

Multi-agent reinforcement learning via distributed MPC as a function approximator

Mallick, Samuel; Airal, Filippo; Dabiri, Azita; De Schutter, Bart

DOI

[10.1016/j.automatica.2024.111803](https://doi.org/10.1016/j.automatica.2024.111803)

Publication date

2024

Document Version

Final published version

Published in

Automatica

Citation (APA)

Mallick, S., Airal, F., Dabiri, A., & De Schutter, B. (2024). Multi-agent reinforcement learning via distributed MPC as a function approximator. *Automatica*, *167*, Article 111803. <https://doi.org/10.1016/j.automatica.2024.111803>

Important note

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

Copyright

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.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Brief paper

Multi-agent reinforcement learning via distributed MPC as a function approximator[☆]Samuel Mallick^{*}, Filippo Airaldi, Azita Dabiri, Bart De Schutter

Delft Center for Systems and Control, Delft University of Technology, Delft, The Netherlands

ARTICLE INFO

Article history:

Received 8 December 2023

Received in revised form 16 April 2024

Accepted 23 May 2024

Available online xxxx

Keywords:

Multi-agent reinforcement learning
Distributed model predictive control
Networked systems
ADMM

ABSTRACT

This paper presents a novel approach to multi-agent reinforcement learning (RL) for linear systems with convex polytopic constraints. Existing work on RL has demonstrated the use of model predictive control (MPC) as a function approximator for the policy and value functions. The current paper is the first work to extend this idea to the multi-agent setting. We propose the use of a distributed MPC scheme as a function approximator, with a structure allowing for distributed learning and deployment. We then show that Q-learning updates can be performed distributively without introducing nonstationarity, by reconstructing a centralized learning update. The effectiveness of the approach is demonstrated on a numerical example.

© 2024 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

1. Introduction

Reinforcement learning (RL) (Sutton & Barto, 2018) has proven to be a popular approach for control of complex processes. For large or continuous state and action spaces, function approximators are commonly used to learn representations of the policy. Deep neural networks (DNNs) (Arulkumaran, Deisenroth, Brundage, & Bharath, 2017) are a prevalent choice in this context; however, they often lack interpretability and are not conducive to safety verification, resulting in traditional DNN-based RL not yet widely being accepted in the control community. Alternatively, model predictive control (MPC) is an extremely successful control paradigm (Borrelli, Bemporad, & Morari, 2016), involving solving a finite-horizon optimal control problem in a receding horizon fashion. Extensive results on the stability and performance of MPC exist (Mayne, Rawlings, Rao, & Scokaert, 2000). However, MPC is entirely model-based, with its performance depending on an accurate system model.

The integration of MPC and RL is a promising direction for achieving safe and interpretable learning-based control (Hewing, Wabersich, Menner, & Zeilinger, 2020; Rosolia & Borrelli,

2018). In particular, MPC has been proposed as a replacement for DNN function approximators in RL (Gros & Zanon, 2020). In this context, the optimal control action and cost of the MPC optimization problem represent the policy and value function respectively. An MPC-based policy facilitates the use of the rich theory underlying MPC for obtaining insights into the policy, and allows to deliver certificates on the resulting behavior. The state of the art (Airaldi, De Schutter, & A., 2023; Gros & Zanon, 2020, 2022), however, relies on a centralized approach with a single learning agent. This is in general prohibitive for multi-agent systems, where centralization requires either a specific topology, with all agents connected to the central agent, or multi-hop communication across intermediate agents, where the number of hops grows with the network size. Additionally, centralized computation can become too complex, and requires the sharing of sensitive information, such as objective functions, with the central agent.

Addressing these challenges, distributed control of multi-agent systems offers computational scalability and privacy, with only neighbor-to-neighbor communication. Many existing works have adapted the MPC methodology to the distributed setting with distributed MPC (Maestre & Negenborn, 2014) and, likewise, RL to the multi-agent setting RL (MARL) setting, in which multiple learning agents act in a common environment. A central challenge in MARL is that the interaction of simultaneously learning agents renders the learning target of each agent nonstationary, degrading the performance and convergence properties of single-agent algorithms (Busoniu, Babuska, & De Schutter, 2008). Several works have tried to circumvent this problem using a centralized training and decentralized execution paradigm (Alqahtani, Scott, & Hu, 2022; Lowe, Wu, Tamar, Harb, Abbeel, & Mordatch,

[☆] This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101018826- CLarinet). The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Luca Schenato under the direction of Editor Christos G. Cassandras.

^{*} Corresponding author.

E-mail addresses: s.h.mallick@tudelft.nl (S. Mallick), f.airaldi@tudelft.nl (F. Airaldi), a.dabiri@tudelft.nl (A. Dabiri), b.deschutter@tudelft.nl (B. De Schutter).

2017). However, centralized training is often either unrealistic or unavailable. Some approaches address the nonstationarity issue through communication across the network of agents during learning (Suttle, Yang, Zhang, Wang, Basar, & Liu, 2019; Wai, Yang, Wang, & Hong, 2018; Zhang, Yang, Liu, Zhang, & Basar, 2018). These works provide theoretical convergence guarantees, but focus on linear function approximation. MARL has also been addressed with DNN function approximators (Foerster et al., 2017; Gupta, Egorov, & Kochenderfer, 2017); however, these approaches do not emphasize information exchange between agents, and suffer from the same drawbacks as DNN-based single-agent RL. Addressing the nonstationarity of learning targets in MARL remains an open challenge.

This paper proposes the following contributions. The use of MPC as a function approximator in RL is extended to the multi-agent setting. To this end we propose a structured convex distributed MPC scheme as an approximator for the policy and value functions, introducing a novel, model-based MARL approach for linear systems, free from nonstationary. The method is distributed in training *and* deployment, with data sharing only between neighbors, irrespective of the network size and topology, thus avoiding centralized computation and multi-hop data communication. Furthermore, privacy of sensitive information in local parameters and functions is preserved, with only state trajectories being shared, in contrast to a centralized approach where local functions are shared with the central agent. Thanks to the MPC-based approximation, insights into the policy can be gained from the learned components, e.g., the prediction model and constraints. Additionally, in contrast to DNN-based approaches, it is possible to inject *a priori* information, e.g., model approximations. Furthermore, we prove a result for consensus optimization; relating the dual variables recovered distributively through the alternating direction method of multipliers (ADMM) to the optimal dual variables of the original problem, that enables the distributed learning.

The paper is structured as follows. Section 2 provides the problem description and background theory. In Section 3 we present a result on the dual variables in ADMM, which will be used later on in the paper. In Section 4 we introduce the structured distributed MPC function approximator. In Section 5 we propose Q-learning as the learning algorithm, and show how the parameter updates can be performed distributively. Section 6 gives an illustrative example.

2. Preliminaries and background

2.1. Notation

Define the index sets $\mathcal{M} = \{1, \dots, M\}$ and $\mathcal{K} = \{0, \dots, N-1\}$. The symbols t , k , and τ represent time steps in an RL context, an MPC context, and iterations of an algorithm, respectively. The symbols s and a refer to states and actions in an RL context, while x and u are used for MPC. We use bold variables to gather variables over a prediction window, e.g., $\mathbf{u} = (u^\top(0), \dots, u^\top(N-1))^\top$ and $\mathbf{x} = (x^\top(0), \dots, x^\top(N))^\top$. A vector stacking the vectors x_i , $i \in \mathcal{M}$, in one column is denoted $\text{col}_{i \in \mathcal{M}}(x_i)$. The term *local* describes components known only by the corresponding agent. For simplicity we write $(x^\tau)^\top$ as $x^{\tau, \top}$.

