen), Mariya Ishteva (KULeuven), Pierre Sacré (Uliege).

Tuesday, March 24, 2026

10:15 – 11:00	Fagus Room *Welcome Coffee*						
11:00 – 11:15	Fagus Room *Opening Ceremony*						
11:15 – 12:15	Plenary - Fagus Room Florian Dörfler – *Data-Enabled Predictive Control (DeePC)*						
12:15 – 13:15	Lunch						
Room TueA	Fagus TueA01 *Neuromorphic systems and control I*	Salix TueA02 *Machine Learning I*	Populus TueA03 *System Identification I*	Prunus TueA04 *Data-Driven Control I*	Hotel Restaurant TueA05 *Optimization I*	Alnus TueA06 *Predictive Control I*	Betula TueA07 *Hybrid Systems I*
13.15 – 13.35	Brandoit	Versmissen	Kivits	Shakouri	Jalsic	Wisse	Pirastehzad
13.35 – 13.55	Jansen	Sandra	Toolhally	Repasova	Bulsink	Boulkroune	Padmalatha
13.55 – 14.15	Scheres	Bassompierre	Drenth	Schietecat	Fernandez-Zapico	Haqshenas	Groot
14.15 – 14.35	Tielens	Pauli	Beij	Wijfjes	Colson	Rojas	Gao
14.35 – 14.55	Brandt	Digge	Habboush	Giusti	Goetz	Mélot	Erlandsen
14.55 – 15.15		Noura	Dozzi	Malladi	Lequeu	Baltussen	Gemert
15.15 – 15.35						Varadarajan	
	Break						
Room TueA	Fagus TueA08 *Neuromorphic systems and control II*	Salix TueA09 *Control Theory and Systems Engineering*	Populus TueA10 *Energy Systems I*	Prunus TueA11 *Robotics I*	Alnus TueA12 *Automated Vehicles I*	Hotel Restaurant TueA13 *Optimal Control I*	Betula TueA14 *Control Applications and Discrete Event*
16.00 – 16.20	Stordeur	Cimpean	Becker	Gerlagh	David	Piengeon	Minkenberg
16.20 – 16.40	Petri	Bainier	Najafabadi	Iregui	Billet	Pak	Marques
16.40 – 17.00	Shahhosseini	Maulana	Shakhesi	Zallaghi	Mun	Ebert	Hulst
17.00 – 17.20	Klip	Bergman		Sathya		Bayat	Weldeghebreal
17:30 – 18:30	Fagus Room Mauro Salazar – *How to Give a Scientific Talk?*						

