

Adaptive Haptic Shared Control

MSc Thesis Report

Max McKenzie



Adaptive Haptic Shared Control

MSc Thesis Report

by

Max McKenzie

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on Monday February 26, 2024 at 13:00

Thesis committee:

Chair: Prof. dr. ir. M. Mulder

Supervisors: Dr. ir. M. M. van Paassen
Prof. dr. ir. M. Mulder

Additional member: Dr. ir. C. Borst

External examiner: Dr. ir. M. J. Ribeiro

Place: Faculty of Aerospace Engineering, Delft

Project Duration: March, 2023 - February, 2024

Student number: 4845730

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.



Copyright © Max McKenzie, 2024
All rights reserved.

Preface

This report is the culmination of my master thesis project, which started in May 2023 and ended in February 2024. It contains three parts: **Part I**) a scientific paper summarising the most important results, **Part II**) a preliminary report containing results of the literature study, and **Part III**) appendices to the scientific paper and preliminary report.

The reader is assumed to have a basic understanding of linear control theory and is expected to be familiar with McRuer's crossover model. Readers who are particularly interested in the results of the human-in-the-loop experiment can refer to the scientific paper in Part I of the report. For those who are interested in the methodology leading up to the final design of the adaptive haptic shared controller, this can be found in Part II. Part III contains appendices to the preceding parts.

This project has been my largest undertaking by far (at the time of writing), and contains work that I am truly proud to share. The only way I could have achieved it is by standing on the shoulders of giants, both inside academia and in my personal life. My two daily supervisors, Professor Max Mulder and Associate Professor René van Paassen, have both been immense inspirations to me throughout the duration of the project.

Max's deep knowledge, unwavering curiosity and meticulous attention to detail continuously motivated me to do my best. Although, beyond his extraordinary dedication and intuition, the thing I am most grateful for is his caring nature. As Carl Jung very wisely said: *"An understanding heart is everything in a teacher, and cannot be esteemed highly enough."* Without Max's genuine encouragement and compassion, I truly would not have been able to develop myself in the way that I have. I sincerely hope that one day I can pay his support forward to others.

René's contribution to the project, as well as towards my personal development, has also been invaluable. His technical expertise is unparalleled, repeatedly saving me from getting lost in dead-ends. The focus and single-mindedness with which René conducts his research is ever-present, and frequently illuminates what would otherwise be dark paths. Even though I have less experience, I felt that René took me and my ideas seriously, always taking care to steer me in the right direction. I have much to thank him for, and am very lucky to have had him as a supervisor.

Over in my personal life, Neve, Sam, Mum and Dad (and our dog, Walter) have always offered me unconditional love and support, for which I will forever be grateful. It is difficult to express the deep affection I have for each member of the McKenzie clan, and how much my family has helped me endure hardships throughout my life. These people are undoubtedly the most diverse set of characters I will ever encounter, and every moment spent with them is a gift.

I also owe a debt of gratitude to my close friends, Sietse and Ari, for always standing by my side and helping me to stay grounded. Sietse's honesty, humour and work ethic are an unstoppable combination, and always make him an absolute delight to be around. Whether it's discussing personal challenges, or knuckling down to do some hard work, Sietse has always been up for the challenge, day or night. And although Ari may never fully realize it, she's the reason that I managed to get through the dark winter months leading up to graduation. Through our chats over Yorkshire tea, she managed to help me stay grounded, and continuously reminded me to never take myself too seriously.

Finally, thank you, dear reader, for your interest in this project. I hope this document serves you well, and I wish you the best of luck with your endeavours!

*M. McKenzie
Antwerp, February 2024*

Contents

List of Figures	ix
List of Tables	xi
I Scientific Article	1
1 Implementing an Adaptive Haptic Shared Controller	3
1.1 Introduction	3
1.2 Background	4
1.3 Haptic Torque Rejection	6
1.4 Controller Design	9
1.5 Controller Simulations	10
1.6 Human-in-the-loop Experiment	11
1.7 Experiment Results	12
1.8 Discussion	15
1.9 Conclusion	16
II Preliminary Thesis Report	19
Summary	21
1 Introduction	23
1.1 Motivation For Adaptive Haptic Shared Control	24
1.2 Research Objective	24
1.3 Project Outline	26
2 Literature Review	27
2.1 Preview Cybernetics	27
2.1.1 Motivation to Study Preview Control	27
2.1.2 Van der El's Cybernetic Model for Preview Control	27
2.1.3 Adaptation to Preview Time	29
2.2 Four Design Choices Architecture	31
2.3 Four Design Choices Haptic Shared Controller	32
2.4 Effects of Controller Look-Ahead Time	35
2.4.1 Simulation and Experiment Results	35
2.4.2 Comparing Experiment Data with Simulation Results	36
2.4.3 HSC Performance for Varying Look-Ahead Times	36
3 Controller Design and Simulations	39
3.1 Motivation and Methodology	39
3.2 Validating FDC-HSC Model	40
3.3 Evaluation of System Identification Results	42
3.4 Conflict and HCR Analysis	45
3.4.1 Root Causes of Conflicts	45
3.4.2 LoHS and SoHF Conflicts	46
3.5 Modelling HO Rejection of HSC Torques	48
3.6 Modelling Time-varying HO Behaviour	52
3.7 Controller Design: Inertial A-HSC	52
3.7.1 Inertial A-HSC Structure	52
3.7.2 Inertial A-HSC Results	55

3.8	Model-Fitting Adaptive HSC	56
3.8.1	Model-Fitting A-HSC Structure.	56
3.8.2	Model-Fitting Results.	58
3.9	Cross-Correlation Adaptive HSC	59
3.9.1	Cross-Correlation A-HSC Structure	59
3.9.2	Cross-Correlation Results	60
3.10	Comparing A-HSC Designs	61
3.11	Cross-Correlation A-HSC: Additional Findings	62
3.11.1	Effects of Transient Behaviour on XC-AHSC Adaptation.	62
3.11.2	XC-AHSC Adaptation for Varying HO Levels of HO Rejection	64
3.11.3	Effects of Fog Direction on XC-AHSC Adaptation	65
3.11.4	Effects of Buffer Size on XC-AHSC Adaptation	66
4	Conclusion	69
	References	72
III	Appendices to Scientific Article	73
A	Computing Cross-correlation	75
A.1	Cross-product, Cross-covariance, and Cross-correlation	75
A.2	Cross-correlation of a Finite Discrete Signal	76
B	XC-AHSC: Robust Protection	79
C	Experiment Design	83
D	Experiment Briefing and Consent Form	85
E	Van der Laan Questionnaire	91

Nomenclature

List of Abbreviations

A-HSC	Adaptive Haptic Shared Controller
ADAS	Advanced Driver-Assistance Systems
CE	Controlled Element
DFT	Discrete Fourier Transform
F0	Fixed pursuit controller
F6	Fixed preview controller
FDC-HSC	Four Design Choices Haptic Shared Controller
FDCA	Four Design Choices Architecture
FM	Full Model
HCR	Human Compatible Reference
HO	Human operator
HSC	Haptic Shared Controller
iDFT	Inverse Discrete Fourier Transform
LoHA	Level of Haptic Authority
LoHS	Level of Haptic Support
MF	Model-Fitting
NH	Non-haptic Condition
PR	Preview
PS	Pursuit
RM	Reduced Model
SI	Single-integrator
SoHF	Strength of Haptic Feedback
VAF	Variance Accounted For
X-corr	Cross-correlation
XC-AHSC	Cross-correlation Adaptive Haptic Shared Controller

List of Symbols

$\Delta\tau_a^{HSC}$	Adaptive look-ahead time increment
ϵ_r	Reference error
ω_{nms}	Neuromuscular dynamics natural frequency
$\omega_{t,d}$	Forcing function frequency
$\phi_{t,d}$	Forcing function phase
τ_a^{HSC}	Adaptive look-ahead time
τ_p	Preview time
τ_v	Human time delay
τ_{HO}	Human look-ahead time
τ_{HSC}	Controller look-ahead time
ζ_m	Time shift at peak cross-correlation
ζ_{nms}	Neuromuscular dynamics damping ratio
$A_{t,d}$	Forcing function amplitude
b	Stick damping
C_{xx_r}	Cross-covariance between state and reference state
e	Error
f_Δ	Simulation update rate
f_d	Disturbance
f_t	Target
H_f^{HO}	Target low-pass filter
H_{nms}	Neuromuscular and stick dynamics
I	Stick inertia
k	Stick stiffness
K_a	Adaptation gain
K_f	Target low-pass filter gain
K_n	Remnant gain
K_R	Rejection gain

K_{e^*}	Van der El aiming point error gain (SI)	T_{HO^*}	Nominal Van der El human torque
K_{stick}	Stick gain	T_{HO}	Human operator torque
K_{xx_r}	Cross-correlation between state and reference state	T_{HSC^*}	Haptic torque feedback
N_b	Buffer length	T_{HSC}	Haptic shared controller torque
q_{yy}	yy -th percentile	$T_{l,f}$	Target low-pass time constant
R_{xx_r}	Cross-product between state and reference state	u	Input to controlled element
T_b	Buffer time	x	State
T_R	Rejection torque	x_r	Reference state
T_r	Reference torque	z^{-1}	Delay one step

List of Figures

1.1	Thesis outline	26
2.1	Intuitive form of Van der El's preview model [17]	28
2.2	Van der El's framework for short and long preview [17]	30
2.3	Look-ahead time adaptation [13]	30
2.4	LP-filter time constant adaptation [13]	30
2.5	FDCA design philosophy [16]	31
2.6	Haptic shared control setup [6]	32
2.7	HO Model [6]	33
2.8	Four Design Choices Haptic Shared Controller [6]	34
2.9	Tracking error and control activity [6]	35
2.10	Torque magnitudes [6]	35
2.11	Conflict time [6]	35
2.12	Conflict torque [6]	35
2.13	Equivalent FDC-HSC model without HO adaptation	36
2.14	HSC error and control activity without HO	37
2.15	Reference state phase shift, HSC without HO	37
2.16	Reference state alignment, HSC without HO	38
2.17	Reference state alignment, HSC without HO, $\tau_v = 0$	38
3.1	Full and simplified inverse-stick model T_{LoHS}	41
3.2	Validation tracking error	41
3.3	Validation conflict torque	41
3.4	Reduced model (RM) for parameter estimation	42
3.5	Full model for parameter estimation	43
3.6	System identification results	43
3.7	Personalised simulations tracking error	44
3.8	Personalised simulations conflict	44
3.9	Experiment target and torque for $\tau_{HSC} = 0.0s$	46
3.10	Experiment T_{SoHF} and T_{LoHS} for $\tau_{HSC} = 0.0s$	47
3.11	Conflict regions between two sine waves with a phase shift	47
3.12	Experiment T_{SoHF} and T_{LoHS} for $\tau_{HSC} = 0.0s, 0.8s$	48
3.13	Adaptive HO model	49
3.14	Rejection block structure	49
3.15	Simulated versus experiment torques for $\tau_{HSC} = 0.0s$	50
3.16	Simulated versus experiment torque for an adaptive HO model, $\tau_{HSC} = 0.0s$	50
3.17	Conflict as a function of varying rejection, $\tau = 0.0s$	51
3.18	VAF for rejection model	51
3.19	Error and conflict for rejection model and varying τ_{HSC}	51
3.20	Test case	52
3.21	Adaptive controller structure	53
3.22	Adaptive controller cost and policy	53
3.23	Inertial A-HSC example run 1	54
3.24	Inertial A-HSC example run 2	55
3.25	Inertial A-HSC look-ahead time, $N=100$	56
3.26	Inertial A-HSC histogram, $N=100$	56
3.27	Model-fitting controller structure	56
3.28	Model-fitting without $0.1s$ restriction	57
3.29	Model-fitting with $0.1s$ restriction	57

3.30 MF-1 look-ahead time, N=100	58
3.31 MF-1 histogram, N=100	58
3.32 MF-2 look-ahead time, N=100	58
3.33 MF-2 histogram, N=100	58
3.34 Cross-correlation A-HSC structure	59
3.35 Cross-correlation adaptive controller block	60
3.36 X-corr look-ahead time, N=100	60
3.37 X-corr histogram, N=100	60
3.38 Adaptive look-ahead time mean and standard distribution per controller	61
3.39 Comparing averaged look-ahead times for A-HSCs (N=3x100)	61
3.40 Convergence of A-HSCs (N=3x100)	62
3.41 Effects of varying transient behaviour, N=3x100	63
3.42 XC-AHSC look-ahead time for slow sigmoid, N=100	64
3.43 XC-AHSC histogram for slow sigmoid, N=100	64
3.44 XC-AHSC look-ahead time for varying K_r (N=5x100)	64
3.45 XC-AHSC look-ahead time for varying K_r , reversed (N=5x100)	65
3.46 XC-AHSC reversed adaptive look-ahead time, N=100	65
3.47 XC-AHSC reversed histogram, N=100	65
3.48 Reference state, phase lag and adaptive look-ahead time during plateau	66
3.49 Effect of buffer size on adaptive look-ahead time	66
A.1 Adaptive look-ahead time loop	75
A.2 Computing estimated cross-product/correlation of simulated x, x_r ($N = 700$)	78
B.1 Normal behaviour	79
B.2 Erratic behaviour	79
B.3 Protected behaviour	79
B.4 Inadvertent plateauing during adaptation because of \hat{K}_{min} protection	80
B.5 Effect of \hat{K}_{min} on average look-ahead time, $N = 2 \times 5 \times 16$, non-haptic	80
B.6 Effect of K_{min} on average look-ahead time, $N = 2 \times 5 \times 16$, $K_{HSC} = 1$	81
C.1 Run segments	83

List of Tables

2.1	HO model settings	29
2.2	Target and disturbance signals [6]	33
2.3	FDC and HO gains	34
3.1	Simplified inverse stick VAF	40
3.2	Grand means RM versus FM	44
3.3	Inertial A-HSC mean and standard deviation	55
3.4	MF-1 mean and standard deviation	58
3.5	MF-2 mean and standard deviation	59
3.6	Cross-correlation mean and standard deviation	60
3.7	Adaptive look-ahead time mean and standard distribution per controller	61
3.8	Convergence of A-HSCs (N=3x100)	62
3.9	Sigmoid parameters	63
3.10	XC-AHSC Slow sigmoid mean and standard deviation	63
3.11	XC-AHSC reversed mean and standard deviation	65
C.1	Experiment conditions	83

Part I

Scientific Article

Implementing an Adaptive Haptic Shared Controller in Pursuit and Preview Tracking Tasks

Max McKenzie

Abstract—Haptic shared controllers (HSCs) are a promising solution to prevent human over-reliance on automation during tasks such as car driving. However, research has shown that if the HSC is tuned incorrectly, then there is a risk of haptic conflicts between the human and HSC. To address this challenge, this paper presents the design and implementation of a novel adaptive HSC that continuously adjusts its look-ahead time. By estimating the time shift between the reference state of the HSC and the actual state, the HSC adapts to the look-ahead time of the human it is interacting with. Results from a human-in-the-loop experiment show that the novel HSC achieves similar subjective ratings as a fixed preview HSC, as well as a significant improvement over a fixed pursuit HSC. Going from pursuit to preview, objective experiment data shows that as the adaptive HSC adjusts its look-ahead time, haptic conflicts are reduced and tracking performance is increased. The presented findings are a step forward in designing haptic support systems with high chances of user acceptance. The proposed adaptive look-ahead algorithm provides a new method for online estimation of human look-ahead time, with or without a HSC in-the-loop.