2.2. Problem description

We define a multi-agent Markov decision process (MDP) for M agents as the tuple $(\mathcal{S}, \{A_i\}_{i \in \mathcal{M}}, P, \{L_i\}_{i \in \mathcal{M}}, \mathcal{G})$. The set \mathcal{S} is the global state set, composed of the local state sets for each agent $\mathcal{S} = \mathcal{S}_1 \times \dots \times \mathcal{S}_M$. Moreover, A_i is the local action set and L_i is the local cost function for agent i , while P describes

the state transition dynamics for the whole system. The graph $\mathcal{G} = (\mathcal{M}, \mathcal{E})$ defines a coupling topology between agents in the network, where edges \mathcal{E} are ordered pairs (i, j) indicating that agent i may affect the cost and state transition of agent j . Define the neighborhood of agent i as $\mathcal{N}_i = \{j \in \mathcal{M} | (j, i) \in \mathcal{E}, i \neq j\}$. Note that an agent is not in its own neighborhood, i.e., $(i, i) \notin \mathcal{E}$. We assume the graph \mathcal{G} is connected. Additionally, agents i and j can communicate if $i \in \mathcal{N}_j$ or $j \in \mathcal{N}_i$.

We consider agents to be linear dynamical systems with state $s_i \in \mathcal{S}_i \subseteq \mathbb{R}^n$ and control input $a_i \in \mathcal{A}_i \subseteq \mathbb{R}^m$. The true dynamics P of the network are assumed to be unknown, and we introduce an approximation of the dynamics for agent i , parametrized by a local parameter θ_i , as

$$\begin{aligned} s_i(t+1) &= f_{\theta_i}(s_i(t), a_i(t), \{s_j(t)\}_{j \in \mathcal{N}_i}) \\ &= A_{i, \theta_i} s_i(t) + B_{i, \theta_i} a_i(t) + \sum_{j \in \mathcal{N}_i} A_{ij, \theta_i} s_j(t) + b_{\theta_i}, \end{aligned} \quad (1)$$

with $b_{\theta_i} \in \mathbb{R}^n$ a constant offset allowing (1) to capture affine relationships.

We consider a co-operative RL setting. At time step t , agent i observes its own state $s_{i,t} \in \mathcal{S}_i$ and takes an action with a local policy parametrized by the local parameter θ_i ; $\pi_{\theta_i}(s_{i,t}) = a_{i,t}$, observing the incurred local cost $L_{i,t}$ and next state $s_{i,t+1}$. The cooperative goal is to minimize, by modifying the parameters θ_i , the discounted cost over an infinite horizon

$$J(\{\pi_{\theta_i}\}_{i \in \mathcal{M}}) = \mathbb{E} \left(\sum_{t=0}^{\infty} \gamma^t L_t \right), \quad (2)$$

with $L_t = \frac{1}{M} \sum_{i \in \mathcal{M}} L_{i,t}$ the average of the agents' local costs, and $\gamma \in (0, 1]$. We define the global parametrization as $\theta = (\theta_1^\top, \dots, \theta_M^\top)^\top$. The joint policy, parametrized by θ , is then $\pi_\theta(s) = \{\pi_{\theta_i}(s_i)\}_{i \in \mathcal{M}}$. The joint action $a = \{a_i\}_{i \in \mathcal{M}}$ is generated by the joint policy; $\pi_\theta(s) = a$, where s is the joint state $s = \{s_i\}_{i \in \mathcal{M}}$. Note that the local policy π_{θ_i} is parameterized with the same parameter θ_i as the dynamics f_{θ_i} because the policy generates an action via an optimization problem in which these dynamics form equality constraints (see Section 4).

2.3. Consensus optimization

Section 4 shows that evaluation of the proposed MPC-based policy and value functions can be posed as a consensus optimization problem and solved using the ADMM and global average consensus (GAC) algorithms. This section provides the relevant background.

ADMM: ADMM solves problems of the form

$$\min_{x \in \mathcal{X}, z \in \mathcal{Z}} \{f_{\text{ADMM}}(x) + g_{\text{ADMM}}(z) : Ax + Bz = c\} \quad (3)$$

by alternating between minimization of the augmented Lagrangian, split over x and z , and maximization of the result with respect to the multipliers y as:

$$x^{\tau+1} = \underset{x \in \mathcal{X}}{\text{argmin}} \mathcal{L}(x, z^\tau, y^\tau) \quad (4a)$$

$$z^{\tau+1} = \underset{z \in \mathcal{Z}}{\text{argmin}} \mathcal{L}(x^{\tau+1}, z, y^\tau) \quad (4b)$$

$$y^{\tau+1} = y^\tau + \rho(Ax^{\tau+1} + Bz^{\tau+1} - c), \quad (4c)$$

with y the Lagrange multipliers, $\rho > 0$, and $\mathcal{L}(x, z, y) = f_{\text{ADMM}}(x) + g_{\text{ADMM}}(z) + y^\top(Ax + Bz - c) + \frac{\rho}{2} \|Ax + Bz - c\|_2^2$ the augmented Lagrangian. We have the following convergence result:

Proposition 1 (Mota, Xavier, Aguiar, & Püschel, 2018). *Assume the optimal solution set of (3) is nonempty and has optimal objective P^* , the functions f_{ADMM} and g_{ADMM} are convex, \mathcal{X} and \mathcal{Z} are convex*

polytopic sets, and A and B are full column rank. Then, $f_{\text{ADMM}}(x^\tau) + g_{\text{ADMM}}(z^\tau) \rightarrow P^*$ as $\tau \rightarrow \infty$. Additionally, $\{(x^\tau, z^\tau)\}_{\tau=1}^\infty$ has a single limit point (x^*, z^*) , which solves (3).

Consider the following optimization problem defined over the graph \mathcal{G} :

$$\min_{x_1, \dots, x_M} \sum_{i \in \mathcal{M}} F_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) \quad (5a)$$

$$\text{s.t. } h_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) \leq 0, \quad i \in \mathcal{M} \quad (5b)$$

$$g_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) = 0, \quad i \in \mathcal{M}, \quad (5c)$$

where F_i , h_i , and g_i are convex and local to agent i , along with its state $x_i \in \mathbb{R}^n$. To solve this problem distributively we introduce the augmented state $\tilde{x}_i = (x_i^\top, \text{col}_{j \in \mathcal{N}_i}^\top(x_j^{(i)}))^\top$ for agent i , where $x_j^{(i)}$ is a local copy of agent j 's state. We then introduce a global copy of each state $z = (z_1^\top, z_2^\top, \dots, z_M^\top)^\top$, with z_i corresponding to x_i . The relevant portion of the global copies for agent i is $\tilde{z}_i = (z_i^\top, \text{col}_{j \in \mathcal{N}_i}^\top(z_j))^\top$. Define the local feasible sets for the augmented states as

$$\tilde{\mathcal{X}}_i = \{\tilde{x}_i | h_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) \leq 0, g_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) = 0\}. \quad (6)$$

Problem (5) can then be reformulated with the addition of a redundant constraint

$$\min_{\{\tilde{x}_i \in \tilde{\mathcal{X}}_i\}_{i \in \mathcal{M}}, z} \sum_{i \in \mathcal{M}} F_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}) \quad (7)$$

$$\text{s.t. } \tilde{x}_i - \tilde{z}_i = 0, \quad i \in \mathcal{M},$$

which is a particular instance of (3) satisfying the assumptions in Proposition 1 (See Appendix A). The steps in (4), when applied to (7), reduce to Boyd (2010, Section 7.2):

$$\tilde{x}_i^{\tau+1} = \underset{\tilde{x}_i \in \tilde{\mathcal{X}}_i}{\text{argmin}} F_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}) + y_i^{\tau, \top} (\tilde{x}_i) + \frac{\rho}{2} \|\tilde{x}_i - \tilde{z}_i^\tau\|_2^2, \quad i \in \mathcal{M} \quad (8a)$$

$$z_i^{\tau+1} = \frac{1}{|\mathcal{N}_i| + 1} (x_i^{\tau+1} + \sum_{j \in \mathcal{N}_i} x_j^{(j), \tau+1}), \quad i \in \mathcal{M} \quad (8b)$$

$$y_i^{\tau+1} = y_i^\tau + \rho(\tilde{x}_i^{\tau+1} - \tilde{z}_i^{\tau+1}), \quad i \in \mathcal{M}. \quad (8c)$$

This is a distributed procedure as each step decouples over the agents, and uses only local and neighboring information. By Proposition 1, at convergence of ADMM the local variable \tilde{x}_i for each agent will contain the minimizers x_i^* of the original problem (5), and copies of the minimizers for neighboring agents $\{x_j^{(i),*}\}_{j \in \mathcal{N}_i}$.