Wednesday, March 25, 2026

08:30 – 09:30	Plenary – Fagus Room Alessandro Abate – *Bits and bobs: sound synthesis with logic and data*						
Room WedM	Fagus WedM01 *Neuromorphic systems and control III*	Salix WedM02 *Machine Learning II*	Populus WedM03 *System Identification II*	Prunus WedM04 *Data-Driven Control II*	Hotel Restaurant WedM05 *Optimization II*	Alnus WedM06 *Predictive Control II*	Betula WedM07 *Hybrid Systems II*
10.00 – 10.20	Vanderheyden	Zutphen	Otterdijk	Xu	Liu	Natu	Ninite
10.20 – 10.40	Heidema	Eysendeyk	Hoekstra	Huang	Mehrnoosh	Vlachos	Mitsos
10.40 – 11.00	Shali	Bogaerts	Alferink	Jong	Perra	Minten	Mattenet
11.00 – 11.20	Petri	Iacob	Winkel	Vuille	Borremans	Willems	Habash
11.20 – 11.40	Wroblowski	Kampen	Kiss	Wind	Wang	Hulst	Wang
11.40 – 12.00	Cecconi	Pino-Santana	Pina	Vanelli	Callens	Orrico	Aribowo
12.00 – 12.20			Moosath	Pei	Meunier		
12:20 – 13:15	Lunch						
Room WedA	Salix WedA01 *Neuromorphic systems and control IV*	Populus WedA02 *Energy Systems II*	Prunus WedA03 *Robotics II*	Alnus WedA04 *Automated Vehicles II*	Hotel Restaurant WedA05 *Optimal Control II*	Betula WedA06 *Discrete Event Systems*	Fagus WedA07 *Motion Planning and Control I*
13.30 – 13.50	Fyon	Monsieur	Dirckx	Zhang	Wu	Wal	Tistaert
13.50 – 14.10	Tong	Rahdarian	Asani	Huisman	Doodeman	Sarsur	Hatmabadi
14.10 – 14.30	Sepulchre	Wang	Hazini	Zhang	Giusti	Zhang	Bos
14.30 – 14.50	Lorden	Singh	Meulendijks	Conjarts	Iacob	Baubekova	Gonzalez-Garcia
14.50 – 15.10	Mendolia	Cicalese	Fei	Bouzid	Florez	Ceelen	Vissers
15.10 – 15.30		Cai	Hadi	Kurian	Roman	Devillers	Santis
	Break						
Room WedA	Betula WedA08 *Recurrence Analysis: Methods and Applications*	Fagus WedA09 *Motion Planning and Control II*	Salix WedA10 *Transportation Systems*	Populus WedA11 *Nonlinear Systems I*	Prunus WedA12 *Mechanisms, Design and Control*	Hotel Restaurant WedA13 *Model Reduction*	Alnus WedA14 *Planning and Control under Uncertainty*
16.00 – 16.20	Martinez	Carli	Vehlhaber	Mohet	Hosseini	Wildeboer	Janssens

16.20 – 16.40	Polakowski-Karol	Hu	Garjani	Krebbekx	Thijssen	Manaa	Bhole
16.40 – 17.00	Kovacs	Massager	Eshof	Vacchini	Jindal	Dou	Keulen
17.00 – 17.20	Marwan	Scheffe	Vries	Olucha	Nouhi	Liu	Vanbecelaere
17.20 – 17.40	Hermann	Jonker	Salazar	Bertollo	Romberg	Hendrikx	Li
17.40 – 18.00			Visterin	Bierwart	Gonzalez	Corrini	Teuwen
18.00 – 18.20			Meerbeeck	Kachaee	Pedroso	Casu	
18:30 – 19:30	Fagus Room *Benelux Networking Drink*						

Thursday, March 26, 2026

08:30 – 10:45	Mini Course – Fagus Room Maryam Kamgarpour – *Decentralized Learning of Equilibria in Multiagent Games*					
	Break					
Room ThuM	Populus ThuM01 *Nonlinear Systems II*	Hotel Restaurant ThuM02 *Adaptive Systems*	Prunus ThuM03 *Data-Driven Control III*	Fagus ThuM04 *Neuromorphic systems and control V*	Alnus ThuM05 *Reinforcement Learning and applications*	Betula ThuM06 *Distributed Control and Estimation*
11.15 – 11.35	Giacomelli	Gavin	Wang	Burger	Toonen	Pedroso
11.35 – 11.55	Hansson	Morales	Torres	Adriaens	Ibrahim	Efron
11.55 – 12.15	Floren		Jafarian	Moreno-Morton	Hadisi	Li
12:15 – 13:15	Lunch					
13:15 – 14:15	Fagus Room *Awards Ceremony*					
14:15 – 14:30	Fagus Room *Closing Ceremony*					