Index Terms—Shared control, preview control, haptic feedback, look-ahead estimation

I. INTRODUCTION

WITH the advent of self-driving cars, societies are becoming increasingly reliant on automatic systems to control vehicles. However, this trend has led to concerns about risks related to system failure, or humans becoming dangerously inattentive [1–4].

As an alternative to full automation, *haptic shared control* aims to offset risks related to over-compliance by making the human operator (HO) partially responsible for steering the vehicle [5]. By requiring that the HSC and HO share the same physical interface (e.g., side-stick), both parties are forced to continuously negotiate with each other in terms of control input. A notable advantage of this method is that the HO is always aware of the intentions of the system through haptic feedback [6, 7]. Should the HO not be satisfied with the felt intentions of the system, then the HSC’s input can *immediately* be inhibited by impeding it from moving the physical interface.

Although a key aim of haptic shared control is to reduce human control effort, HSCs are liable to cause conflict and annoyance on the part of the HO [8, 9]. Paradoxically, it is possible that even when a HSC improves overall control performance, significant discomfort still occurs as a consequence of misaligned intentions or control strategy [10].

Recently developed human operator models have paved the way for the design of HSCs with a high chance of user acceptance. Specifically, Van der El’s human controller model for preview tracking tasks [11] has led to a better understanding of HO behaviour during car driving. Through

the use of HO models, the HSC can be configured to mimic the HO’s control strategy, hereby reducing the chance of haptic conflicts [10]. For preview tracking tasks, HOs have been found to use an aiming point situated up ahead along the previewed target signal [11]. The amount of time that the aiming point is situated up ahead into the future is referred to as the *look-ahead* time.

To investigate human sensitivity to varying controller settings, a previous study by Span et. al aimed to quantify the impact of varying HSC look-ahead time on conflicts and tracking performance. Using the Four Design Choices Architecture (FDC-HSC) [12], a region of optimal HSC look-ahead times between 0.7 – 0.8s was identified for which conflicts are minimised and performance is maximised [10].

Despite the discovery of a range of optimal look-ahead times for a FDC-HSC controller, it is still unclear exactly why haptic conflicts arise, and how they should best be resolved [13, 14]. In general, the challenge of investigating HO-HSC interaction lies in the intertwined dynamics that occur between the HSC and HO through physical feedback [15]. Since both agents are either directly or indirectly influenced by one another, it is impossible to study HO adaptation without considering the potential effects of these dynamics on the entire system. On the other hand, understanding human adaptation could play a crucial role in designing a new type of HSC, which adjusts its settings to resolve conflicts automatically.

The goal of this study is to develop an adaptive haptic shared controller (A-HSC) that adapts to its human counterpart in real time to reduce haptic conflicts. As a first step, simulations were used to test an A-HSC design that varies its look-ahead time based on the time shift between the HSC’s reference state and the actual state. Secondly, the A-HSC was tested in a human experiment with 16 participants. Subjects performed a tracking task during which the displayed preview time switched from zero (i.e., pursuit), to one second of preview, or vice versa. For reference, the A-HSC was tested alongside two types of fixed FDC-HSCs with (1) pursuit settings, and (2) preview settings.

The paper is structured as follows. In Section II the control task and context surrounding the paper is explained. Thereafter, Section III describes the effects of look-ahead time on HO-HSC interaction and investigates the nature of haptic conflicts. The design methodology of the A-HSC is given in Section IV. Thereafter, Section V shows the results of offline simulations that test the A-HSC structure under varying conditions. Section VI provides an overview of the human-in-the loop experiment, after which experiment results are presented in Section VII. Finally, the results and conclusions are presented in the final two sections.

II. BACKGROUND

This section provides the necessary background and context for the project by describing (A) the control task, (B) the human operator model, and (C) the Four-Design-Choices architecture for HSCs.

The general setup of a HSC in a preview tracking task is illustrated in Fig. 1, where the dashed line indicates the influence of haptic feedback on the HO. To extend the applicability of simulations used by Span [10], Section III proposes a method to explicitly model HO rejection of HSC forces, using simulated haptic feedback (dashed line in Fig.1).




Fig. 1: Haptic shared control in a preview tracking task [10]

A. Control Task

The control task consists of both the HSC and HO using the previewed target signal, f_t , displayed up until a preview time of τ_p seconds into the future, as illustrated by Fig. 1. HOs are to minimise the lateral error between the target and state, $e(t) = f_t(t) - x(t)$. This can be considered a rudimentary form of car driving, making this task a relevant case study for the development of (haptic) driver support systems [16].

1) *Signals*: The range of preview information used by the HO between the current time t and τ_p is denoted by $f_t([t, t + \tau_p])$, as shown in Fig. 1. On the other hand, the HSC is configured to use a single aiming point τ_{HSC} seconds up ahead along the previewed target, $f_t(t + \tau_{HSC})$. This is related to the HSC's design and how HOs use preview information, as is described in Section II-B.

$T_{HO}(t)$ and $T_{HSC}(t)$ are the torques applied to the physical interface by the HO and HSC, respectively. The resulting position of the side-stick or steering wheel is thereafter used

to drive the controlled element (CE) through the input $u(t)$. The HO is fed back information on the state by observing its position on the display, whereas the HSC uses the raw signal itself to compute T_{HSC} . Finally, an unobservable disturbance $f_d(t)$ is applied to the CE, making the activity both a target-following- and disturbance rejection task.

2) *Forcing Functions*: The target and disturbance functions are defined to be the sum of sines as determined by Eq. 1, with amplitude A , frequency ω and phase shift ϕ .

$$f_{t,d}(t) = \sum_{i=1}^{10} A_{t,d}[i] \sin(\omega_{t,d}[i]t + \phi_{t,d}[i]) \quad (1)$$

The frequencies of the sines are selected from the same set used by Van der El to study manual control behaviour in preview tracking tasks [11, 17, 18]. However, for this study, single bands of input frequencies are used as opposed to double bands, which are commonly used in order to measure a model's linearity [11]. Each sine is an integer multiple k_i of the base frequency of 0.0524 rad/s, which corresponds to 120s of measurement time.

Table I lists the amplitudes, frequencies, and phase shifts for one realization. A total of four realizations of the target function were used to prevent pattern recognition during the experiment, each varying only in terms of phase shift.

3) *Bandwidth*: The bandwidth of the target signal used in this study is $\omega_i = 1.5$ rad/s, corresponding to the highest frequency at which the signal has significant power. This aligns with the target signal characteristics of both Van der El's and McRuer's landmark experiments [11, 19], facilitating a comparison between solo HO behaviour and a shared control situation. Furthermore, the target bandwidth is lower than the expected crossover frequency for single integrator (SI) CE dynamics [19]. As a result, the HO is expected to be capable of adequately tracking components of the target that lie within the signal bandwidth.

4) *Disturbance and Controlled Element*: By injecting a disturbance at the output of the controlled element, the HO is forced into using a feedback control strategy as opposed to tending towards a predominantly feedforward approach [11]. This is due to the fact that the disturbance cannot be

TABLE I: Multisine target and disturbance functions parameters

		Target signal f_t			Disturbance signal f_d			
i	k_t	A_t cm	ω_t rad/s	ϕ_t rad	k_d	A_d cm	ω_d rad/s	ϕ_d rad
1	3	0.731	0.157	4.488	4	0.292	0.209	0.241
2	5	0.731	0.262	5.699	7	0.292	0.367	1.669
3	8	0.731	0.419	1.373	9	0.292	0.471	1.899
4	11	0.731	0.576	5.472	13	0.292	0.681	1.295
5	19	0.731	0.995	1.331	22	0.292	1.152	3.982
6	29	0.731	1.518	5.257	31	0.292	1.623	4.496
7	47	0.073	2.461	5.399	51	0.029	2.670	3.365
8	77	0.073	4.032	3.289	79	0.029	4.136	0.469
9	143	0.073	7.488	2.999	147	0.029	7.697	0.964
10	263	0.073	13.77	5.591	267	0.029	13.98	4.296

anticipated, and can therefore only be corrected *after* it has influenced the position of the state. Throughout this study the controlled element dynamics are single-integrator with $H_{CE}(j\omega) = 1.5/(j\omega)$.

5) *Physical Interface*: The HO and HSC both share control authority of a side-stick, modelled as a mass-spring-damper system according to Eq. 2, with $I = 0.01$ kg/m², $b = 0.22$ Nms/rad, $k = 3.58$ Nm/rad and $K_{stick} = 25.4$ cm/rad.

$$H_{stick} = \frac{K_{stick}}{I(j\omega)^2 + bj\omega + k} \quad (2)$$

B. Human Operator Model

To simulate human behaviour in a preview tracking task, work by Van der El et al. [11] is used to model the HO as a quasi-linear controller. A key benefit of this model is that it can capture both pursuit and preview HO control behaviour using one single framework, illustrated by Fig. 2.

1) *Preview*: In preview tracking tasks ($\tau_p = 1.0$ s), HOs effectively track an aiming point τ_{HO} seconds into the future along a low-pass filtered version of the target signal. By minimising the error between the state and this self-generated aiming point, HOs are able to cancel out their inherent delays without additional control effort [17].




Fig. 2: Human operator model [10]

Fig. 2 shows how a low-pass filter H_f^{HO} is applied to a point $f(t + \tau_{HO})$ along the original target, where τ_{HO} is the *look-ahead time*. The low-pass filter is defined by Eq. 3 with gain K_f and time constant $T_{l,f}$.

$$H_f^{HO}(j\omega) = \frac{K_f}{1 + T_{l,f}(j\omega)} \quad (3)$$

For SI CE dynamics, the error between the resulting aiming point $f_{i,f}^*$ and the state x is minimised by a gain K_{e^*} [20]. Then, $H_{nms-stick}^{HO}$ models the entire system's mechanical limitations (stick + neuromuscular) by a second-order filter. HO time delays are accounted for by $H_{nms-stick}^{HO}$ in combination with a pure time shift τ_v , represented by $e^{-\tau_v j\omega}$ in the frequency domain. The second order filter representing lumped neuromuscular and stick dynamics is given by Eq. 4, with natural frequency ω_{nms} and damping ratio ζ_{nms} .

$$H_{nms-stick}^{HO}(j\omega) = \frac{\omega_{nms}^2}{(j\omega)^2 + 2\zeta_{nms}\omega_{nms}j\omega + \omega_{nms}^2} \quad (4)$$

Since Van der El's model aims to predict HO control activity in terms of CE input $u(t)$ as shown in Fig. 2, a simplified inverse-stick model $H_{stick}^{-1} = k/K_{stick}$ is applied to the output u to determine the underlying torques T_{HO} .

The model parameters for a preview HO model are given by Table. II, with $\tau_{HO} = 0.68$ s for $\tau_p = 1.0$ s. Preview

information beyond the *critical look-ahead time* of around $\tau_{crit}^{HO} = 0.6 - 0.7$ s is not used by the HO, since only a certain amount of look-ahead is needed to cancel out their lags. Therefore, when a full preview display is used with $\tau_p = 1.0$ s, the HO's look-ahead time tends to be approximately $\tau_{HO} = 0.68$ s.

2) *Pursuit*: When the display's preview time τ_p is reduced to zero, the control task reverts to a pursuit tracking task. This means that $\tau_{HO} = \tau_p = 0$, which causes the HO to adopt a more aggressive control strategy due to the lack of preview information. This is reflected by the higher gain K_{e^*} in Table II for the pursuit model relative to the preview model. Furthermore, since HOs are unable to filter the upcoming target in pursuit, $T_{l,f}$ reduces to 0.01s, essentially making the low-pass filter inactive.

3) *Parameter Settings*: Table II shows the settings that were used to simulate both the HO and HSC. In cases where the situation changed from pursuit to preview or vice versa, model settings were switched discontinuously. Finally, for both pursuit and preview models a remnant n is injected by low pass filtering white noise at a power of 1 cm², with a break frequency of 3.5 rad/s and a gain of $K_n = 0.223$.

TABLE II: Model and display parameters

	τ_p [s]	τ_{HO} [s]	$T_{l,f}$ [s]	K_{e^*} [-]	K_f [-]	τ_v [s]	ω_{nms} [rad/s]	ζ_{nms} [-]
Pursuit	0	0	0.01	2.0	1.0	0.26	10.5	0.35
Preview	1.0	0.68	0.2	1.25	1.0	0.26	10.5	0.35

C. Four-Design-Choices Haptic Shared Controller

This project makes extensive use of the FDC architecture [12] to implement a controller that has the best chances of user acceptance. The structure of the FDC-HSC used in this study is given in Fig. 3, and consists of two main parts. The first pertains to the generation of a *human compatible reference* (HCR), which can be seen as the controller's internal model for how it expects the HO to act. The second part, *haptic torque generation*, uses inputs from the HCR to generate a haptic feedback torque that is then sent to the physical interface. In this study, Van der El's model was used to generate the HCR with the parameters given in Table II, and the FDC gains were set to $K_{LoHS} = 0.6$, $K_{SoHF} = 0.8$, $K_{LoHA} = 1.0$. The four design choices are:

- 1) **Human compatible reference (HCR)** is the HSC's internal model for the HO's expected behaviour. Based on the HCR, the HSC determines the reference state x_r , and reference torque T_r .
- 2) **Level of haptic support (LoHS)** uses the HCR's reference torque T_r to provide open-loop guidance to the user. Since this type of support is only dependent on the HCR (and independent of the HO), it is useful to think of this torque as being 'suggestive' rather than corrective.
- 3) **Strength of haptic feedback (SoHF)** serves as a corrective torque that minimises the difference between the reference and state $\epsilon_r = x_r - x$, shown in red in Fig. 3. This difference is denoted the *reference error*.




Fig. 3: Four Design Choices HSC [10]

- 4) **Level of haptic authority (LoHA)** determines the HSC's control authority, i.e., to what extent it is capable of overriding the HO's inputs. For the FDC-HSC structure shown in Fig. 3, the LoHA is tuned by changing the gain K_{LoHA} .

Previous experiment results have shown that the SoHF channel is significantly more prone to inducing conflicts than the LoHS channel [10]. To explain why this is, consider the reference error ϵ_r upon which K_{SoHF} acts. The reference error can be influenced by three key factors:

- 1) The HO's remnant, influencing x
- 2) The disturbance $f_d(t)$, influencing x
- 3) Misaligned or sub-optimal HCR, influencing x_r

For factors (1) and (2), activation of the SoHF channel is actually desirable, since it will suppress effects of the remnant and disturbance signal, which are both detrimental to performance. However, for factor (3) it is not desirable for the SoHF channel to become active, since in this case the HO's control strategy is misaligned with the HSC. The resulting difference in HO-HSC intentions will inevitably manifest itself in haptic conflicts, with the root cause being the HCR itself. The solution is to develop methods for identifying *sustained, systematic conflicts* that occur as a consequence of a misaligned HCR, so that the HCR settings can be adjusted.