GAC: The GAC algorithm allows a network of agents to agree on the average value of local variables $v_{i,0} \in \mathbb{R}$, $i \in \mathcal{M}$, communicating over the graph \mathcal{G} . For each agent, the algorithm updates values as $v_i^{\tau+1} = \mathbf{P}(i, i)v_i^\tau + \sum_{j|(i,j) \in \mathcal{E}} \mathbf{P}(i, j)v_j^\tau$, where $\mathbf{P} \in \mathbb{R}_+^{M \times M}$ is a doubly stochastic matrix, i.e., entries in each row and column sum to 1. The iterates converge as $\lim_{\tau \rightarrow \infty} v_i^\tau = M^{-1} \sum_{i \in \mathcal{M}} v_{i,0}$, $i \in \mathcal{M}$, with M the number of agents (Olfati-Saber, Fax, & Murray, 2007).

3. Local recovery of optimal dual variables from ADMM

In this section we provide a result linking the dual variables from the local minimization step (8a) to a subset of the dual variables in the original problem (5). This will be used later to construct the distributed learning update. The Lagrangian of (5) is

$$\mathcal{L}(\{x_i\}_{i \in \mathcal{M}}, \{\lambda_i\}_{i \in \mathcal{M}}, \{\mu_i\}_{i \in \mathcal{M}}) = \sum_{i \in \mathcal{M}} (f_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \lambda_i^\top h_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \mu_i^\top g_i(x_i, \{x_j\}_{j \in \mathcal{N}_i})), \quad (9)$$

where λ_i and μ_i are the multipliers associated with the inequality and equality constraints, for each agent. We stress that these multipliers are related to the local constraints, and differ from the consensus multipliers $\{y_i\}_{i \in \mathcal{M}}$ used in the ADMM iterations. Denote the optimal multipliers as $(\{\lambda_i^*\}_{i \in \mathcal{M}}, \{\mu_i^*\}_{i \in \mathcal{M}})$. At iteration τ of ADMM, the Lagrangian of the local minimization step (8a) for agent i is

$$\begin{aligned} \mathcal{L}_i^\tau(x_i, \lambda_i^\tau, \mu_i^\tau) &= f_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}) \\ &+ \lambda_i^{\tau, \top} h_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}) + \mu_i^{\tau, \top} g_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}) \\ &+ y_i^{\tau, \top} (\tilde{x}_i) + \frac{\rho}{2} \|\tilde{x}_i - \tilde{z}_i^\tau\|_2^2. \end{aligned} \quad (10)$$

Denote the optimal multipliers for iteration τ as $(\lambda_i^{*, \tau}, \mu_i^{*, \tau})$.

Proposition 2. Assume that the assumptions in Proposition 1 hold for problem (5). Additionally assume that the functions F_i are strictly convex. Then, at convergence of ADMM, the optimal multipliers from the local minimizations (8a) converge to the corresponding subset of optimal multipliers for the original problem (5), i.e.,

$$(\lambda_i^{*, \tau}, \mu_i^{*, \tau}) \rightarrow (\lambda_i^*, \mu_i^*), \quad i \in \mathcal{M}, \quad \text{as } \tau \rightarrow \infty. \quad (11)$$

Proof. See Appendix B.

4. Distributed MPC as a function approximator

In Gros and Zanon (2020) it was shown how an MPC scheme can capture the RL policy, state-value, and action-value functions. This section introduces a distributed counterpart, parametrized in $\theta = (\theta_1^\top, \dots, \theta_M^\top)^\top$, approximating these quantities for a multi-agent system.

4.1. Parametrized distributed MPC scheme

Consider the distributed MPC-based action-value function approximation:

$$\begin{aligned} Q_\theta(s, a) &= \min_{(\{x_i\}, \{u_i\}, \{\sigma_i\})_{i \in \mathcal{M}}} \sum_{i \in \mathcal{M}} \left(\beta_{\theta_i}(x_i(0)) \right. \\ &+ \sum_{k \in \mathcal{K}} \left(\gamma^k (l_{\theta_i}(x_i(k), u_i(k), \{x_j(k)\}_{j \in \mathcal{N}_i}) \right. \\ &\left. \left. + w^\top \sigma_i(k)) \right) + \gamma^N V_{f, \theta_i}(x_i(N)) \right) \end{aligned} \quad (12a)$$

s.t. $\forall i \in \mathcal{M}, \forall k \in \mathcal{K}$:

$$u_i(0) = a_i, \quad x_i(0) = s_i \quad (12b)$$

$$x_i(k+1) = f_{\theta_i}(x_i(k), u_i(k), \{x_j(k)\}_{j \in \mathcal{N}_i}) \quad (12c)$$

$$h_{\theta_i}(x_i(k), u_i(k), \{x_j(k)\}_{j \in \mathcal{N}_i}) \leq \sigma_i(k). \quad (12d)$$

The functions β_{θ_i} , l_{θ_i} , and V_{f, θ_i} are the initial, stage, and terminal cost approximations respectively, f_{θ_i} is a model approximation, and h_{θ_i} is an inequality constraint function. Each is parametrized by a local parameter θ_i and known only to the corresponding agent, such that knowledge of the components of the MPC scheme is distributed across the agents. Parameterized terminal constraints could be included, but are omitted here for simplicity. For conciseness, we summarize (12) as

$$\begin{aligned} Q_\theta(s, a) &= \min \sum_{i \in \mathcal{M}} F_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \boldsymbol{\sigma}_i) \\ \text{s.t. } \forall i \in \mathcal{M}, \\ &G_{\theta_i}(s_i, a_i, \mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i) = 0 \\ &H_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \boldsymbol{\sigma}_i) \leq 0. \end{aligned} \quad (13)$$

The joint policy is then computed as

$$\begin{aligned} \pi_\theta(s) &= \underset{\{u_i(0)\}_{i \in \mathcal{M}}}{\operatorname{argmin}} \sum_{i \in \mathcal{M}} F_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \sigma_i) \\ \text{s.t. } \forall i \in \mathcal{M}, \forall k \in \mathcal{K} : \\ (12c)-(12d), \quad x_i(0) &= s_i \end{aligned} \quad (14)$$

and the global value function $V_\theta(s)$ can be obtained as the optimal value of (14), satisfying the fundamental Bellman equations (Gros & Zanon, 2020).

Intuitively, this structured MPC scheme approximates the local cost of each agent through the local cost functions. Interactions between coupled agents enter the scheme via the parametrized stage costs F_{θ_i} , the dynamics in G_{θ_i} , and the inequality constraints in H_{θ_i} . Inexact knowledge of agents' dynamics, constraints, inter-agent interactions, and local costs can be encoded in an initial guess of the local parameters θ_i . Through learning, the cost and model representations will be modified to improve the global performance. The constraints in (13) are chosen to be affine and the functions F_{θ_i} to be strictly convex in their arguments \mathbf{x}_i , $\{\mathbf{x}_j\}_{j \in \mathcal{N}_i}$, \mathbf{u}_i , and σ_i , such that (13) satisfies the assumptions for Propositions 1 and 2. This is not a strong assumption for linear systems, as the true optimal value function is convex when the stage costs are convex, as is common in the literature (Borrelli et al., 2016). Indeed, for quadratic stage costs, with a sufficiently long prediction horizon, a convex optimization problem can capture the true infinite horizon cost (Chmielewski & Manousiouthakis, 1996).