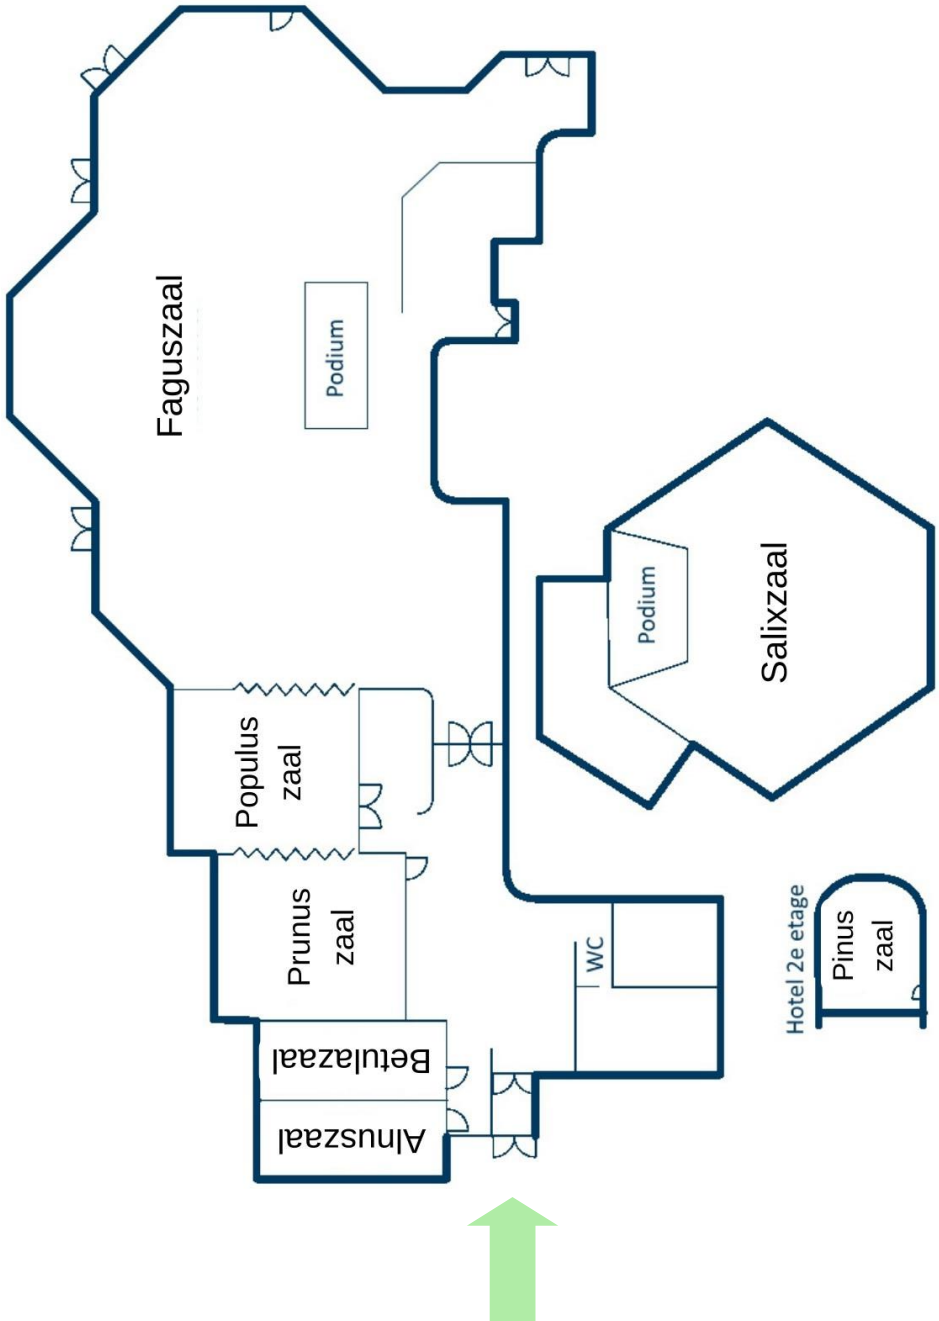
Faguszaal
Podium
Salixzaal
Podium
Populus zaal
Prunus zaal
WC
Betulazaal
Alnuszaal
Hotel 2e etage
Pinus zaal