It is crucial to understand that the SoHF channel can only ever steer the state towards the reference x_r , and that this happens independently of the target signal. Therefore, any significant mismatch between the reference state x_r and the target f_t will inevitably lead to a reduction in HSC tracking performance, since the HSC is steering towards a sub-optimal reference state. For example, if the HCR is setup in such a way that it lags the target, the only way this can be resolved is through 1) HO influence, or 2) adaptation of controller settings. Therefore, an *external* feedback loop outside of the FDC-HSC is required for effective HSC adaptation.

Ensuring that the reference state is designed in such a way that it 1) behaves in a way that is compatible with the HO, and 2) aligns with the target signal, is the primary responsibility of designers. The goal of the novel controller is to automate this process in real-time by aligning the reference state with the HO, both in terms of target signal and control behaviour.

III. HAPTIC TORQUE REJECTION

Given that the goal of the novel controller is to reduce HO-HSC conflicts by varying the HSC's look-ahead time,

first an analysis is performed on the effects of varying this parameter and its influence on HO-HSC conflicts. Thereafter, a structure is proposed for modelling haptic torque rejection. The proposed model is validated by comparing its result with experimental data gathered by Span [10].

A. Effects of Look-ahead Time

Work by Span et. al [10] investigated the effects of varying the look-ahead time of an FDC-HSC controller with Van der El's model as HCR, as shown in Fig. 3. The resulting experimental data is given in Fig. 4, compared with simulations that assumed no haptic interaction.

Both experimental and simulated results show that when using a fixed preview display, if the HSC's look-ahead time varies significantly from $0.7 - 0.8s$, then conflicts increase and tracking performance decreases. These results correspond with the HO's critical preview time as estimated by Van der El's model, which predicts that the HO look-ahead time for a preview display will be around $0.6 - 0.7s$ [17].

The simulation results shown in Fig. 4 were generated using a model that assumed no direct interaction between the HO and HSC (inactive dashed line in Fig. 1). This is equivalent to assuming that both the HO and HSC operate as if they have their own side-stick, influencing the CE independently of one another. Despite this approximation, the model successfully predicts the optimal values of HSC look-ahead time ($\tau_{HSC} = 0.7 - 0.8s$) in terms of conflict resolution and tracking performance. It is only when τ_{HSC} deviates substantially from $0.7 - 0.8s$ that the model underestimates the magnitude of conflicts, as well as tracking performance. This suggests that through physical interaction with the side-stick, HOs are capable of increasing performance at the cost of conflict.




Fig. 4: Tracking error and conflicts for a FDC-HSC (experimental data taken from [10]), NH refers to 'non-haptic'

Fig. 5 shows the effect of τ_{HSC} on the HSC's own internal reference state x_r . For $\tau_{HSC} = 1.2s$, the reference state leads

the target signal f_t , whereas for $\tau_{HSC} = 0.0s$ the opposite happens, causing x_r to lag behind the target. These results clarify why tracking performance gets worse when τ_{HSC} deviates from $\tau_{HSC} = 0.7 - 0.8s$. In these cases, the HSC is steering towards a sub-optimal reference state x_r .

Fig. 4 indicates that, in practice, HOs intuitively correct for the sub-optimal reference state to maximise tracking performance at the cost of conflict torques. The experimental data in Fig. 6 further supports this hypothesis by revealing that when the reference state x_r is aligned with the target f_t , conflicts significantly reduce.




Fig. 5: Effect of HSC look-ahead time on reference state

In the event of a misaligned reference state as shown in Fig. 5, HO and HSC torques have a tendency to amplify each other. This is due to the nature of the SoHF channel and how HOs correct for the HSC's sub-optimal inputs. As soon as a HO steers the state away from the misaligned reference, the reference error $\epsilon_r = x - x_r$ increases. As a consequence, the HSC's SoHF channel will apply torques to steer the state back towards the misaligned reference, in the opposite direction to the HO's intentions. The more successful the HO is in correcting for the sub-optimal reference state by minimising the true error $e = f_t - x$, the more the HSC will try to steer the state x back towards the reference state x_r .

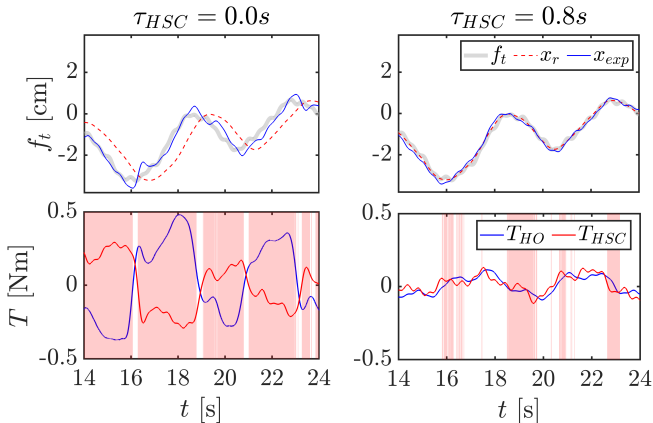


Fig. 6: Measured experiment state, reference state and conflicts for fixed preview display ($\tau_p = 1.0s$), data taken from [10]

B. Modelling Rejection Torques

1) *Torque Rejection Loop*: To account for the discrepancy between the modelled and experimental data in Fig. 4, an amendment is made to the structure shown in Fig. 2. By constructing a *torque-rejection* loop as shown in Fig. 7, a

rejection torque T_R is superimposed onto the nominal torque T_{HO*} , predicted by Van der El's model according to Fig. 2 (T_{HO}). The goal of this structure is to allow the simulated HO to cancel out the torques of the HSC, but only under certain conditions. Note: the rejection torque T_R is not to be confused with the reference torque T_r of the HCR.




Fig. 7: Torque rejection model

The contents of the rejection block are shown in Fig. 8. $T_{HSC*} - T_{HO*}$ represents the difference between the HO's original plan of action, and the HSC forces as felt through haptic feedback. A moving root-mean-square of $T_{HSC*} - T_{HO*}$ is calculated using a sliding window of $5s$, so that only *sustained* differences in intentions will lead to rejection. When a sustained difference in intentions occurs, a negative gain K_R is used to cancel out the original HSC torques by applying force in the opposite direction to the HSC's inputs.

To prevent over-compensation, a saturation limiter is used to ensure that the rejection gain K_R is always between 0 (no rejection) and -1 (all HSC torques are rejected). The saturated gain is then multiplied by the original HSC torque and passed through H_{nms} (Eq. 4), resulting in the final HO rejection torque to cancel the HSC's inputs. In this model the HO is assumed to have perfect knowledge about HSC inputs through haptic feedback, i.e., $T_{HSC} = T_{HSC*}$.




Fig. 8: Torque rejection loop

2) *Haptic Conflicts for a Lagging HSC*: The initial effects of modelling rejection are shown in Fig. 9, where the look-ahead times are set to $\tau_{HSC} = 0.0s$ and $\tau_{HO} = 0.6s$. For the model without rejection, the state is situated somewhere between the misaligned reference and target, causing sub-optimal performance. Furthermore, conflicts are underestimated relative to experimental data. These two results are consistent with Fig. 4, and fully explain the observed discrepancy between experimental and modelled data.

By accounting for HO rejection of HSC torques, the simulated HO and HSC torques both become larger in magnitude, due to their tendency to amplify each other. Furthermore, the simulated state tracks the target more closely. By using the rejection model, the simulated torques and state correspond much better to measured experimental data than the original model, as shown in Fig. 9.

3) *Haptic Conflicts for a Leading HSC*: Fig. 10 shows the difference between the two models for $\tau_{HSC} = 1.2s$. Although

the HO torques between the models are roughly the same in magnitude, both undershoot the experimental data. This is why the variance accounted for becomes negative as shown in Fig. 13: the difference between the modelled torques and experimental data is large relative to the magnitude of the modelled signals themselves.

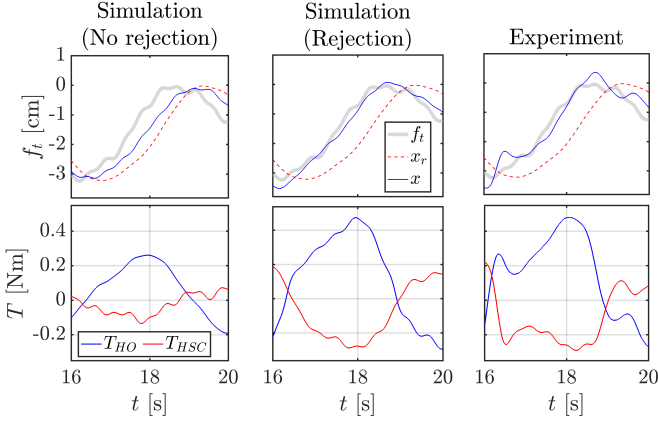


Fig. 9: Effects of rejection loop on modelled state, reference state, and torques ($\tau_{HSC} = 0.0s$, $\tau_{HO} = 0.6s$, $\tau_p = 1.0s$, $K_n = 0$, $K_R = -4$), experimental data taken from [10]

Comparing Fig. 10 with Fig. 9 shows that for $\tau_{HSC} = 0$, HO rejection is more pronounced than for $\tau_{HSC} = 1.2s$, suggesting that it is easier for HOs to reject a lagging HSC ($\tau_{HSC} = 0$), compared to a leading HSC ($\tau_{HSC} = 1.2s$). This is further supported by the tracking error, for $\tau_{HSC} = 0$ the average error is $RMS_e = 0.278cm$, whereas for $\tau_{HSC} = 1.2s$ the error is $RMS_e = 0.319cm$. Because HOs reject the HSC more strongly when it is lagging, this also leads to a higher quality-of-fit as shown in Fig. 13.

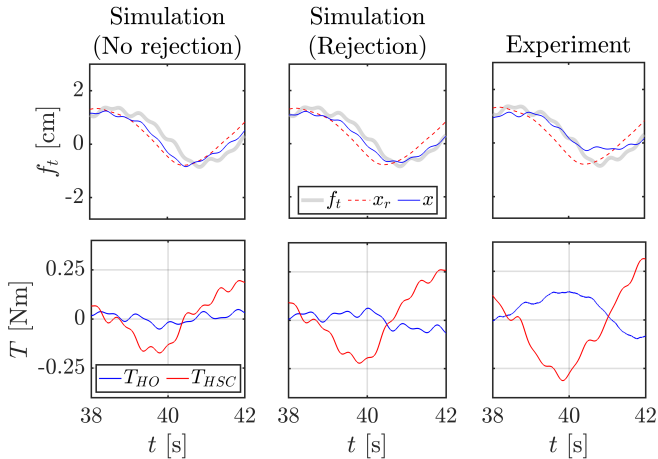


Fig. 10: Effects of rejection loop on modelled state, reference state, and torques ($\tau_{HSC} = 1.2s$, $\tau_{HO} = 0.6s$, $\tau_p = 1.0s$, $K_n = 0$, $K_R = -4$), experimental data taken from [10]

4) *Comparing Models*: Fig. 11 again shows the match between experimental and simulated results, now using the rejection loop with $K_R = -4$ and a 5s moving RMS window. For all settings of τ_{HSC} , both the average error and conflict

match experimental data better, or similar to the original model (no rejection). The remnant gain for the original model is $K_n = 0.223$, whereas for the rejection model it reduces to $K_n = 0.169$. This effectively means that a portion of the remnant (assumed to be entirely arbitrary) has now been explained, hence reducing its contribution.




Fig. 11: FDC-HSC tracking error and conflicts (rejection-loop model), experimental data taken from [10]

The Variance Accounted For (VAF) is used to determine a model's quality-of-fit, and is calculated according to:

$$VAF = \left(1 - \frac{\sum_{k=1}^N |x(k) - \hat{x}(k)|^2}{\sum_{k=1}^N \hat{x}^2(k)} \right) \times 100\% \quad (5)$$

Fig. 12 shows the VAF of the state x , both for the rejection model and the original. The experimental data consisted of a total of 270 experiment runs of 120s, performed by 10 subjects. The rejection model's quality of fit is either similar or slightly better than the model without rejection. The only significant difference is for $\tau_{HSC} = 0$, where the VAF of the rejection model is slightly higher.




Fig. 12: VAF for state ($N = 9 \times 3 \times 10$, $K_R = -4, 0$)

Fig. 13 further investigates how well the models are able to capture the underlying HO torques and control activity. For $\tau_{HSC} > 0.8s$ the VAF becomes negative, indicating that the variance of the mismatch is higher than the variance of the modelled signal. For these settings both models fall short and are unable to fully capture the dynamics of HO-HSC interaction.

Between $\tau_{HSC} = 0.3 - 0.8s$ both models perform roughly the same. However, for $\tau_{HSC} = 0$ the rejection model's VAF is 79%, which is significantly higher than the original model's 38% quality-of-fit for $\tau_{HSC} = 0$.

The takeaway is that both models are capable of adequately predicting the state x , but are less suited for modelling specific HO torques when $\tau_{HSC} > 0.8s$. This suggests that HO's use an unknown strategy when it comes to HSC's with a higher look-ahead time than their own.




Fig. 13: VAF for HO torques ($N = 9 \times 3 \times 10$, $K_R = -4, 0$)

IV. CONTROLLER DESIGN

This section describes the methodology behind the design of the adaptive haptic shared controller. The fundamental hypothesis is that conflicts can be solved by aligning the controller's reference state with the intentions of the human. If there is a difference in intentions between the HO and HSC, it is assumed that the HO will always attempt to override the HSC's influence by countering its torques. By analysing differences between the reference state and the actual state, the look-ahead time can be adjusted accordingly. If x and x_r are in alignment, then the assumption is that the HO agrees with the inputs of the HSC. By designing the controller in this way, the root cause of conflicts is addressed head-on, hence circumventing the need for conflict metrics and thereby avoiding false positives (factors 1 & 2 in Section II-C).

A. Adapting Look-ahead Time

Fig. 14 shows the structure of the novel controller. The same FDC-HSC structure from Span et. al is retained, except that now its look-ahead time varies according to the 'adaptive X-corr' block. Taking x and x_r as inputs, this block determines what *adaptive* look-ahead time, τ_a^{HSC} , the FDC-HSC should use. From this point forwards τ_a^{HSC} refers to the look-ahead time used by an adaptive HSC, whereas τ_f^{HSC} refers to a fixed HSC look-ahead time.

At each instant there are only three possibilities for adaptation: (1) τ_a^{HSC} should stay constant, (2) τ_a^{HSC} should increase, (3) τ_a^{HSC} should decrease. If x and x_r are in alignment, then no conflicts are expected, hence condition (1) applies. Alternatively, if the reference state x_r lags behind the actual state x as shown in Fig. 5 ($\tau_f^{HSC} = 0.0s$), then this indicates that the look-ahead should be increased. By analogy, the look-ahead time should decrease when x_r leads x .

To know in which direction to adapt τ_a^{HSC} , the phase shift between x_r and x is used to measure lead or lag between the two signals. To this end, cross-correlation is used to estimate the similarity between the two signals as a function of varying relative time shift. This is illustrated by Fig. 15, which expands the 'adaptive X-corr' block of Fig. 14.