4.2. Distributed evaluation

The ADMM and GAC algorithms can be used to solve (12) distributively. Each agent stores a local copy of the predicted states of neighboring agents over the prediction horizon, and constructs the augmented state $\tilde{\mathbf{x}}_i = (\mathbf{x}_i^\top, \operatorname{col}_{j \in \mathcal{N}_i}^\top(\mathbf{x}_j^{(i)}))^\top$, where $\mathbf{x}_j^{(i)}$ is agent i 's local copy of agent j 's predicted state over the prediction horizon. Global copies are introduced of each state prediction $\mathbf{z} = (\mathbf{z}_1^\top, \mathbf{z}_2^\top, \dots, \mathbf{z}_M^\top)^\top$, with the relevant components of \mathbf{z} for agent i denoted $\tilde{\mathbf{z}}_i = (\mathbf{z}_i^\top, \operatorname{col}_{j \in \mathcal{N}_i}^\top(\mathbf{z}_j))^\top$.

As in Section 2.3, ADMM solves (13) by the following iterations:

$$\begin{aligned} \tilde{\mathbf{x}}_i^{\tau+1}, \mathbf{u}_i^{\tau+1}, \sigma_i^{\tau+1} &= \underset{\{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \sigma_i}{\operatorname{argmin}} F_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \sigma_i) \\ &\quad + \sum_{k \in \mathcal{K}} \gamma_i^{\tau+1}(k) (\tilde{x}_i(k) - \tilde{z}_i(k)^\tau) \|^2 \\ \text{s.t. } G_{\theta_i}(s_i, a_i, \mathbf{x}_i, \{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i) &= 0 \\ H_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \sigma_i) &\leq 0 \end{aligned} \quad (15a)$$

$$\mathbf{z}_i^{\tau+1} = \frac{1}{|\mathcal{N}_i| + 1} (\mathbf{x}_i^{\tau+1} + \sum_{j \in \mathcal{N}_i} \mathbf{x}_j^{(j), \tau+1}) \quad (15b)$$

$$\mathbf{y}_i^{\tau+1} = \mathbf{y}_i^\tau + \rho(\tilde{\mathbf{x}}_i^{\tau+1} - \tilde{\mathbf{z}}_i^{\tau+1}). \quad (15c)$$

By Proposition 1, as $\tau \rightarrow \infty$, the outputs of the local minimization (15a) converge to the minimizers of the original problem (13), i.e., $\tilde{\mathbf{x}}_i^{\tau+1}, \mathbf{u}_i^{\tau+1}, \sigma_i^{\tau+1} \rightarrow \tilde{\mathbf{x}}_i^*, \mathbf{u}_i^*, \sigma_i^*$. Agents then evaluate their local objective component $F_{\theta_i}^* = F_{\theta_i}(\mathbf{x}_i^*, \{\mathbf{x}_j^{(i),*}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i^*, \sigma_i^*)$. The global action-value $Q_\theta(s, a) = \sum_{i \in \mathcal{M}} F_{\theta_i}^*$ can then be agreed upon across the network via the GAC algorithm. Each agent makes the naive initial guess $Q_\theta(s, a) = MF_{\theta_i}^*$, i.e., each agent assumes that all other agents have the same local cost. The GAC algorithm gives convergence of these values, as per Section 2.3, to the average $\sum_{i \in \mathcal{M}} F_{\theta_i}^*$, which is the global action-value.

Evaluation of the global value $V_\theta(s)$ and the joint policy follows from applying the same steps to (14). From the optimal control

sequence \mathbf{u}_i^* agents evaluate their local component $a_i = \pi_{\theta_i}(s_i) = u_i^*(0)$ of the joint policy $\pi_\theta(s)$. We highlight that, while the joint policy is a function of the global state s , the local policies are functions only of the local state s_i , i.e., $\pi_{\theta_i}(s_i)$, as (15a) requires knowledge only of s_i .

4.3. Summary of distributed MPC as function approximator

We briefly summarize the key points of the distributed-MPC-based approximator. Each agent stores local knowledge of its learnable parameter θ_i , and its local parametrized functions β_{θ_i} , l_{θ_i} , V_{f, θ_i} , f_{θ_i} , and h_{θ_i} . In addition, it maintains its own state prediction \mathbf{x}_i , copies of predictions for its neighbors' states $\{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i}$, and an additional set of copies required for agreement. Finally, each agent stores the local Lagrange multipliers \mathbf{y}_i used in ADMM. Evaluating Q_θ (or V_θ), from the perspective of agent i , is summarized in Algorithm 1. For simplicity, both ADMM and GAC use a fixed number of iterations in the algorithm, T_A and T_C respectively, with the consequences discussed in Section 5.2.

Algorithm 1 Evaluation of $Q_\theta(s, a)$ for agent i .

- 1: **Inputs:** s_i and a_i (not a_i if evaluating (14)).
 - 2: **Initialize:** $\{\mathbf{x}_j^{(i)}\}_{j \in \mathcal{N}_i} \leftarrow 0, \mathbf{z}_i \leftarrow 0, \mathbf{y}_i \leftarrow 0$.
 - 3: **for** $\tau = 0, 1, \dots, T_A$ **do**
 - 4: Get $\tilde{\mathbf{x}}_i^{\tau+1}, \mathbf{u}_i^{\tau+1}, \sigma_i^{\tau+1}$ via (15a).
 - 5: For $j \in \mathcal{N}_i$, send $\mathbf{x}_j^{(i), \tau+1}$ and receive $\mathbf{x}_j^{(j), \tau+1}$.
 - 6: Perform averaging step (15b).
 - 7: For $j \in \mathcal{N}_i$, send $\mathbf{z}_i^{\tau+1}$ and receive $\mathbf{z}_j^{\tau+1}$.
 - 8: Perform local multipliers update (15c).
 - 9: **end for**
 - 10: $\bar{Q}_\theta(s, a) \leftarrow MF_{\theta_i}(\mathbf{x}_i^{T_A}, \{\mathbf{x}_j^{(i), T_A}\}_{j \in \mathcal{N}_i}, \mathbf{u}_i^{T_A}, \sigma_i^{T_A})$
 - 11: Perform T_C iterations GAC, with initial guess $\bar{Q}_\theta(s, a)$, to agree on $Q_\theta(s, a)$.
 - 12: **Outputs:** Global state-value $Q_\theta(s, a)$ (local action $a_i \leftarrow u_i^{T_A}(0)$ if evaluating (14)).
-

5. Distributed Q-learning

In this section we show that using Q-learning as the RL algorithm to learn the local parameters θ_i enables a distributed learning update that avoids nonstationarity. Q-learning (Sutton & Barto, 2018) adjusts, at each time step t , the global parameters $\theta = (\theta_1^\top, \dots, \theta_M^\top)^\top$ as

$$\begin{aligned} \delta_t &= L_t + \gamma V_\theta(s_{t+1}) - Q_\theta(s_t, a_t) \\ \theta &\leftarrow \theta + \alpha \delta_t \nabla_\theta Q_\theta(s_t, a_t), \end{aligned} \quad (16)$$

where $\alpha \in \mathbb{R}$ is a learning rate and $\delta_t \in \mathbb{R}$ is the temporal-difference (TD) error at time step t .

5.1. Separable Q-learning updates

The update (16) appears to be centralized due to the global components a_t, L_t, s_t, s_{t+1} , and θ . However, the structure of (12) allows decomposition into local updates of θ_i for each agent. Addressing first the TD error δ_t , we have shown in Section 4 that $V_\theta(s)$ and $Q_\theta(s, a)$ can be evaluated distributively. The globally averaged cost L_t is the average of all agents' locally incurred costs $L_{i,t}$, and can be shared across the network using the GAC algorithm. This can be incorporated into line 11 of Algorithm 1, with no additional communication overhead.