To estimate the time shift between x and x_r at each instant, two rectangular windows are used to store the last 7 seconds of data for the two signals ($T_b = 7s$). The cross correlation between the windowed signals, $K_{xx_r}(\zeta)$, is then computed as a function of time shift ζ . The estimated time shift corresponding to the maximum cross-correlation is denoted τ_m according to Eq. 6.

$$\zeta_m = \arg \max_{\zeta} [K_{xx_r}(\zeta)] \quad (6)$$

In turn, ζ_m is used to drive changes in adaptive look-ahead time, $\Delta\tau_a^{HSC}$. This is based on the principle that if there is a sustained, non-zero time shift between x and x_r , then the look-ahead time should either be increased or decreased.

A positive time shift ζ_m corresponds to x_r lagging x , which means that the HSC's look-ahead time should be *increased* by a positive $\Delta\tau_a^{HSC}$. Therefore, changes to the HSC's look-ahead time $\Delta\tau_a^{HSC}$ are defined to be proportional to the measured time shift, ζ_m . The gain K_a controls the rate of adaptation and determines the magnitude of $\Delta\tau_a$ for a given time shift. The larger K_a , the faster the look-ahead time will be updated.




Fig. 14: X-corr adaptive haptic shared controller

Note that the rate at which τ_a^{HSC} is updated by $\Delta\tau_a^{HSC}$ is directly proportional to the adaptation speed. This is because for a constant time-shift, if τ_a^{HSC} is updated by τ_a^{HSC} twice as often, then the final adaptive look-ahead time will change twice as fast. For this study, $K_a = 0.002$ and the update rate is 100 Hz.




Fig. 15: Adaptive X-corr block

Fig. 16 shows the simulation results for a single run, during which the HO model was set to discontinuously switch from pursuit to preview according to the parameters in Table II. As a consequence of varying HO behaviour, the look-ahead time of the HSC adapts from $\tau_a = 0s$ to $\tau_a = 0.64s$ by the end of the run. Note how adaptation is delayed by approximately seven seconds after $t = 50s$, which is due to the buffer length.

To mitigate the risk of the algorithm adapting its look-ahead time in situations where no clear time shift can be measured, a minimum threshold for correlation is set at $K_{xx_r} > 0.26$. This ensures that adaptation will only take place if a time shift between x and x_r can be measured reliably. The downside of this approach is that adaptation might unnecessarily be hindered, causing stagnation as seen around $t = 68s$ in Fig. 16.




Fig. 16: Simulated X-corr adaptive look-ahead time, cross-correlation peak, time shift, state and reference state ($K_a = 0.002$, $K_R = -4$, $K_n = 0.223$, $K_{xx_r} > 0.26$, $T_b = 7s$, $\tau_{HO} = 0s$ ($t < 50$), $\tau_{HO} = 0.68s$ ($t > 50$))

B. Adapting $T_{l,f}$ and K_{e^*}

Because the low-pass filter $T_{l,f}$ and gain K_{e^*} both significantly vary between pursuit and preview situations, it is important that changes to these parameters are also accounted for during adaptation. This is done by varying the contributions of the models given in Table II according to complementary sigmoid functions, as shown in Fig. 17. The preview model gain K_{prev} is given by Eq. 7 as a function of the adaptive look-ahead time, the pursuit model gain is defined by $K_{purs} = 1 - K_{prev}$.

$$K_{prev} = \frac{1}{1 + \exp(-33 \cdot (\tau_a^{HSC} - 0.186))} \quad (7)$$




Fig. 17: Model contributions

Fig. 18 shows the results of simulating a HSC that dynamically varies $T_{l,f}^{HSC}$ and $K_{e^*}^{HSC}$, compared with fixed HSCs with pursuit (PS) and preview (PR) settings according to Table II. For $t < 50s$ a PS-HO model was used with $\tau_{HO} = 0s$, after $t > 50s$ the HO model switched to PR with $\tau_{HO} = 0.68s$.

If τ_a^{HSC} is adapted whilst keeping the other HSC parameters at their PR values ('Fixed preview'), then τ_a^{HSC} tends to be overestimated as shown in Fig. 18. This is due to the fact that the adaptive controller is compensating for the lag introduced by its own low-pass filter and low gain. As a result,

for fixed PR parameters, the adaptive algorithm converges to $\tau_a^{HSC} = 0.25s$ during PS.

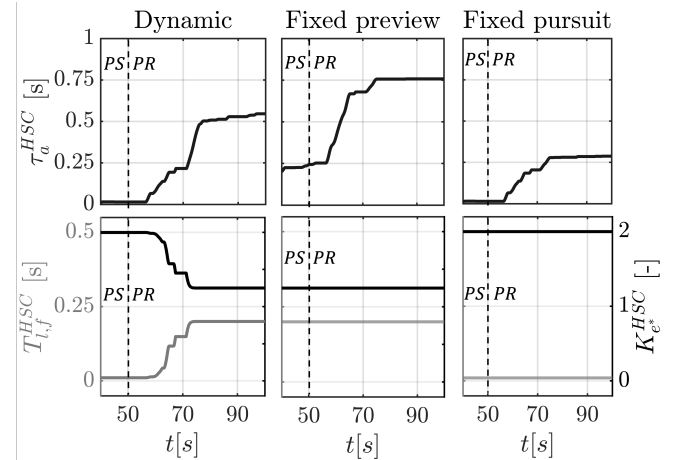


Fig. 18: Simulated effects of $T_{l,f}$ and K_{e^*} adaptation ($K_R = -4$, $K_n = 0$, $K_{xx_r} > 0.26$, $T_b = 7s$, $\tau_{HO} = 0s$ ($t < 50$), $\tau_{HO} = 0.68s$ ($t > 50$))

For fixed PS parameters ('Fixed pursuit') the opposite happens, as the adaptive controller now requires less lead compensation through look-ahead time. This is due to the fact that the HSC's low-pass filter is now essentially inactive ($T_{l,f}^{HSC} = 0.01s$), in addition to the HSC using a more aggressive control strategy ($K_{e^*}^{HSC} = 2.0$). As a result, look-ahead time stagnates prematurely around $\tau_a = 0.291s$ after transition from PS to PR.

V. CONTROLLER SIMULATIONS

This section briefly evaluates the results of Monte Carlo simulations designed to test the novel structure under varying conditions. First, the general distribution of the adaptive look-ahead time is analysed for random remnant seeds. This is to confirm that the adaptive HSC will converge to approximately the same look-ahead time for HO with varying remnants. Thereafter, the effects of adaptation gain are considered to gauge how sensitive results are to changes in K_a . Finally, the influence of rejection is assessed, to determine if HO rejection of HSC torques could significantly affect adaptation.

A. Distribution

Fig. 19 shows the distribution of 2×48 runs for three different forcing function realizations with random remnant seeds. At $t = 50s$ the HO model discontinuously switches from PS to PR, or vice versa, according to Table II. For PS to PR the maximum spread is encountered at $t = 80s$, where the look-ahead time ranges between $\tau_a^{HSC} = 0.12 - 0.66s$ across all realizations. By the end of the run all adaptive look-ahead times converge around a mean of $\tau_a = 0.64s$ between $\tau_a^{HSC} = 0.55 - 0.82s$.

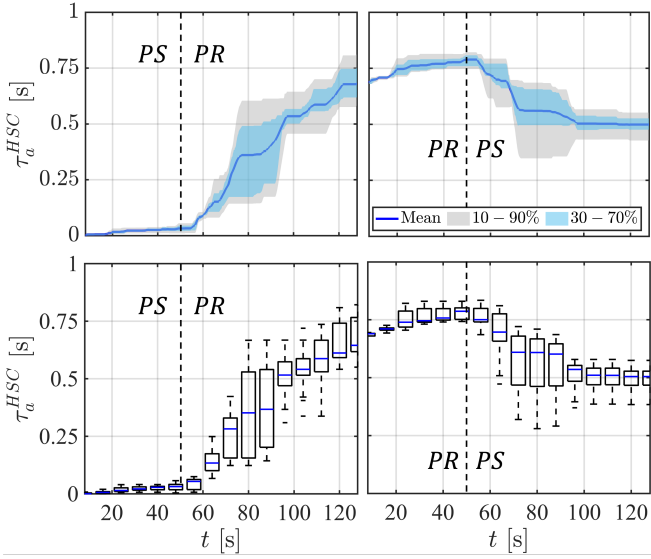


Fig. 19: Simulated adaptive look-ahead time ($N = 2 \times 48$, $T_b = 7s$, $K_R = -4$, $K_n = 0.169$, $K_{xx_r} > 0.26$, $K_a = 0.002$)

For PR to PS the maximum variance is also encountered at $t = 80s$, with $\tau_a^{HSC} = 0.28 - 0.71s$ around a mean of $\tau_a^{HSC} = 0.61s$. By $t = 128s$ all runs converge around a mean of $\tau_a = 0.50s$ between $0.38 - 0.57s$.

The asymmetry between PS-PR and PR-PS may be attributed to how the pursuit model acts upon the error. Once the HO model has switched to PS, the only option it has is to minimise the current difference between the target f_t and the state x (no preview). Just after $t = 50s$, the adaptive controller still retains its relatively high look-ahead time. As a result, the HSC will *outperform* the HO in terms of reaction time and will always ‘get there first’ in terms of minimising the error. Since this aligns with strategy of the HO model, there is less incentive on the part of the HO to influence the state, which would cause deviations from x_r (and adaptation).

The adaptation that does still occur in the PR-PS case in Fig. 19 is likely the result of rejection torques ($K_R = -4$), which explicitly superimpose conflict forces to cancel out effects of the HSC. This may cause the HO model to become ‘surprised’ when the HSC proposes torques that do not (yet) align with its own actions, hence causing it to reject the HSC’s input according to the structure laid out in Fig. 8. However, in practice, it is likely that HOs will recognise that HSC torques can be trusted to benefit tracking performance, since the HSC will be reacting faster than the HO. Trusting the HSC equates to complying with HSC torques, even though they may be unexpected (relative to the HO’s own, slower control efforts).

B. Adaptation Gain

Fig. 20 shows the effects of increasing the adaptation gain. Note that each line represents the averaged mean of 16 realizations, with higher K_a values resulting in faster adaptation. For PS to PR, Fig. 20 suggests that the adaptation gain has a significant effect on the final mean, with higher adaptation gains resulting in not only faster convergence, but also higher adaptive look-ahead times (τ_a^{HSC}). The same trends apply

going from PR to PS, where larger values for K_a result in larger τ_a^{HSC} deviations from the initial adaptive look-ahead time.

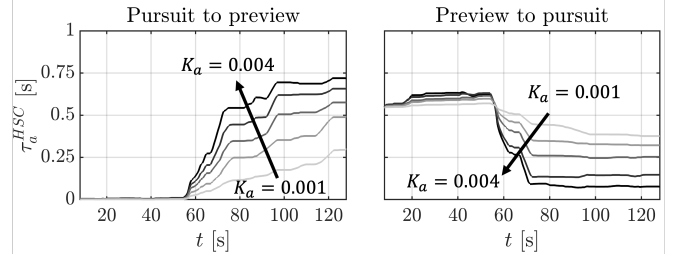


Fig. 20: Simulated effect of adaptation gain ($N = 2 \times 5 \times 16$, $K_R = -4$, $K_n = 0.223$, $K_{xx_r} > 0.26$)

C. Rejection

Fig. 21 shows the effect of increasing the magnitude of the rejection gain described in Section III. Going from PS to PR, increasing the rejection gain of the HO model has little effect, aside from a marginal improvement in adaptation speed. For the reverse condition, PR to PS, rejection plays a much larger role. Fig. 21 shows that after transition, τ_a^{HSC} stays approximately constant when rejection is low ($K_R = 0$). This is because the HO model is mostly complying with HSC torques, such that look-ahead adaptation does not take place. For $K_R = -8$, the adaptive look-ahead time significantly reduces from $\tau_a^{HSC} = 0.60s$ before transition, to $\tau_a^{HSC} = 0.27$ after transition.

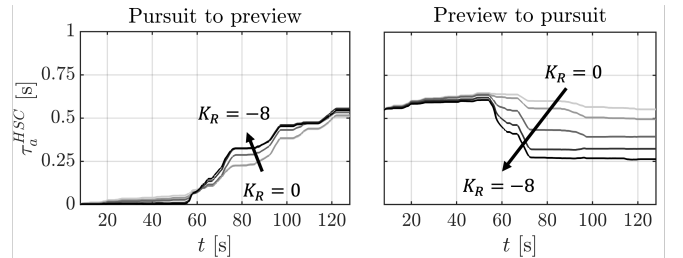


Fig. 21: Simulated effect of rejection gain ($N = 2 \times 5 \times 16$, $K_a = 0.002$, $K_n = 0.223$, $K_{xx_r} > 0.26$)

VI. HUMAN IN-THE-LOOP EXPERIMENT

In this section the experiment methodology and hypotheses are described. Sixteen right-handed subjects, aged 23-52 years old, were invited to participate in the tracking task laid out in Section II. The goal of the experiment is to evaluate the adaptive controller by using two fixed FDC-HSC’s as a benchmark.

A. Experiment Design

1) *Apparatus*: The experiment was performed at the Human-Machine Interface Lab, located within the faculty of Aerospace Engineering at the Delft University of Technology. The research and experiment was approved by the TU Delft Human Ethics Research Committee. A 36×29.5 cm display with a resolution of 1280×1024 pixels was used to show one

second of preview information, updated at a 100 Hz refresh rate. A right-handed electro-hydraulic servo-controlled side-stick was used to perform the task and provide haptic force-feedback, updated at 2,500 Hz. Only the roll axis of the stick had freedom of movement, its parameters are defined in Section II-A.

2) *Controllers*: Two independent variables are varied between conditions: (1) type of HSC, and (2) display settings. For the first, three types of HSC's are tested: a fixed FDC-HSC configured for full preview (F6), a fixed FDC-HSC configured for pursuit (F0), and the novel adaptive controller (A-HSC). In addition to the three HSCs, two non-haptic (NH) conditions are also performed for reference.

The settings for the fixed HSCs are given in Table II. For the A-HSC, a buffer time of $T_b = 7s$ was used, its update rate was $f_\Delta = 100 Hz$, its gain was $K_a = 0.002$, and $T_{l,f}$, K_{e^*} were dynamically adjusted according to Fig. 18. The remaining HCR parameters were set according to Table II. Finally, the initial pursuit and preview adaptive look-ahead times were 0.0s and 0.6s, respectively.

3) *Display*: In terms of display settings, each condition went from pursuit (PS, $\tau_p = 0s$) to preview (PR, $\tau_p = 1.0s$), or vice versa, at $t = 50s$. It is important to note that the HSC is *not* subject to this limitation, e.g., a HSC could have a look-ahead time of $\tau_{HSC} = 0.6s$ when the HO is in pursuit. For example, when using F6, the HSC will inevitably have a higher look-ahead time than the HO for approximately half of the run (during PS).

4) *Training*: Each participant was offered a set of minimum three training runs, starting without haptic feedback to get acquainted with the control task. The other two training conditions consisted of at least one run of each fixed-controller type. After a training run with haptic feedback, subjects were asked to fill in a Van der Laan questionnaire [21], so that they could establish a personal subjective baseline for evaluating HSCs during the subsequent conditions. Throughout training runs, the preview time was kept constant at $\tau_p = 1.0s$, hereby ensuring subjects experienced what it was like to conflict with F0 before measurements were taken.

5) *Conditions*: Each of the three HSCs was tested in two situations: PS-PR and PR-PS. In addition to two NH conditions (PS-PR + PR-PS), a total of eight experiment conditions were tested per subject. Each condition consisted of four runs lasting 128 seconds each, discarding the first 8 seconds of each run. The first run of each condition was discarded to mitigate learning effects, the remaining three runs used three different forcing function realizations.

6) *Method*: After each haptic condition, participants were asked to rate their subjective experience of the HSC by filling in a Van der Laan questionnaire. During non-haptic conditions, the same cross-correlation algorithm was used to calculate τ_a in the background, without applying any haptic feedback.

Three metrics were chosen to evaluate each condition: tracking error, conflicts, and control activity. Before transition at $t = 50s$, the root-mean-square (RMS) was calculated per run between $t = 18-48s$ for each metric. After transition the RMS was taken between $t = 98-128s$, to avoid measuring transient effects while the A-HSC is still adapting. A distinction is

made between preview before transition (PR1) and preview after transition (PR2), *mutatis mutandis* for pursuit.

To aggregate and compare pairwise results, first a Shapiro-Wilk normality test was performed on the given metrics. Thereafter, an F-test was applied to investigate the significance of pairwise variances. Depending on the result of the F-test, either a standard t-test was performed or Welch's t-test.

B. Hypotheses

The experiment aims to test six hypotheses:

H.I *Tracking performance and conflicts will be better for the A-HSC compared to F0.*

Previous results shown in Fig. 4 indicate that HSCs with $\tau_{HSC} < 0.5s$ cause conflict and reduce performance.

H.II *Tracking performance and conflicts will be similar for the A-HSC compared to F6.*

Because the F6 HSC will be able to react faster than the HO in pursuit, F6 is expected to score higher or equal to the adaptive HSC in terms of error and conflicts.

H.III *The adaptive look-ahead time will increase from 0.0s to 0.6 – 0.8s for the A-HSC (PS-PR).*

The adaptive HSC is designed to adjust itself according to the HO's look-ahead time, which is expected to increase.

H.IV *The adaptive look-ahead time will stay constant at 0.6 – 0.8s for the A-HSC (PR-PS).*

Results in Fig. 21 indicate that if the HO decides to trust the HO and comply with its inputs, then the adaptive HSC won't be influenced by changes in the state.

H.V *The estimated adaptive look-ahead time for NH (PS-PR) will increase from 0.0s to 0.6 – 0.8s throughout runs.*

In the non-haptic condition, the adaptive look-ahead time is expected to be an estimate of τ_{HO} .

H.VI *The estimated adaptive look-ahead time for NH (PR-PS) will decrease from 0.6 – 0.8s to 0.0s throughout runs.*

Because the HSC is no longer providing haptic feedback, the measured adaptive look-ahead time is expected to reduce relative to the A-HSC (PR-PS), reflecting τ_{HO} .

VII. EXPERIMENT RESULTS

This section lays out the results of the experiment: aggregate results, control activity, look-ahead adaptation, and subjective results.

A. Aggregate Results

Fig. 22 compares the distribution of the RMS for each of the three metrics before and after transition.

1) *Tracking Performance*: For each of the three HSCs, the mean error RMS_e is smaller than the corresponding mean error in the non-haptic case, shown by dashed blue lines in Fig. 22. The F6 HSC scores the highest in terms of tracking performance, with the error for F6 PS being approximately equal to the NH error in PR. F0 results in the worst tracking performance out of the three HSCs, although still providing a marginal improvement in error relative to NH.

For PR1-PS2, the A-HSC yields similar tracking performance to F6, whereas for PS1-PR2 the tracking performance

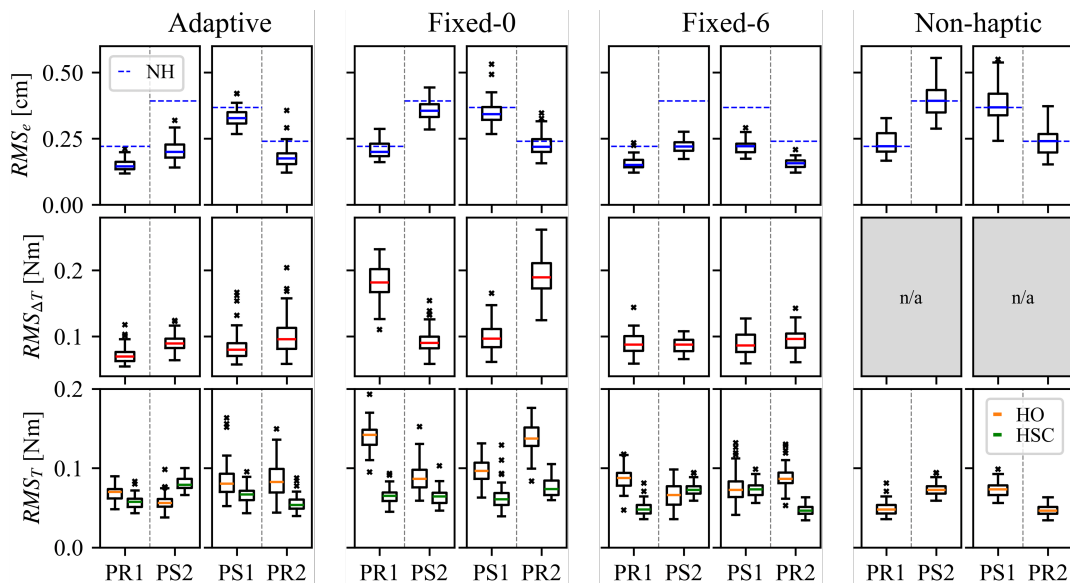


Fig. 22: Root-mean-square for tracking performance, haptic conflict, and control activity for pursuit and preview per condition

of the A-HSC is worse, and therefore closer to F0. This indicates a certain asymmetry: when the look-ahead time increases, the tracking performance of the A-HSC tends to more in line with F0, whereas when the look-ahead time decreases the A-HSC yields similar performance for F6.

2) *Conflicts*: Fig. 22 shows that F0 results in significantly more haptic conflict compared to the other two HSCs, hereby performing the worst out of the three. The A-HSC results in similar levels of conflict compared to F6, although the variance of conflict is higher for PR2. The minimum conflict recorded during the experiment occurred during PR1 for the A-HSC. In this case the median conflict RMS was $RMS_{\Delta T} = 0.068 Nm$ for the A-HSC, whereas for F6 it was $0.084 Nm$.

3) *Control activity*: The last row of Fig. 22 shows that HO control activity is highest for F0, which is a result of increased conflict. For all three HSCs, going from PR1 to PS2 consistently leads to a reduction in HO control activity. The opposite happens in the NH condition when going from PR1 to PS2, where HO control activity tends to *increase* when HOs perform the control task on their own. This indicates that HOs tend to become more passive going from PR to PS, corresponding with increased HSC control activity between PR1 and PS2 for the A-HSC and F6.

Going from PS1 to PR2, on average, HO activity either stays constant (A-HSC) or increases (F0, F6), whereas in the NH condition HO control activity reduces. This is approximately the opposite of what occurred for PR1-PS2, which is due to the fact that the display settings have simply been switched (pursuit to preview vs. preview to pursuit).

4) *Statistical Tests*: Table III shows the results of the statistical tests corresponding with Fig. 22, comparing the A-HSC with the other two HSCs. Less significant differences were found between A and F6 than were found when comparing A F0. For A-F0, 11/16 highly significant differences were found across all metrics, whereas for A-F6 only 7/16 highly significant differences were found. This corresponds to the fact

that, by and large, A and F6 show more similarities in Fig. 22 than A and F0 do. This is primarily because of the significant amount of conflict caused by F0.

Furthermore, Table III shows that A and F0 are least different during PS1, relative to the other three cases. In fact, PS1 is the only case for which the differences between A-F6 are more significant than they are for A-F0 as shown in Table III. This is because both the A-HSC and F0 happen to be using a similar look-ahead time for PS1, as is discussed in Section VII-C.

TABLE III: Statistical test results for A-F0, A-F6

* indicates significant ($p < 0.05$),
** indicates highly significant ($p < 0.001$)

	A-F0				A-F6			
	PR1	PS2	PS1	PR2	PR1	PS2	PS1	PR2
RMS_{THO}	** , W	** , W	* , W	** , t	** , W	* , W	* , t	- , W
RMS_{THSC}	** , t	** , W	- ,	** , t	** , t	** , t	** , t	** , W
$RMS_{\Delta T}$	** , W	- , W	* , t	** , W	* , t	* , t	- , W	- , W
RMS_e	** , W	** , t	* , t	** , t	* , t	* , W	** , W	** , W

B. Control Activity

Fig. 23 shows how control activity, haptic conflict, and tracking performance evolve throughout haptic conditions. The moving average (MA) of the RMS was calculated for each metric using a 10-second sliding window, and then averaged across all runs. The use of a MA allows underlying trends to be revealed, which are less dependent on high-frequency content of the target signal.

1) *Tracking Performance*: Fig. 23 shows that in terms of error during PR-PS, the A-HSC scores similarly to F6, corresponding to results shown in Fig. 22. For PS-PR the tracking error for the A-HSC is in line with F0, which is due to the initial settings of the HSC ($\tau_a^{HSC} = 0$). This can also be seen in Fig. 22. During PR2 the tracking error reduces for all three HSCs. However, although tracking performance for

the A-HSC is initially similar to F0, after transition it tends more towards F6. This is a consequence of adaptation taking place in the A-HSC.




Fig. 23: Moving averages (MA) of root-mean-square for tracking performance, haptic conflict, and HO control activity

2) *Conflicts*: For PR-PS, Fig. 23 shows that the A-HSC performs slightly better than F6 in terms of conflicts during PR1. After transition to PS2, all three HSCs result in approximately the same amount of conflict. This is a remarkable finding, given that tracking performance is significantly different between F0 and F6, and at least two of the HSCs (F0 and F6) are using different look-ahead times. This result provides further evidence that HOs adapt their compliance to HSC torques depending on if the HSC is positively contributing to tracking performance.

For PS-PR a brief period of conflict occurs after transition for the A-HSC, tending towards the baseline level of conflict set by F6. Out of all three HSCs, F6 is the only controller for which the amount of HO-HSC conflict remained invariant throughout the experiment conditions.

3) *HO Control Activity*: Fig. 23 shows that the HO was least active for the A-HSC during PR-PS. Furthermore, for all three HSCs, HO control activity reduced after transitioning from PR to PS. This corresponds to an earlier finding as shown in Fig. 22, and discussed in Section VII-A, suggesting that HOs become passive when going from PR to PS under the influence of haptic feedback. For PS to PR, control activity of the A-HSC increases due to the conflict, but then tends towards F6 as adaptation takes place.

4) *Adaptive HSC Statistical Tests*: Table IV shows the statistical test results pertaining only to the A-HSC. Two types of pairwise tests are performed: 1) within-run pairwise comparison (e.g., PS-PR), and 2) between-run pairwise comparison (e.g., PS1 with PS2).

For the within-run tests, the only pairwise comparison that had no significance was $RMS_{T_{HO}}$, going from PS to PR.

A hypothesis for this is that since the A-HSC consistently aligns itself with the HO's behaviour, the HO's contributions remained approximately constant. For PR-PS a significant effect was found for $RMS_{T_{HO}}$, with Fig. 22 showing a reduction in HO control activity. This is due to the A-HSC operating with a higher look-ahead time than the HO, as described in Section VII-C.

TABLE IV: Statistical test results for A-HSC

* indicates significant ($p < 0.05$),
** indicates highly significant ($p < 0.001$)

	PR-PS	PS-PR	PS1-PS2	PR1-PR2
$RMS_{T_{HO}}$	** , t	- , t	** , W	** , W
$RMS_{T_{HSC}}$	** , t	** , t	** , t	- , t
$RMS_{\Delta T}$	** , t	* , t	- , W	** , W
RMS_e	** , W	** , t	** , t	** , t

For the between-run comparison, no significant difference was found between the pursuit cases for $RMS_{\Delta T}$, even though it is known that in PS2 the HO tends to reduce their control activity. This simply means that the level of conflict for a complying HO in PS2 ($\tau_a^{HSC} > 0.6s$) happens to be similar to the level of conflict in PS1 ($\tau_a^{HSC} = 0s$). Furthermore, for PR1-PR2, no significant effect is observed for T_{HSC} , indicating that the level of control activity during PR is approximately the same regardless of the order of pursuit/preview.

C. Look-Ahead Adaptation

Fig. 24 shows the adaptive look-ahead time τ_a^{HSC} for the A-HSC and NH conditions, PS to PR. With the A-HSC active in-the-loop, the mean look-ahead time across all subjects increases from $\tau_a = 0.08s$ just before transition, to $\tau_a = 0.70$ by the end of the condition. The maximum variance occurs at $t = 80s$ as predicted by simulations, the 25th and 75th percentile equal $q_{25}, q_{75} = 0.25 - 0.58s$, respectively. At the end of the condition $q_{25}, q_{75} = 0.660 - 0.780s$.

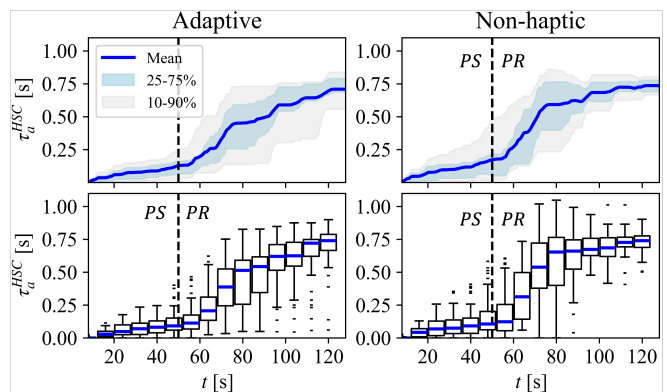


Fig. 24: Experiment adaptive look-ahead time (PS-PR)

For NH the mean just before transition equals $\tau_a = 0.09s$, increasing to $0.73s$ by the end of the condition. At $t = 80s$, the mean look-ahead time is $\tau_a = 0.64s$ and $q_{25}, q_{75} = 0.39 - 0.76s$. By the end of the condition the 25th and 27th percentile become $q_{25}, q_{75} = 0.69 - 0.77s$.