We now address the gradient term $\nabla_{\theta} Q_{\theta}(s_t, a_t)$. The Lagrangian of (12) is

$$\begin{aligned} \mathcal{L}_{\theta}(s, a, p) = \sum_{i \in \mathcal{M}} & \left(F_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \boldsymbol{\sigma}_i) \right. \\ & + \boldsymbol{\lambda}_i^{\top} G_{\theta_i}(s_i, a_i, \mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i) \\ & \left. + \boldsymbol{\mu}_i^{\top} H_{\theta_i}(\mathbf{x}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \mathbf{u}_i, \boldsymbol{\sigma}_i) \right), \end{aligned} \quad (17)$$

where $\boldsymbol{\lambda}_i$ and $\boldsymbol{\mu}_i$ are the multiplier vectors associated with the equality and inequality constraints, respectively, and the primal and dual variables are grouped as

$$p = (\{\mathbf{x}_i\}, \{\mathbf{u}_i\}, \{\boldsymbol{\sigma}_i\}, \{\boldsymbol{\lambda}_i\}, \{\boldsymbol{\mu}_i\})_{i \in \mathcal{M}}. \quad (18)$$

We stress again that the multipliers in p are associated with the constraints of (12), and are unrelated to the multipliers \mathbf{y}_i used in the ADMM procedure. Via sensitivity analysis (Büskens & Maurer, 2001), the gradient $\nabla_{\theta} Q_{\theta}(s, a)$ coincides with the gradient of the Lagrangian at the optimal primal and dual variables p^* :

$$\nabla_{\theta} Q_{\theta}(s, a) = \frac{\partial \mathcal{L}_{\theta}(s, a, p^*)}{\partial \theta}. \quad (19)$$

Observe that (17) is separable over parameters θ_i , states s_i , actions a_i , and subsets of the primal and dual variables p_i , i.e., $\mathcal{L}_{\theta}(s, a, p) = \sum_{i \in \mathcal{M}} \mathcal{L}_{\theta_i}(s_i, a_i, p_i)$, where \mathcal{L}_{θ_i} is the i th term within the sum (17), and $p_i = (\mathbf{x}_i, \mathbf{u}_i, \boldsymbol{\sigma}_i, \{\mathbf{x}_j\}_{j \in \mathcal{N}_i}, \boldsymbol{\lambda}_i, \boldsymbol{\mu}_i)$. We then express (19) as

$$\frac{\partial \sum_{i \in \mathcal{M}} \mathcal{L}_{\theta_i}(s_i, a_i, p_i^*)}{\partial \theta} = \begin{bmatrix} \frac{\partial \mathcal{L}_{\theta_1}(s_1, a_1, p_1^*)}{\partial \theta_1} \\ \vdots \\ \frac{\partial \mathcal{L}_{\theta_M}(s_M, a_M, p_M^*)}{\partial \theta_M} \end{bmatrix}, \quad (20)$$

and the centralized parameter update (16) can be hence unpacked into M local updates:

$$\theta_i \leftarrow \theta_i + \alpha \delta_t \frac{\partial \mathcal{L}_{\theta_i}(s_{i,t}, a_{i,t}, p_i^*)}{\partial \theta_i}, \quad i \in \mathcal{M}. \quad (21)$$

What then remains to be shown is that the subset of optimal primal and dual variables p_i^* for (12) is available locally to agent i . By Proposition 1, at convergence of ADMM, the local minimization (15a) returns the minimizers $(\mathbf{x}_i^*, \mathbf{u}_i^*, \boldsymbol{\sigma}_i^*)$ and local copies of the minimizers of neighboring agents $\{\mathbf{x}_j^{(i),*}\}_{j \in \mathcal{N}_i} = \{\mathbf{x}_j^*\}_{j \in \mathcal{N}_i}$. Each agent hence has local knowledge of the optimal *primal* values in p_i^* . For the optimal *dual* values, by Proposition 2, at convergence of ADMM, the optimal dual variables of the local minimization (15a) for agent i are equal to the relevant subset of optimal dual variables for the original problem, i.e., $(\boldsymbol{\lambda}_i^*, \boldsymbol{\mu}_i^*)$. Each agent hence has local knowledge of the optimal *dual* values in p_i^* as well. Agents can therefore perform the local update (21). Nonstationarity is avoided as the local updates reconstruct the centralized update of the whole network (16).

5.2. Implementation details

In this section we discuss some auxiliary details in the implementation of our proposed approach.

- Exploration can be added to the approach in, e.g., an epsilon-greedy fashion, in which case the agent adds a random perturbation to the objective (12a) with some probability ϵ_i , where ϵ_i decreases as training progresses (Airoldi et al., 2023; Zanon & Gros, 2021). This causes a perturbation in the joint policy. It is assumed that the system, with exploration injected, is sufficiently persistently exciting.

- The finite termination of ADMM and the GAC algorithm introduces errors in the evaluation of V_{θ} and Q_{θ} , π_{θ} , and p^* , and could lead to instability in the learning. To counter this, large numbers of ADMM and GAC iterations may be used in a simulated learning phase, when there are no constraints on computation time between actions. This is desirable as primal and dual variables with high precision are needed to reliably compute the sensitivity (19). On the contrary, at deployment, sensitivities are not required and only the policy must be evaluated. Thus, the iteration numbers can be reduced, as ADMM often converges to modest accuracy within few iterations (Boyd, 2010). Additionally, experience replay (ER) (Airoldi et al., 2023; Lin, 1992) can improve learning stability by using an average of past observations when calculating the gradient and the TD error. In our method ER requires no extra mechanism as agents can maintain a local history of values for δ and $\nabla_{\theta_i} \mathcal{L}_{\theta_i}(s, p_i^*)$. In our numerical experiments, we found the learning to succeed with ER and modest numbers of ADMM and GAC iterations.
- The variables $\alpha > 0$ and $\gamma \in (0, 1]$ are hyperparameters; as the distributed update fully reconstructs the centralized update, existing methods for selecting these parameters in centralized learning apply directly, e.g., see Sutton and Barto (2018, Chapter 9.6).

6. Example

This section presents a numerical example. Source code and simulation results can be found at <https://github.com/SamuelMallick/dmpcrl-concept>. Additionally, the longer online version of this manuscript includes a power systems case study. We modify the system from Gros and Zanon (2020), forming a three-agent system with state coupling in a chain, i.e., $1 \leftrightarrow 2 \leftrightarrow 3$, with real (unknown) dynamics $s_i(t+1) = A_i s_i(t) + B_i a_i(t) + \sum_{j \in \mathcal{N}_i} A_{ij} s_j(t) + [e_i(t), 0]^{\top}$ where

$$A_i = \begin{bmatrix} 0.9 & 0.35 \\ 0 & 1.1 \end{bmatrix}, \quad \{A_{ij}\}_{j \in \mathcal{N}_i} = \begin{bmatrix} 0 & 0 \\ 0 & -0.1 \end{bmatrix}, \quad (22)$$

with $B = [0.0813, 0.2]^{\top}$ and $e_i(t)$ uniformly distributed over the interval $[-0.1, 0]$. The RL task is to drive the states and control towards the origin while avoiding violations of the state constraints $\underline{s} \leq s_i \leq \bar{s}$, with local costs as:

$$\begin{aligned} L_{i,t}(s_i, a_i) = & \|s_i\|_2^2 + \frac{1}{2} \|u_i\|_2^2 \\ & + \max(0, \omega^{\top}(\underline{s} - s_i)) + \max(0, \omega^{\top}(s_i - \bar{s})), \end{aligned} \quad (23)$$

with $\omega = [10^2, 10^2]^{\top}$, $\underline{s} = [0, -1]^{\top}$, and $\bar{s} = [1, 1]^{\top}$. Non-positive noise on the first state biases that term towards violating the lower bound of zero. This renders the task challenging, as the agents must regulate the state to zero, and yet driving the first dimension to zero will result in constraint violations due to the noise. The true model is unknown, with agents instead knowing a uniform distribution of models, containing the true model:

$$\hat{A}_i = \begin{bmatrix} 1 + a_{1,i} & 0.25 + a_{2,i} \\ 0 & 1 + a_{3,i} \end{bmatrix}, \quad \{\hat{A}_{ij}\}_{j \in \mathcal{N}_i} = \begin{bmatrix} 0 & 0 \\ 0 & c_i \end{bmatrix}, \quad (24)$$

with $\hat{B}_i = [0.0312 + b_{1,i}, 0.25 + b_{2,i}]^{\top}$ and, for all i , the random variables distributed uniformly as $a_{1,i}, a_{2,i}, a_{3,i} \in [-0.1, 0.1]$, $b_{1,i} \in [0, 0.075]$, $b_{2,i} \in [-0.075, 0]$ and $c_i \in [-0.1, 0]$. We