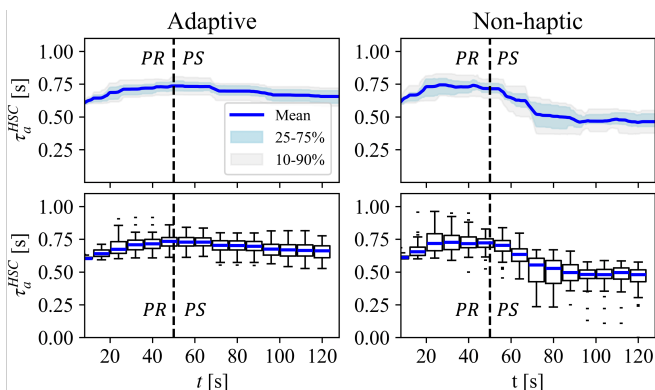


Fig. 25: Experiment adaptive look-ahead time (PR-PS)

Fig. 25 shows the adaptive look-ahead time for the A-HSC and NH conditions, now PR to PS. For the A-HSC the mean look-ahead time marginally decreases from $\tau_a = 0.72s$ just before transition, to $\tau_a = 0.65s$ by $t = 128s$. The variance is approximately equal throughout the condition, with $q_{75} - q_{25} = 0.09s$. For NH the mean just before transition equals $\tau_a = 0.71s$, decreasing to $0.47s$ by $t = 128s$. At the end of the run $q_{25}, q_{75} = 0.42 - 0.52s$.

D. Subjective Results

Fig. 26 shows the subjective results gathered from the Van der Laan questionnaires. It shows that the worst performing HSC in terms of subjective experience is the fixed HSC with pursuit settings. The results of the corresponding statistical tests are listed in Table V, showing that the F0 does indeed score significantly worse than the A-HSC. Comparing A-F6, no significant difference was found between the subjective ratings. These findings align with results summarised in Table III, and shown in Fig. 22, where A-F6 are found to be more similar than A-F0. Finally, no significant effect was observed between the two conditions for the A-HSC (PS-PR vs. PR-PS).




Fig. 26: Subjective results

TABLE V: Statistical test results for subjective ratings

* indicates significant ($p < 0.05$), t indicates a normal t -test, W indicates Welch's t -test, note: 6/16 forms were lost

	A-A		A-F0		A-F6	
	PS-PR vs PR-PS	PS-PR PR-PS	PS-PR PR-PS	PS-PR PR-PS	PS-PR PR-PS	PS-PR PR-PS
Usefulness	-,W	*,t	*,W	-,t	-,t	-,t
Satisfaction	-,W	*,t	*,W	-,t	-,t	-,t

VIII. DISCUSSION

The goal of this research was to design and implement an adaptive haptic shared controller (A-HSC) capable of adjusting its look-ahead time according to time-variant human behaviour. To this end, the A-HSC was tested alongside traditional fixed HSCs with constant parameter settings.

1) *Hypotheses*: H.I states that ‘tracking performance and conflicts will improve when comparing the A-HSC to F0’. Table III shows that in 3/4 cases conflict was found to have significantly reduced, and that in 4/4 cases error significantly reduced. Therefore, H.I is accepted.

H.II posits that ‘tracking performance and conflicts will be similar for the A-HSC compared to F6’. In 2/4 cases no significant difference was found in conflict between the two controllers (A-F6). In terms of average error, all 4/4 cases were found to be significantly different. In 2/4 cases the average error was smaller for the A-HSC. Since the overall differences in error and conflict are small relative to A-F0, H.II is accepted.

H.III claims that ‘the adaptive look-ahead time will increase from 0s to 0.6 – 0.8s for the A-HSC (PS-PR)’. On average the adaptive look-ahead increased from 0.08s to 0.70s with confidence intervals of $q_{10}, q_{90} = 0.02 - 0.22s$ and $q_{10}, q_{90} = 0.55 - 0.84s$, respectively. Therefore, H.III is accepted.

H.IV claims that ‘the adaptive look-ahead time will stay constant at 0.6 – 0.8s for the A-HSC (PR-PS)’. Results show that during the experiment the look-ahead time remained bounded between $q_{10}, q_{90} = 0.60 - 0.77s$ for the A-HSC (PR-PS), therefore, H.IV is accepted.

H.V expects that ‘the estimated adaptive look-ahead time for NH (PS-PR) will increase from 0s to 0.6 – 0.8s throughout runs’. Fig. 24 shows that for NH (PS-PR), τ_a increased from 0.09s to 0.73s with confidence intervals of $q_{10}, q_{20} = 0.01 - 0.42s$ and $q_{10}, q_{20} = 0.66 - 0.80s$, respectively. Therefore, H.V is accepted.

Finally, H.VI states that ‘the estimated adaptive look-ahead time for NH (PR-PS) will decrease from 0.6 – 0.8s to 0s throughout runs’. Although a decrease in adaptive look-ahead time was measured from $\tau_a = 0.71s$ to $0.47s$, a larger difference was expected, such that the HO’s actual look-ahead time (0s in PS) would be reflected by the end of the run. Therefore, H.VI is rejected. This effect is likely due to the minimum threshold that was set for K_{xx} , as discussed in Section IV-A.

2) *Trust and Acceptance*: The HO became significantly more passive going from preview to pursuit than vice versa, as shown in Fig. 23. For all three HSCs, HO control activity reduced going from PR to PS. Some participants reported that, although the F6 HSC was comfortable and trustworthy, it also led to complacency and inattentiveness in pursuit situations. This is likely due to the fact that the HSC was operating with more information than was available to the HO. Furthermore, it was observed that under no circumstance was there sustained conflict with the F6 HSC.

Despite the A-HSC having similar subjective ratings as F6, Fig. 23 shows a period of brief haptic conflict after transition for PS-PR. Furthermore, the error of the A-HSC during PS1

was significantly larger ($RMS_e \approx 0.33cm$) relative to F6 ($RMS_e \approx 0.20cm$). Remarkably, despite experiencing these drawbacks, on average subjects rated the adaptive and F6 HSCs similarly. This indicates that by successfully adapting to the HO intentions, some of the A-HSC's drawbacks are negated and overall acceptance is still possible.

3) *Recommendations*: As a result of this study, new opportunities emerge for the development of adaptive haptic shared controllers. The first step would be to explore possible avenues for optimising the design laid out in this paper. Since this study endeavoured to test an entirely novel design, the A-HSC's settings were set to be conservative in terms of adaptation speed, favouring robustness instead. It is therefore likely that the buffer size could further be reduced without sacrificing stability, which would minimise delays related to the use of a sliding window. Furthermore, the adaptation speed could be improved by increasing K_a , in addition to removing the minimum threshold for K_{x_r} : this protection was found to be redundant, except under highly exceptional circumstances.

Looking further ahead, the next steps for this research are to expand the application of the proposed novel A-HSC towards control tasks that more closely resemble car driving. In order to achieve this, the same steps outlined by Van der El [16] can be used to move away from pure preview tracking, and closer to real-life car driving. This could include the incorporation of 1) a linear perspective, 2) motion feedback, 3) multiple visual feedback cues (e.g., sky and winding road), and 4) boundary avoidance [22–25]. Additionally, the effect of using second order controlled element dynamics should be investigated, as these dynamics will more closely resemble car driving. These steps could pave the way for the use of adaptive HSCs in real-life car driving, hereby improving both the safety and ease-of-use of advanced driver assistance systems.

As the experiment conditions are refined to more closely reflect real car-driving, HO behaviour is expected to become more complex. This will further complicate the design of an adaptive controller, since it is likely that only adapting a single parameter in an otherwise fixed controller structure will not suffice. Despite these challenges, the results of this paper show that, in principle, reliable methods exist for adapting look-ahead time. Given the importance of look-ahead time and its impact on HO control behaviour, the results of this paper can be considered a successful step towards adaptive HSCs for car-driving.

IX. CONCLUSION

In this paper, a novel design for an adaptive haptic shared controller (A-HSC) is designed and implemented. A human-in-the-loop tracking task was performed, during which the display's preview time transitioned from pursuit to preview, or vice versa, at 50 seconds after the start of each run. Going from pursuit to preview, the look-ahead time of the A-HSC increased from 0.08 to 0.70s, resulting in a reduction of haptic conflicts and an increase in tracking performance. For preview to pursuit, subjects' control activity reduced, and the look-ahead time of the A-HSC stayed approximately constant between 0.6–0.8s. Because the A-HSC had a higher

look-ahead time than the subjects at the point of transition, participants reported complying with the HSCs inputs, hereby preventing further adaptation. In terms of subjective rating, the A-HSC scored equally well compared to a fixed preview HSC, despite a brief period of conflict during adaptation, as well as lower tracking performance during pursuit-to-preview. This paper demonstrates the implementation of an adaptive haptic shared controller that adjusts its look-ahead time according to the human it shares control with. The adaptive look-ahead algorithm can be used to determine an online estimate of human look-ahead time during preview tracking tasks, with or without haptic support.

REFERENCES

- [1] E. Aria, J. Olstam, and C. Schwietering. "Investigation of Automated Vehicle Effects on Driver's Behavior and Traffic Performance". In: *Transportation Research Procedia* 15 (2016), pp. 761–770.
- [2] M. Körber, E. Baseler, and K. Bengler. "Introduction matters: Manipulating trust in automation and reliance in automated driving". In: *Applied Ergonomics* 66 (2018), pp. 18–31.
- [3] D. de Waard, M. van der Hulst, M. Hoedemaeker, and K. A. Brookhuis. "Driver behavior in an emergency situation in the Automated Highway System". In: *Transportation human factors* 1.1 (1999), pp. 67–82.
- [4] L. Bainbridge. "Ironies of automation". In: *Automatica* 19.6 (1983), pp. 775–779.
- [5] D.A. Abbink, M. Mulder, and E.R. Boer. "Haptic shared control: Smoothly Shifting Control Authority?" In: *Cognition, Technology and Work* 12.1 (2011), pp. 1–10.
- [6] T. J. Gordon and M. Lidberg. "Automated driving and autonomous functions on road vehicles". In: *Vehicle System Dynamics* 53.7 (2015), pp. 958–994.
- [7] R. Kondo, T. Wada, and K. Sonoda. "Use of Haptic Shared Control in Highly Automated Driving Systems". In: vol. 52. 19. 2019, pp. 43–48.
- [8] Boink, M. M. Van Paassen, M. Mulder, and D. A. Abbink. "Understanding and reducing conflicts between driver and haptic shared control". In: *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (2014), pp. 1510–1515.
- [9] D. A. Abbink and M. Mulder. "Neuromuscular Analysis as a Guideline in designing Shared Control". In: *Advances in Haptics*. Intech, 2010, pp. 499–516.
- [10] J. Span, D. M. Pool, M. M. van Paassen, and M. Mulder. *Effects of Look-Ahead Time in a Haptic Shared Controller for Preview Tracking*. Unpublished MSc thesis. 2021.
- [11] K. van der El, D. M. Pool, H. J. Damveld, M. M. van Paassen, and M. Mulder. "An Empirical Human Controller Model for Preview Tracking Tasks". In: *IEEE Transactions on Cybernetics* 46.11 (2016), pp. 2609–2621.

- [12] S. Wietske, S. Barendswaard, D.M. Pool, M. M. Van Paassen, and D.A. Abbink. "A New Haptic Shared Controller Reducing Steering Conflicts". In: *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2018, pp. 2705–2710.
- [13] M. Itoh, F. Flemisch, and D.A. Abbink. "A hierarchical framework to analyze shared control conflicts between human and machine". In: *13th IFAC Symposium on Analysis, Design, and Evaluation of Human-Machine Systems HMS 2016*. Vol. 49. 19. Kyoto, Japan, 2016, pp. 96–101.
- [14] M. M. van Paassen, R. P. Boink, D. A. Abbink, M. Mulder, and M. Mulder. "Haptic guidance, interaction between the guidance model and tuning". In: *18th International Symposium on Aviation Psychology, ISAP2015*. Dayton, United States, May 2015, pp. 410–415.
- [15] J. C. F. de Winter and D. Dodou. "Preparing drivers for dangerous situations: A critical reflection on continuous shared control". In: *2011 IEEE International Conference on Systems, Man, and Cybernetics*. Anchorage, AK, United States, Oct. 2011, pp. 1050–1056.
- [16] K. van der El, D.M. Pool, and M. Mulder. "Measuring and Modeling Driver Steering Behavior: From Compensatory Tracking to Curve Driving". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 61 (2019), pp. 337–346.
- [17] K. van der El, S. Padmos, D. Pool, M. M. van Paassen, and M. Mulder. "Effects of Preview Time in Manual Tracking Tasks". In: *IEEE Transactions on Human-Machine Systems* 48.5 (2018), pp. 486–495.
- [18] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder. "Effects of Target Trajectory Bandwidth on Manual Control Behavior in Pursuit and Preview Tracking". In: *IEEE Transactions on Human-Machine Systems* 50.1 (2020), pp. 68–78.
- [19] D. T. McRuer and H. R. Jex. *A review of quasi-linear pilot models*. Vol. HFE-8. 3. 1967, pp. 231–249.
- [20] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder. "Effects of Preview on Human Control Behavior in Tracking Tasks With Various Controlled Elements". In: *IEEE Transactions on Cybernetics* 48.4 (2018), pp. 1242–1252.
- [21] J. D. Van der Laan, A. Heino, and D. De Waard. "A simple procedure for the assessment of acceptance of advanced transport telematics". In: *Transportation Research Part C: Emerging Technologies* 5.1 (1997), pp. 1–10.
- [22] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder. "Effects of Linear Perspective on Human Use of Preview in Manual Control". In: *IEEE Transactions on Human-Machine Systems* 48.5 (2018), pp. 496–508.
- [23] K. van der El, J. Morais Almeida, D. M. Pool, M. M. van Paassen, and Max M. "The Effects of Motion Feedback in Manual Preview Tracking Tasks". In: *AIAA Modeling and Simulation Technologies Conference*.
- [24] K. van der El, D. M. Pool, M. M. Van Paassen, and M. Mulder. "Identification and Modeling of Driver Multiloop Feedback and Preview Steering Control". In: *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*. Myazaki, Japan, Oct. 2018, pp. 1227–1232.
- [25] K. van der El, D. M. Pool, M. M. van Paassen, and Max M. "Modeling driver steering behavior in restricted-preview boundary-avoidance tasks". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 94 (2023), pp. 362–378.

Part II

Preliminary Thesis Report

*This part has been assessed for the course AE4020 Literature Study.

Summary

With the advent of self-driving cars, new types of human-machine interfaces are emerging to support drivers in terms of performance, safety, and workload [1, 2]. Even though it may be attractive to replace humans in tasks such as driving vehicles, hidden dangers such as over-reliance and loss of situation awareness prove that full automation comes at a cost [3]. A promising solution to these challenges are haptic shared controllers (HSCs), a form of automation that offers to combine the benefits of machines with the strengths of human operators (HOs), whilst simultaneously offsetting risks related to complacency and over-reliance [4, 5].

Work by Span proved that by configuring a haptic shared controller to mimic human behaviour through a cybernetic model, the risk of conflicts is reduced and tracking performance increases [6]. However, in reality the real strength of humans lies in their versatility in *adapting* to changing circumstances such as weather, which is often the reason why they are kept in the control loop [7]. In such conditions a time-invariant HSC with a fixed structure may become obsolete, since it cannot adapt to the varying behaviour of the HO who it is trying to assist.

This report aims to propose a structure for an adaptive HSC which is capable of dynamically adjusting its settings in real-time to reduce conflicts between a HO. To this end, three novel controllers are proposed, each designed to adjust the look-ahead time of a Four-Design-Choices (FDC) HSC with the goal of minimising conflict torques between the HO and HSC. Simulations show that each of the three adaptive controllers is capable of adjusting its look-ahead time as a consequence of adapting to time-variant human behaviour.

Furthermore, system identification is performed on previously gathered experiment data to evaluate if individual differences can be captured, potentially increasing alignment between controller behaviour and the HO. It was found that individual HO differences are not significant enough to result in increased performance of a FDC-HSC by using subjects' parameter estimation results.

To expand the applicability of existing models, a structure for HSC torque rejection is proposed and tested against experimental data. The novel architecture superimposes a torque-rejection loop to model HO suppression of HSC torques through force feedback. The rejection model was used in the design process of the adaptive HSC to more accurately simulate HO behaviour, particularly for extreme values of HSC look-ahead time.

The three adaptive haptic shared controllers (A-HSCs) are designed and tested in a simulated, time-varying HO scenario. During runs of 200 seconds the HO's look-ahead time is discontinuously changed from $0.6s$ to $0.1s$ at $t = 38s$, triggering adaptation of the controllers.

- (1) The inertial A-HSC uses a memory buffer (chunks) and decision policy to determine if adaptation steps yield a reduction in HO-HSC conflict. If subsequent chunks indicate that conflict is reducing, then the controller continues to adapt in the same direction. Conversely, if conflict is increasing, then the controller reverses adaptation direction.
- (2) The model-fitting A-HSC employs multiple HSCs running in the background. Every 5 seconds a selector decides which HSC is the prime candidate based on its proposed torques. This candidate is then selected to be in-the-loop together with the HO.
- (3) The cross-correlation A-HSC is designed to minimise the phase difference between the FDC-HSC's reference state, x_r , and the actual state, x . In doing so, the controller reduces the mismatch between the HSC's internal model and reality, hereby resolving the root cause of HSC-HO conflicts.

The cross-correlation A-HSC shows the most promising results based on adaptation speed, adaptation range, and robustness against disturbances. Simulation results averaged over 100 realizations show that the controller is capable of adapting from $\tau_a = 0.78s$ to $\tau_a = 0.29s$ during a period of $16s$. After $199s$, on average the controller adapts its look-ahead time from $\tau_a = 0.76s$ to $\tau_a = 0.30s$. The average standard deviation throughout simulation runs is $\sigma = 0.07s$.

Introduction

As society becomes increasingly dependent on automated machines due to their economic and practical benefits, there exists a strong impulse to replace human beings with autonomous systems. However, even once a task or process has been automated humans are nonetheless often required to be present in a supervisory role for redundancy purposes [7]. The transition from active participant to observer may cause the user to adopt a passive mindset, become over-reliant on the system, or lose situation awareness [3, 8]. Although these issues may not reveal themselves under normal circumstances, they frequently prove to be crucial factors in the case of accidents resulting from human-machine error [9, 10]. Therefore, there is a strong case to be made that as the complexity of automation increases, the more important the design of its human interface becomes [11].

An alternative to completely automating a given task is *shared control*, which proposes that a human operator (HO) is assisted by automation in such a way that control authority is shared between the system and the user. In the context of steering vehicles, shared control entails that the human is required to *actively participate* in driving the vehicle whilst being assisted by a guidance system.

By ensuring that the operator is partially responsible for controlling the vehicle, risks related to over-reliance and complacency are offset in a way that retains the benefits of automation [8]. The difficulty in shared control lies in configuring the guidance system such that the HO trusts and accepts its inputs [12], which implies that the driver has a minimum level of understanding about the system's intent. One method to achieve this is through haptic force-feedback.

The objective of this report is to propose a novel *adaptive haptic shared controller (HSC)* that adjusts its settings in real-time to reduce conflicts between a HO in a preview tracking task. To this end, three novel haptic shared controllers were designed and evaluated. Simulations investigated each controllers' ability to adapt to time-varying human behaviour.

Through its adaptive structure, the proposed architecture promises to accommodate time-varying human behaviour. This could be beneficial in developing Advanced Driver-Assistance Systems (ADAS), particularly in the case of reduced visibility when the driver's control behaviour is expected to change due to the now-limited preview information [13].

The report is presented in the following structure. The remainder of Chapter 1 describes in detail the motivation for adaptive haptic shared control, as well as the research objective of the project and its outline. Thereafter, Chapter 2 summarizes the relevant literature concerning human behaviour in a preview tracking task, the Four Design Choices (FDC) Architecture, and the effects of preview time on conflicts for an FDC-HSC. Chapter 3 details the design and testing of the novel adaptive controller, as well as an evaluation of system identification results and proposal of an HO HSC-rejection model. Finally, conclusions are laid out in Chapter 4.

1.1. Motivation For Adaptive Haptic Shared Control

In *haptic shared control* the guidance system provides force inputs to the same control device that the HO uses, hereby allowing the driver to sense the system's intentions by means of touch or force feedback. Because both the guidance system and operator influence the position of the same physical object, control authority is negotiated by the operator's continuous acceptance or rejection of the guidance system's forces. For example, if the operator senses that the system's forces are *not* in line with his/her own expected control inputs, the operator can override the automation by counteracting its forces.

Neuromuscular feedback in humans is significantly faster than visual cues [14], making haptic feedback a highly effective method of interaction and communication between guidance systems and HOs. Furthermore, since the operator is required to be in direct physical contact with the control device to apply steering forces, he/she is always optimally positioned to correct or reject the guidance system. Compared with traditional on-off automation this allows users to respond more swiftly to unexpected situations (even before situation awareness is considered).

For example, in the case of an unexpected event in an aircraft flying on autopilot, the pilot would typically first be required to place his/her hand on the disengage switch to disable the automation. Thereafter control is abruptly surrendered by the system and it is left to the operator to rapidly transition from passive observer to active controller, a challenge that is compounded by potential startle & surprise effects of the original unexpected event.

Since it is largely unknown how humans adapt to haptic guidance systems [15], it is difficult to predict the likelihood that a given controller will be accepted by the user. A promising method is to use a cybernetic model to simulate the HO, hereby configuring the HSC to mimic the HO's behaviour.

By simulating novel HSC designs and observing the effects of controller settings on error and control inputs, a prediction can be made as to which settings will minimize human-machine conflict during experiments. Furthermore, insights can be gained as to which configurations will have the highest chance of user acceptance. The ultimate goal is to converge towards the optimal settings in real-time such that performance and conflicts are optimised.

1.2. Research Objective

Previous research by Span [6] has proven that the driver's workload can be reduced by configuring the guidance system to mimic the human's control strategy. Through the use of cybernetics a so-called human compatible reference (HCR) can be implemented in the guidance system, which allows it to model and replicate the user's control behaviour [16]. In the case of a preview display, the human's response is characterised by a 'corner-cutting' approach and the use of a look-ahead time, whereby the user aims to track a point on the signal τ seconds into the future [17].

Span showed that if the look-ahead time of the HCR deviates significantly from the human's actual control strategy in a preview tracking task, then the duration and magnitude of conflicts between the automation and user increase [6]. These conflicts defeat the entire purpose of the guidance system, which is to offer collaborative support to the human.

Although a region of optimal look-ahead times has been identified in which conflicts between the guidance system and human are reduced, the root causes of human-haptic conflicts is yet to be investigated. Moreover, most of the parameters of current HSCs are fixed and tuned heuristically (e.g. level of haptic authority), meaning that such controllers are not capable of accommodating time-varying human behaviour. This is likely because the effects of changing parameters within the haptic controller are currently not sufficiently understood, possibly made worse by the knowledge gap concerning the user's adaptation to haptic forces.

In conclusion, current HSCs are to an extent limited by their fixed structure and parameters, rendering them unable to adapt to time-varying human behaviour and dynamic environments. The most pertinent case study can be found in driver-assist systems in cars, which are currently limited in their ability to accommodate time-varying human behaviour.

In light of these challenges, the primary focus of this project is therefore given by the following research objective:

Research Objective

To design, simulate, and validate an adaptive haptic shared controller that adjusts itself in real-time to reduce conflicts between a human operator in a preview tracking task, without undermining control performance.

An important first step is to define under which circumstances the controller should adapt itself, and to what end. A common assumption is that conflicting control inputs between the human and automation are undesirable and should be reduced through the design or tuning of the controller. However, other indicators may also be suitable to drive adaption. Therefore, the first research question is:

Research Question 1

Which metrics should be used to trigger or drive adaptation of the haptic shared controller?

Once metrics have been identified and selected to drive changes within the controller, the next step is to further clarify exactly what those changes are and how they should relate to the adaptation drivers. For example, if conflict torques occur increase, to what extent should the adaptive haptic shared controller adapt itself and how? This is the focus of the following research question:

Research Question 2

How should the system adapt itself in relationship to the chosen metrics?

Finally, the ultimate goal of the project is to prove that the proposed adaptive controller is successful in reducing conflicts in a real-life setting. This implies the design and execution of a human-in-the-loop experiment, leading to the final research question:

Research Question 3

To what extent does the adaptive shared controller lead to a reduction in conflict torques during human-in-the-loop experiments?

1.3. Project Outline

Figure 1.1 illustrates the general outline of the project. During the *literature study* phase the focus lies on understanding the state of the art of the following three areas of research, as described in Chapter 2:

1. **Preview Cybernetics:** understanding how humans use preview information in a manual control task is key to successfully designing a haptic controller that mimics the user's behaviour in such a task.
2. **Four Design Choices Architecture (FDCA):** this project will make use of an FDCA haptic controller, which has proven to reduce conflict torques by segregating two key components of haptic guidance: a feedforward and feedback component.
3. **Effects of Look-Ahead Time:** the known effects of look-ahead time on conflict torques can be compared with the results of the novel adaptive controller for validation.

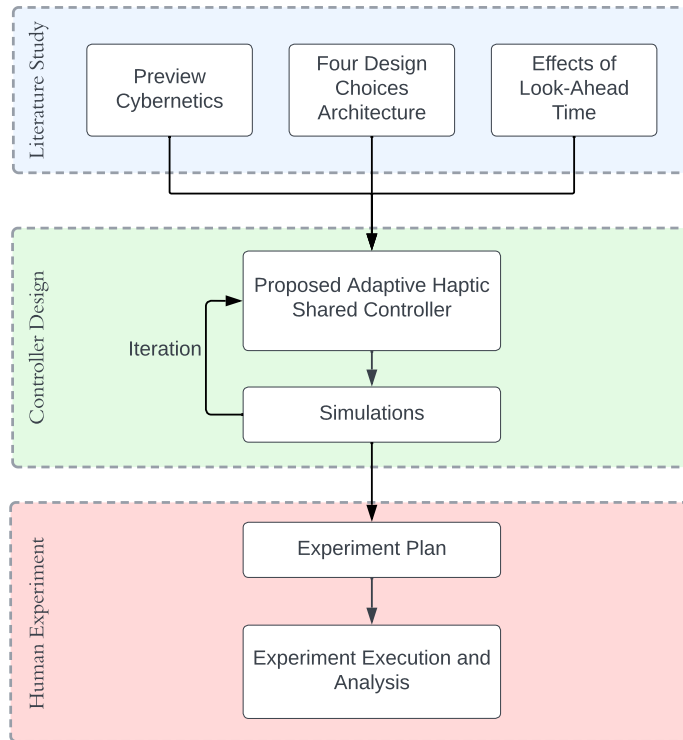


Figure 1.1: Thesis outline

In the *controller design* phase a novel adaptive controller is proposed by leveraging simulations to conduct an iterative design process, of which the results are laid out in Chapter 3. During this process previous work by Span [6] is expanded upon by proposing a model to capture human rejection of guidance forces. Furthermore, Span's system identification results are evaluated to investigate if they could lead to improvements within the shared controller. The final step after this process is to develop an experiment plan to validate and test the novel controller in a human-in-the-loop experiment.

2

Literature Review

This chapter aims to recapitulate the core literature relevant to the project, starting with a summary of Van der El's work on human adaptation to preview information in Section 2.1. Thereafter a brief recap of the Four Design Choices Architecture is given in Section 2.2, followed by a description of an FDC-HSC in Section 2.3. Finally, effects of varying HSC look-ahead times are summarised in Section 2.4.

2.1. Preview Cybernetics

2.1.1. Motivation to Study Preview Control

There are three primary motivations to implement haptic shared control in a *preview* tracking task as done in a previous study by Span. Firstly, it is known that human beings are adapted to use preview information when available to increase control performance in everyday situations [18]. The most prominent example of this is car-driving, where human beings heavily use the visual cues of the road ahead to optimize their control strategy. Designing a shared control system specifically for preview situations will therefore likely be more widely applicable in the real world than e.g. compensatory control as studied by McRuer & Jex in 1967 [19].

Secondly, implementing haptic shared control for a preview control task allows for the use and validation of novel human-operator models by Van der El, building on existing cybernetic foundations that were previously limited to compensatory control [7, 19]. By using these models to design controllers that are better adapted to the user's behaviour, the power and validity of cybernetic models for preview information can be demonstrated.

Finally, a preview display offers the driver information about the *future* of the target signal and therefore affords the driver the possibility of anticipating the upcoming route, something that is not possible for a compensatory tracking task. It is possible that the ability to anticipate future control inputs plays a key role in deciding when to *reject* the guidance system's inputs, making preview tracking a compelling case study for haptic shared control.

For example, if the upcoming route clearly shows a left turn and the guidance system is steering to the right, the driver will most likely counteract the system. By comparing his/her own (planned) intentions with the guidance system, the driver can decide to reject the system's inputs by impeding the movement of the control device. In this way, the user can dynamically determine the level of control authority the guidance system has by the degree to which its inputs are admitted.

2.1.2. Van der El's Cybernetic Model for Preview Control

Up until recently the study of manual cybernetics was limited to compensatory control, a form of steering that at first glance bears little resemblance to common real world tracking tasks such as driving a car or flying an aircraft. Pioneering work by Van der El has proven successful in expanding the field of manual cybernetics to accommodate human behaviour in a preview tracking task, opening up a vast landscape of opportunities for modelling human behaviour.

Through a black-box system identification method analogous to McRuer's landmark experiment, a cybernetic model for preview information was derived as shown in Figure 2.1. This model identified two channels, the first one describing the human controller's (HC) response to a so-called 'near-viewpoint' some τ_n

seconds along the previewed target signal into the future, $f_t(t + \tau_n)$. The second channel describes the HC's response to a far-viewpoint further ahead, $\tau_f (> \tau_n)$ seconds along the previewed signal.