Fig. 1. Accuracy of the dual variables recovered by Proposition 2 as a function of the ADMM iteration index τ .

implement the following distributed MPC scheme:

$$\begin{aligned} \min_{\{(x_i, u_i, \sigma_i)\}_{i \in \mathcal{M}}} & \sum_{i \in \mathcal{M}} \left(V_{i,0} + \sum_{k \in \mathcal{K}} f_i^\top \begin{bmatrix} x_i(k) \\ u_i(k) \end{bmatrix} \right) \\ & + \frac{1}{2} \gamma^k \left(\|x_i(k)\|^2 + \frac{1}{2} \|u_i(k)\|^2 + \omega^\top \sigma_i(k) \right) \\ \text{s.t. } & \forall i \in \mathcal{M}, \forall k \in \mathcal{K}: \\ & x_i(k+1) = A_i x_i(k) + B_i u_i(k) \\ & \quad + \sum_{j \in \mathcal{N}_i} A_{ij} x_j(k) + b_i \\ & \underline{s} + x_i - \sigma_i(k) \leq x_i(k) \leq \bar{s} + \bar{x}_i + \sigma_i(k) \\ & -1 \leq u_i(k) \leq 1, x_i(0) = s_i \end{aligned}$$

where $M = 3$, $\mathcal{K} = \{0, \dots, 9\}$, and $\gamma = 0.9$. The learnable parameters for each agents are then $\theta_i = (V_{i,0}, \underline{x}_i, \bar{x}_i, b_i, f_i, A_i, B_i, \{A_{ij}\}_{j \in \mathcal{N}_i})$. The initial values for $A_i, B_i, \{A_{ij}\}_{j \in \mathcal{N}_i}$ are the inaccurate model (24) with the random variables set to zero. All other learnable parameters are initialized to zero.

To illustrate Proposition 2, for a given global state s , Fig. 1 shows the error between the true optimal dual variables and those recovered from evaluating the MPC scheme with ADMM, as a function of the ADMM iteration index τ . The locally recovered dual variables are close to the true values, with the error initially decreasing with the iteration index.

We now compare the learning performance of our distributed approach with the centralized approach inspired from Gros and Zanon (2020), using the same MPC scheme, and the same learning hyperparameters for both. We use an exploration probability of $\epsilon_i = 0.7$, exponentially decaying with rate 0.99, with local costs perturbed uniformly over the interval $[-1, 1]$. We use a learning rate of $\alpha = 6e^{-5}$, exponentially decaying with rate 0.9996, and ER with an average over 15 past samples at an update rate of every 2 time steps. For the distributed approach we use 50 iterations in ADMM and 100 iterations in GAC. These values were tuned to keep the iterations low without introducing significant approximation errors. Fig. 2 shows the state and input trajectories for the three agents during training. Fig. 3 shows the evolution of the global TD error and the collective cost. Fig. 4 shows the learnable parameters of the second agent during training (similar convergence profiles are observed for the other agents). It is seen that the behavior of the distributed approach is similar to that of the centralized approach, reducing the TD error and costs incurred, with the centralized approach converging slightly faster. The costs are reduced by maintaining the first state of each agent above zero to prevent expensive violations of the state constraint due to coupling and noise.

We compare the performance of the distributed policy against both a distributed nominal MPC controller (NMPC) and a distributed stochastic MPC controller (SMPC) based on the scenario approach (Schildbach, Fagiano, Frei, & Morari, 2014). The nominal MPC controller uses the inexact model (24), with all random

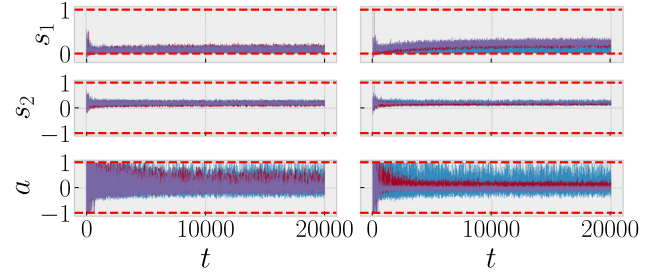


Fig. 2. Centralized (left) and distributed (right). Evolution of the states and inputs during training. Agent 1 (blue), agent 2 (red), agent 3 (purple), and bounds (dashed). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 3. Evolution of TD errors (top) and stage costs (bottom) during training.




Fig. 4. Evolution of learnable parameters for agent 2 during training. Distributed (blue) and centralized (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

variables set to 0. For the SMPC controller we consider the case where the true model is unknown and $N = 25$ samples of the model distribution (24) and the noise distribution are used to account for the uncertainty. The number of samples was manually tuned to balance conservativeness and robustness (Schildbach et al., 2014). We also consider the case where the exact model (22) is known, and only the noise distribution is sampled. The controllers are compared in Fig. 5, showing the closed-loop cost accumulated over 100 time steps. The learned policy can be seen to improve from a performance initially comparable to NMPC, learning to outperform SMPC, and approaching the performance of SMPC with a perfect model. We highlight that our approach retains the computational complexity of NMPC, while SMPC is significantly more complex, optimizing over N copies of the state trajectories.




Fig. 5. NMPC and SMPC compared against policy at training time steps t (log scale).

7. Conclusions

This paper has extended the idea of using MPC-based RL to the multi-agent setting. We have proposed a novel approach to MARL via the use of distributed MPC as a distributed function approximator within Q-learning. A result on the optimal dual variables in ADMM is first presented. The structure of the distributed MPC function approximator is then detailed, and is shown to enable a distributed evaluation of the global value functions and the local policies. Finally, by using Q-learning to update the parametrization of the MPC scheme, the parameter updates can also be performed fully distributively. The effectiveness of the newly proposed method is demonstrated on a numerical example.

A limitation of the proposed approach is the convexity requirement on the MPC scheme, restricting the class of systems for whom the optimal policy can be captured to linear systems with convex cost functions. Additionally, as both the ADMM and GAC algorithms use iterative communication between neighbors, the approach is less suited for applications with high communication costs. Future work will look at extending the idea to other RL methods, such as policy-based approaches, and a formal analysis on the propagation of errors from the finite termination of the ADMM and GAC algorithms.

Appendix A. Equivalence of (3) and (7)

Define the aggregate of the local augmented states as $\bar{x} = \text{col}_{i \in \mathcal{M}}(\tilde{x}_i)$ and the feasible set as $\bar{\mathcal{X}} = \tilde{\mathcal{X}}_1 \times \dots \times \tilde{\mathcal{X}}_M$. Problem (7) can then be written in the form of (3) as

$$\min_{\bar{x} \in \bar{\mathcal{X}}, z} \{f_{\text{ADMM}}(\bar{x}) + g_{\text{ADMM}}(z) : A\bar{x} + Bz = c\}, \quad (\text{A.1})$$

where $f_{\text{ADMM}}(\bar{x}) = \sum_{i \in \mathcal{M}} F_i(\tilde{x}_i)$, $g_{\text{ADMM}}(\cdot)$ is the zero map, and $c = \mathbf{0}$. Furthermore, $A = \mathbf{I}_{q \times q}$, $q = \sum_{i \in \mathcal{M}} n(|\mathcal{N}_i| + 1)$, and $B \in \mathbb{R}^{q \times nM}$ is a block matrix that contains negative identity matrix entries picking out the corresponding global states, with all other entries being zero matrices. Clearly, matrix A is full column rank. For B , we note that every column has at least one identity entry, and that by rearranging its rows we can express it as $[-\mathbf{I}_{nM \times nM}, B_+^T]^T$, where B_+ contains the remaining rows. Therefore, B is full column rank. The assumptions in Proposition 1 are hence satisfied.

Appendix B. Proof of Proposition 2

The proof involves showing that the Karush–Kuhn–Tucker (KKT) conditions for the local minimization problems (8a), when aggregated for all agents, are equivalent to the KKT conditions of the original problem (5). Then, due to strict convexity of both (5) and (8a), the primal and dual solutions are unique, and the equivalence between the optimal dual variables follows. We write the optimal primal and dual variables of (5) as $(\{x_i^*\}_{i \in \mathcal{M}}, \{\lambda_i^*\}_{i \in \mathcal{M}}, \{\mu_i^*\}_{i \in \mathcal{M}})$, and the optimal primal and dual variables of (8a) for agent i , at convergence of ADMM, as $(x_i^*, \{\tilde{x}_j^{(i),*}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i^*, \hat{\mu}_i^*)$. Primal variable equivalence is given by

Proposition 1, i.e., the x_i^* 's are equivalent for both problems. Also, optimal local copies are equal to their true values

$$\{x_j^{(i),*}\}_{j \in \mathcal{N}_i} = \{x_j^*\}_{j \in \mathcal{N}_i} \quad i \in \mathcal{M}. \quad (\text{B.1})$$

First, we state the KKT conditions of (5). These are

$$h_i(x_i^*, \{x_j^*\}_{j \in \mathcal{N}_i}) \leq 0 \quad i \in \mathcal{M} \quad (\text{B.2a})$$

$$g_i(x_i^*, \{x_j^*\}_{j \in \mathcal{N}_i}) = 0 \quad i \in \mathcal{M} \quad (\text{B.2b})$$

$$\lambda_i^* \geq 0 \quad i \in \mathcal{M} \quad (\text{B.2c})$$

$$\lambda_i^{*\top} h_i(x_i^*, \{x_j^*\}_{j \in \mathcal{N}_i}) = 0 \quad i \in \mathcal{M} \quad (\text{B.2c})$$

$$\nabla_x \left(\sum_{i \in \mathcal{M}} F_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \lambda_i^\top h_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \mu_i^\top g_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) \right) \Big|_{\substack{(x_i, \lambda_i, \mu_i)_{i \in \mathcal{M}} \\ = (x_i^*, \lambda_i^*, \mu_i^*)_{i \in \mathcal{M}}}} = 0, \quad (\text{B.2d})$$

where $x = (x_1^\top, \dots, x_M^\top)^\top$. Now consider the composition of the KKT conditions of each agent's local minimization (8a). First, primal feasibility, dual feasibility, and complementary slackness for agent i :

$$h_i(x_i^*, \{\tilde{x}_j^{(i),*}\}_{j \in \mathcal{N}_i}) \leq 0, \quad g_i(x_i^*, \{\tilde{x}_j^{(i),*}\}_{j \in \mathcal{N}_i}) = 0, \quad (\text{B.3})$$

$$\hat{\lambda}_i^* \geq 0, \quad \hat{\lambda}_i^{*\top} h_i(x_i^*, \{\tilde{x}_j^{(i),*}\}_{j \in \mathcal{N}_i}) = 0.$$

Inserting the equivalence (B.1) into (B.3), we find that the composition of these conditions for all agents are the original problem's conditions (B.2a)–(B.2c).

For ease of exposition in showing the equivalence of the stationarity conditions, let us define

$$\mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \lambda_i, \mu_i) = F_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \lambda_i^\top h_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}) + \mu_i^\top g_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}). \quad (\text{B.4})$$

Moreover, we group the primal and dual variables as $p = (\{x_i\}_{i \in \mathcal{M}}, \{\lambda_i\}_{i \in \mathcal{M}}, \{\mu_i\}_{i \in \mathcal{M}})$, such that the stationarity condition (B.2d) reads

$$\nabla_x \left(\sum_{i \in \mathcal{M}} \mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \lambda_i, \mu_i) \right) \Big|_{p=p^*} = 0. \quad (\text{B.5})$$

By decomposing the derivative operator and considering each row l of (B.5) we have

$$\left(\frac{\partial}{\partial x_l} \sum_{i \in \mathcal{M}} \mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \lambda_i, \mu_i) \right) \Big|_{p=p^*} = 0. \quad (\text{B.6})$$

Observing that $\frac{\partial}{\partial x_l} \mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \lambda_i, \mu_i) = 0$ if $i \notin \mathcal{N}_l \cup l$, from (B.6) we have

$$\left(\frac{\partial}{\partial x_l} \mathcal{F}_l(x_l, \{x_j\}_{j \in \mathcal{N}_l}, \lambda_l, \mu_l) \right) \Big|_{p=p^*} + \sum_{i \in \mathcal{N}_l} \left(\frac{\partial}{\partial x_l} \mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \lambda_i, \mu_i) \right) \Big|_{p=p^*} = 0. \quad (\text{B.7})$$

We highlight that the partial derivative terms within the sum are not zero, as x_l is one of the elements of $\{x_j\}_{j \in \mathcal{N}_i}$. We will now show that the condition (B.7) for each l can be reconstructed by a unique subset of the KKT conditions arising in the agents' local minimizations. Consider the stationarity conditions of the agents' local minimizations. The primal and dual variables for agent i are grouped as $p_i = (\tilde{x}_i, \hat{\lambda}_i, \hat{\mu}_i)$. The stationarity condition for agent i 's local minimization can be written as

$$\nabla_{\tilde{x}_i} \left(\mathcal{F}_i(x_i, \{\tilde{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) + y_i^{\top} \tilde{x}_i + \frac{\rho}{2} \|\tilde{x}_i - \tilde{z}_i^{\top}\|_2^2 \right) \Big|_{p_i=p_i^*} = 0. \quad (\text{B.8})$$

Evaluating the derivative of the second term, we have

$$\nabla_{\tilde{x}_i} \left(\mathcal{F}_i(x_i, \{\tilde{x}_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \Big|_{p_i=p_i^*} + y_i^{\top} + \rho(\tilde{x}_i^* - \tilde{z}_i^{\top}) = 0. \quad (\text{B.9})$$

From [Proposition 1](#), we have that at convergence of ADMM, $\tilde{x}_i^* - \tilde{z}_i^* = 0$. Decomposing the derivative operator we have

$$\left(\begin{array}{c} \left[\text{col}_{j \in \mathcal{N}_i} \left(\frac{\partial}{\partial x_i} \mathcal{F}_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \right]_{p_i=p_i^*} \\ + \left[\text{col}_{j \in \mathcal{N}_i} (y_j^{(i), \tau}) \right] \end{array} \right) = 0 \quad (\text{B.10})$$

We denote this as $\Psi_i = \left[\Psi_i^{(i), \top} \quad \text{col}_{j \in \mathcal{N}_i}^{\top} (\Psi_j^{(i)}) \right]^{\top}$, where $\Psi_j^{(i)}$ is the row of [\(B.10\)](#) corresponding to the partial derivative with respect to $x_j^{(i)}$. Let us sum all rows of the agents' local stationarity conditions which correspond to the partial derivative with respect to x_i , or a copy of x_i , to get

$$\Psi_i^{(l)} + \sum_{i \in \mathcal{N}_l} \Psi_i^{(l)} = 0. \quad (\text{B.11})$$

The result is equal to zero as each term in the sum is zero. The sum expression reads

$$\begin{aligned} & \left(\frac{\partial}{\partial x_l} \mathcal{F}_l(x_l, \{x_j^{(l)}\}_{j \in \mathcal{N}_l}, \hat{\lambda}_l, \hat{\mu}_l) \right) \Big|_{p_l=p_l^*} \\ & + \sum_{i \in \mathcal{N}_l} \left(\frac{\partial}{\partial x_l^{(i)}} \mathcal{F}_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \Big|_{p_i=p_i^*} \\ & + y_l^{\tau} + \sum_{i \in \mathcal{N}_l} y_i^{(i), \tau} = 0, \end{aligned} \quad (\text{B.12})$$

where again we highlight that the partial derivative terms in the sum are not zero, as $l \in \mathcal{N}_l$. It is shown in [Boyd \(2010\)](#) that the sum of the dual ADMM variables corresponding to the same global variables is zero after the first iteration of the procedure, i.e., $y_l^{\tau} + \sum_{i \in \mathcal{N}_l} y_i^{(i), \tau} = 0$, $\tau > 1$. Hence, at convergence, we have

$$\begin{aligned} & \left(\frac{\partial}{\partial x_l} \mathcal{F}_l(x_l, \{x_j^{(l)}\}_{j \in \mathcal{N}_l}, \hat{\lambda}_l, \hat{\mu}_l) \right) \Big|_{p_l=p_l^*} \\ & + \sum_{i \in \mathcal{N}_l} \left(\frac{\partial}{\partial x_l^{(i)}} \mathcal{F}_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \Big|_{p_i=p_i^*} = 0. \end{aligned} \quad (\text{B.13})$$

Substituting the equivalence of the local copies [\(B.1\)](#) and observing that

$$\begin{aligned} & \left(\frac{\partial}{\partial x_l^{(i)}} \mathcal{F}_i(x_i, \{x_j^{(i)}\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \Big|_{p_i=p_i^*} = \\ & \left(\frac{\partial}{\partial x_l} \mathcal{F}_i(x_i, \{x_j\}_{j \in \mathcal{N}_i}, \hat{\lambda}_i, \hat{\mu}_i) \right) \Big|_{p_l=p_l^*}, \end{aligned} \quad (\text{B.14})$$

we obtain the equivalent stationarity condition from the l th row of [\(B.5\)](#). Repeating the sum in [\(B.11\)](#) for each agent, we reconstruct the equivalent rows of [\(B.5\)](#). We then have equivalence between the stationarity conditions on $(\{\hat{\lambda}_i^*\}_{i \in \mathcal{M}}, \{\hat{\mu}_i^*\}_{i \in \mathcal{M}})$ and on $(\{\lambda_i^*\}_{i \in \mathcal{M}}, \{\mu_i^*\}_{i \in \mathcal{M}})$. This concludes the proof.

References

- Airoldi, F., De Schutter, B., & A., D. (2023). Learning safety in model-based reinforcement learning using MPC and Gaussian processes. *IFAC-PapersOnLine*, 56(2), 5759–5764, 22nd IFAC World Congress.
- Alqahtani, M., Scott, M. J., & Hu, M. (2022). Dynamic energy scheduling and routing of a large fleet of electric vehicles using multi-agent reinforcement learning. *Computers & Industrial Engineering*, 169, Article 108180.
- Arulkumar, K., Deisenroth, M. P., Brundage, M., & Bharath, A. N. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6), 26–38.
- Borrelli, F., Bemporad, A., & Morari, M. (2016). *Predictive control for linear and hybrid systems*. Cambridge University Press.

- Boyd, S. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Büsken, C., & Maurer, H. (2001). Sensitivity analysis and real-time optimization of parametric nonlinear programming problems. In *Online optimization of large scale systems* (pp. 3–16). Berlin, Heidelberg: Springer.
- Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2), 156–172.
- Chmielewski, D., & Manousiouthakis, V. (1996). On constrained infinite-time linear quadratic optimal control. *Systems & Control Letters*, 29(3), 121–129.
- Foerster, J., Nardelli, N., Farquhar, G., Afouras, T., Torr, P. H. S., Kohli, P., et al. (2017). Stabilising experience replay for deep multi-agent reinforcement learning. In *International conference on machine learning*. PMLR (pp. 1146–1155).
- Gros, S., & Zanon, M. (2020). Data-driven economic NMPC using reinforcement learning. *IEEE Transactions on Automatic Control*, 65(2), 636–648.
- Gros, S., & Zanon, M. (2022). Learning for MPC with stability & safety guarantees. *Automatica*, 146, Article 110598.
- Gupta, J. K., Egorov, M., & Kochenderfer, M. (2017). Cooperative multi-agent control using deep reinforcement learning. *Autonomous Agents and Multiagent Systems*, 66–83.
- Hewing, L., Wabersich, K. P., Menner, M., & Zeilinger, M. N. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1), 269–296.
- Lin, L.-J. (1992). Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning*, 8, 293–321.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems*, 30.
- Maestre, J. M., & Negenborn, R. R. (Eds.). (2014). *Distributed Model Predictive Control Made Easy*. Dordrecht, The Netherlands: Springer.
- Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Sckaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.
- Mota, J. F. C., Xavier, J. M. F., Aguiar, P. M. Q., & Püschel, M. (2018). A proof of convergence for the alternating direction method of multipliers applied to polyhedral-constrained functions. *arXiv:1112.2295*.
- Olfati-Saber, R., Fax, J. A., & Murray, R. M. (2007). Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1), 215–233.
- Rosolia, U., & Borrelli, F. (2018). Learning model predictive control for iterative tasks. A data-driven control framework. *IEEE Transactions on Automatic Control*, 63(7), 1883–1896.
- Schildbach, G., Fagiano, L., Frei, C., & Morari, M. (2014). The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12), 3009–3018.
- Suttle, W., Yang, Z., Zhang, K., Wang, Z., Basar, T., & Liu, J. (2019). A multi-agent off-policy actor-critic algorithm for distributed reinforcement learning. *arXiv:1903.06372*.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. MIT Press.
- Wai, H., Yang, Z., Wang, Z., & Hong, M. (2018). Multi-agent reinforcement learning via double averaging primal-dual optimization. *Advances in Neural Information Processing Systems*, 31.
- Zanon, M., & Gros, S. (2021). Safe reinforcement learning using robust MPC. *IEEE Transactions on Automatic Control*, 66(8), 3638–3652.
- Zhang, K., Yang, Z., Liu, H., Zhang, T., & Basar, T. (2018). Fully decentralized multi-agent reinforcement learning with networked agents. In *International conference on machine learning*. PMLR (pp. 5872–5881).



Samuel Mallick received the B.Sc. and M.Sc. degrees from The University of Melbourne in 2020 and 2022, respectively. He is currently a Ph.D. candidate at the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. His research interests include model predictive control, reinforcement learning, and distributed control of large-scale and hybrid systems.



Filippo Airoldi received the B.Sc. and M.Sc. degrees from the Polytechnic University of Turin, Italy, in 2017 and 2019, respectively. He is currently a Ph.D. candidate at the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. His research interests include model predictive control, reinforcement learning, and other machine learning techniques, and in particular in their combination in learning-based control.



Azita Dabiri received the Ph.D. degree from the Automatic Control Group, Chalmers University of Technology, in 2016. She was a Post-Doctoral Researcher with the Department of Transport and Planning, Delft University of Technology, from 2017 to 2019. In 2019, she received an ERCIM Fellowship and also a Marie Curie Individual Fellowship, which allowed her to perform research at the Norwegian University of Technology (NTNU), as a Post-Doctoral Researcher, from 2019 to 2020, before joining the Delft Center for Systems and Control, Delft University of Technology, as

an Assistant Professor. Her research interests are in the areas of integration of model-based and learning-based control and its applications in transportation networks.



Bart De Schutter received the Ph.D. degree (summa cum laude) in applied sciences from KU Leuven, Belgium, in 1996. He is currently a Full Professor and Head of Department at the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. His research interests include multi-level and multi-agent control, model predictive control, learning-based control, and control of hybrid systems, with applications in intelligent transportation systems and smart energy systems. Prof. De Schutter is a Senior Editor of the IEEE Transactions on Intelligent Transportation

Systems.