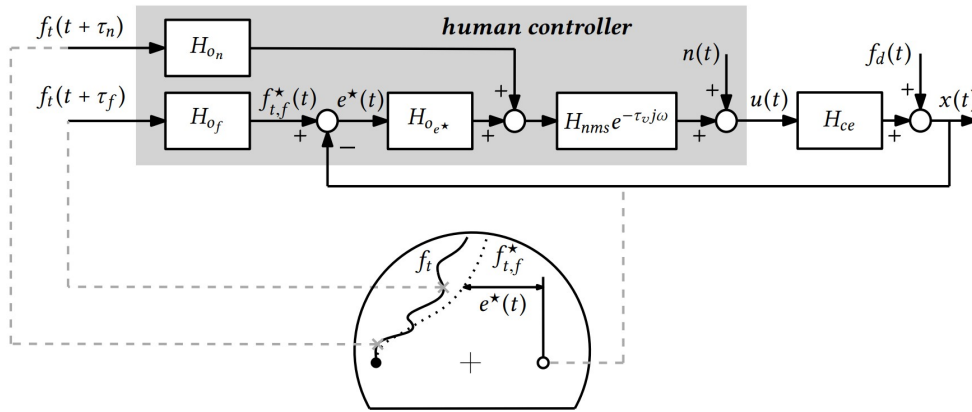


Figure 2.1: Intuitive form of Van der El's preview model [17]

The *near-viewpoint response* is characterised by a feedforward open-loop control strategy, whereby the human aims to track high-frequency components of the signal. Van der El posits that this response is made possible by the fact that full periods of the high-frequency components are visible on the display, allowing the HO to recognise cyclic patterns. This behaviour is modelled by a high-pass filter $H_{o_n}(j\omega)$:

$$H_{o_n} = K_n \frac{j\omega}{1 + T_{l,n}(j\omega)} \quad (2.1)$$

For the purposes of this project the near-viewpoint response will not be considered, since it is likely a special case of human adaptation which only occurs in a laboratory setting. In most real-world control tasks the chances of a pure sine wave occurring in the target signal are extremely small (consider the geometry of the average road), hence the motivation to discard this channel and focus only on the relatively strong far-viewpoint response.

The *far-viewpoint response* involves a feedback control loop that minimises the error between the current state and a low-pass filtered point on the signal τ_f seconds ahead into the future. The way that humans use the far-viewpoint can be characterised by two phenomena:

1. **Corner cutting approach:** high frequency components of the signal are neglected, similar to the behaviour described by McRuer's crossover model. This *corner cutting* is captured in the model by low-pass filtering a point on the signal τ_f seconds ahead into the future.
2. **Look-ahead anticipation:** by minimising the error between the current state and a point on the target τ_f seconds *into the future*, the HO effectively cancels out the inherent delays resulting from his/her physical limitations.

These two strategies are captured by the feedback loop illustrated in Figure 2.1 [17]. This channel minimises the error between the current state and a low-pass filtered point on the signal (corner-cutting) situated τ_f seconds ahead (delay cancellation). Equation 2.2 shows the transfer function of the low-pass filter.

$$H_{o_f} = \frac{K_f}{1 + T_{l,f}(j\omega)} \quad (2.2)$$

Neuromuscular dynamics are modelled by a second order filter H_{nms} as described by Equation 2.3, after which a remnant is injected as shown in Equation 2.2. A time delay term $e^{-\tau_v j\omega}$ accounts for any other inherent HO delays.

$$H_{nms} = \frac{\omega_{nms}^2}{(j\omega)^2 + 2\omega_{nms}\zeta_{nms}(j\omega) + \omega_{nms}^2} \quad (2.3)$$

The remnant is modelled by low-pass filtering white noise according to the filter described by Equation 2.4, as proposed by Levison et al. (1969).

$$H_n(j\omega) = \frac{K_n}{1 + T_{l,n}(j\omega)} \quad (2.4)$$

Table 2.1 shows the model settings for single-integrator (SI) dynamics, as will be used throughout this project. For SI dynamics, $H_{oe} = K_{e^*} = 1.25$.

Table 2.1: HO model settings

	H_{CE}	τ_f (s)	K_f (-)	$T_{l,f}$ (s)	K_{e^*} (-)	τ_v (s)	ω_{nm} (rad/s)	ξ_{nm} (-)	K_n (-)	$1/T_{l,n}$ (s)
SI	$1.5/(j\omega)$	0.60	1.0	0.20	1.25	0.26	10.5	0.35	8	3.5

2.1.3. Adaptation to Preview Time

To provide context for designing a novel controller capable of successfully adapting to a time-varying HO, this section summarizes how human behaviour changes as a function of preview time. In a preview tracking task it is known that HOs adapt themselves to the following three task variables [17]:

1. Controlled element dynamics
2. Preview time
3. Target trajectory bandwidth

For this project preview time is selected as the driver for time-varying human behaviour for the following reasons. In reality, the chances that the controlled element dynamics change during a tracking task are small, since it would imply that the dynamics of the vehicle are not constant (e.g. mechanical failure). As for target trajectory bandwidth, Van der El concluded that:

“Humans do not systematically adapt their control behavior to the target trajectory bandwidth (between 1.5-4 rad/s) in preview tracking tasks” [17]

This leaves **preview time as the primary task variable which drives time-varying human behaviour in real-world tracking tasks**, for example in situations of reduced visibility due to weather or time of day.

Furthermore, work done by Span investigated the effects of the HSC’s look-ahead time on conflicts and control performance. By varying the preview time in an experiment with an *adaptive* HSC in a similar way, the hope is to shed light on the underlying relationships between the human’s look-ahead time, preview time, and HSC look-ahead time.

Figure 2.2 [17] shows how a limited preview time τ_p affects the target information on the display, i.e. a short preview time will limit the amount of the upcoming signal that is shown to the operator. Note however that the underlying cybernetic structure (as described in Section 2.1.2) is assumed to be constant regardless of preview time. As a consequence, human adaptation to varying preview time can in theory be fully captured by changing the model parameters.

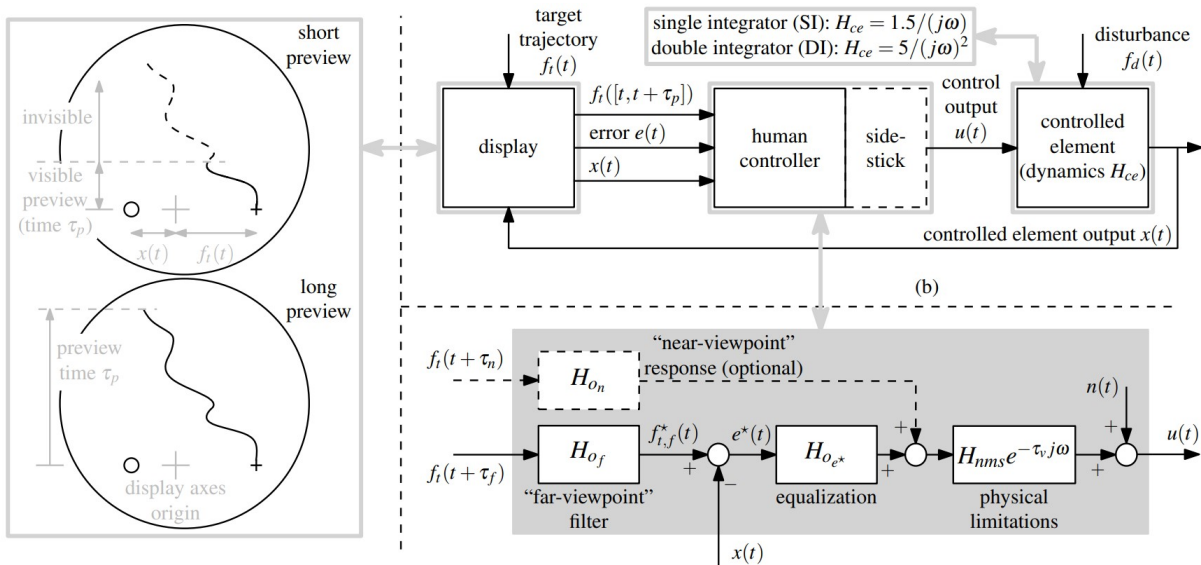


Figure 2.2: Van der El's framework for short and long preview [17]

Hypothetically, if humans somehow employed a different control strategy in light of shortened preview time, this would make the development of an adaptive shared controller vastly more difficult. In this case the controller would likely have to adjust its entire structure to accommodate varying human behaviour, whereas Van der El's work suggests that to account for changing human behaviour in light of varying preview time, a simple adaptation of model parameters would suffice.

Figure 2.3 [13] and Figure 2.4 [13] show the two HO model parameters that are most dependent on preview time for single integrator (SI) and double integrator (DI) CE dynamics.




Figure 2.3: Look-ahead time adaptation [13]

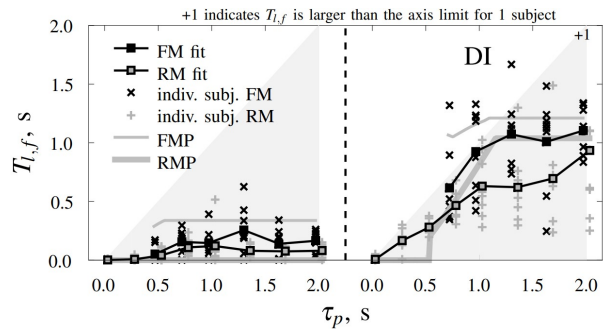


Figure 2.4: LP-filter time constant adaptation [13]

Figure 2.3 [13] shows that the HO look-ahead time τ_f does not increase beyond a critical preview time of approximately $0.6s$ for SI dynamics, and approximately $1.0s$ for DI dynamics. This suggests that around this point the human has optimally cancelled out his/her own phase lags, hence additional preview information is no longer used. For preview times less than critical, HOs seem to use the maximal possible look-ahead time available to them, i.e. $\tau_f = \tau_p$. In this case the operator will *not be able to generate the optimal amount of phase lead to cancel out his/her delay*, hence the output will tend to lag the target [17]. This result is crucial to bear in mind for the novel shared controller design.

Figure 2.4 [13] shows that the low-pass filter time constant $T_{l,f}$ is on average higher for DI dynamics, indicating that higher frequency components of the target signal are ignored. This is similar to behaviour described by McRuer & Jex's crossover model. Furthermore, for zero preview time the entire task regresses to a pursuit tracking task for which the HO is no longer afforded the possibility of filtering out high-frequency components of the target (they are no longer visible), hence $T_{l,f} = 0$.

Main Finding of Section 2.1

The theoretical cornerstone of this project is Van der El's Cybernetic Model for manual preview control tasks. It captures two phenomena:

1. **Corner-cutting:** high-frequency components of the target signal are neglected
2. **Look-ahead time:** using a point up ahead, lead is generated to compensate for HO lags

Concerning human adaptation to preview time, there are two takeaways:

1. For **preview times less than critical**, the HO will no longer be able to generate the optimal amount of phase lead to match the target.
2. As **preview time increases**, subjects respond less to the higher frequency components of the target signal.

2.2. Four Design Choices Architecture

As this project makes extensive use of the Four Design Choices Architecture (FDCA), this section will briefly summarize its philosophy and implications. Span proved that implementing a FDC-HSC minimises conflict torques relative to other types of controllers, such as a meshed HSC [6]. One key strength of the proposed architecture lies in the fact that each of the four design choices play a distinct and physically-interpretable role in the behaviour of the controller [16]:

1. **Level of haptic authority (LoHA)** determines to what extent the controller is capable of overriding inputs from the human and can be interpreted as the 'stiffness' of the guidance system. A rigid controller will yield less to human inputs than a soft controller.
2. **Human compatible reference (HCR)** can be interpreted as the internal trajectory that the controller is steering towards, denoted x_r . This can be taken to be the average of multiple human runs without haptic guidance, but more recently cybernetic models have been used to generate a HCR reference state without the need for prior experiments [6].
3. **Strength of haptic feedback (SoHF)** serves as a corrective force that minimises the difference between the state and human compatible reference. Note that this force does *not* act upon deviations from the target signal: *any mismatch between the HCR reference state x_r itself and the target can never be corrected by this force.*
4. **Level of haptic support (LoHS)** is designed to support the user by providing guidance forces based on upcoming parts of the target signal. These forces are open-loop as they are independent of the state, error, or human input.

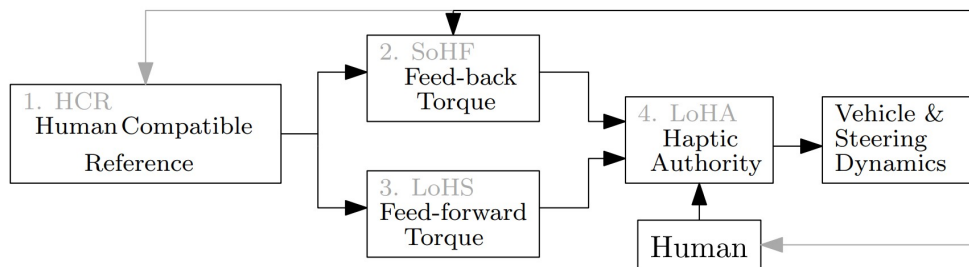


Figure 2.5: FDCA design philosophy [16]


Note that since LoHS, SoHF and LoHA are each simple gains, increasing the LoHA by a factor of two is equivalent to increasing SoHF and LoHS by a factor of two (see Figure 2.5). Therefore, when tuning LoHS and SoHF, it is the ratio between the two that is important, the absolute values only have meaning relative to the level of haptic authority.

Furthermore, in the case that there are no disturbances and no HO, the LoHS can be tuned such that the SoHF channel becomes inactive. In this case the state tracks the HCR perfectly by feed-forward control alone, hence there is no need for correction by the SoHF channel. In practice this is not only undesirable since it would imply no input from the HO (meaning full automation instead of shared control), but also impossible for the following reasons:

Table 2.2: Target and disturbance signals [6]

i	Target signal f_t				Disturbance signal f_d			
	k_t	A_t cm	ω_t rad/s	ϕ_t rad	k_d	A_d cm	ω_d rad/s	ϕ_d rad
1	3	0.731	0.157	4.488	4	0.292	0.209	0.241
2	5	0.731	0.262	5.699	7	0.292	0.367	1.669
3	8	0.731	0.419	1.373	9	0.292	0.471	1.899
4	11	0.731	0.576	5.472	13	0.292	0.681	1.295
5	19	0.731	0.995	1.331	22	0.292	1.152	3.982
6	29	0.731	1.518	5.257	31	0.292	1.623	4.496
7	47	0.073	2.461	5.399	51	0.029	2.670	3.365
8	77	0.073	4.032	3.289	79	0.029	4.136	0.469
9	143	0.073	7.488	2.999	147	0.029	7.697	0.964
10	263	0.073	13.77	5.591	267	0.029	13.98	4.296

Figure 2.7 shows the HO model used by Span to simulate the FDC-HSC setup. This model is based on Van der El's preview model, using the settings given in Table 2.1 where $H_f^{HO} = H_{of}$, $H_e^{HO} = K_e^*$. Note that torque from the HSC is assumed to have no influence on HO behaviour in this model. Furthermore, K_{HO} is assumed to be 1, indicating that the HO does not adjust the magnitude of his/her response relative to a 'vanilla' preview tracking task (no HSC). The settings used in simulations are given by Table 2.1. A remnant is modelled according to $K_n = 8$ and $w_{b,n} = 3.5$ according to Equation 2.4 with $w_{b,n} = 1/T_{l,n}$.

**Figure 2.7:** HO Model [6]

In simulating a shared control scenario with HO and HSC, it is beneficial to model the expected forces on the stick¹. To determine the expected forces an adjustment must be made to the model shown in Figure 2.7. The model was derived according to a black-box method which describes HO behaviour in terms of the relationship between HO input (state variables) and stick output ($u(t)$ in Figure 2.6) [17]ts, meaning that it includes both neuromuscular and stick dynamics. This means that to determine the (simulated) HO forces in the HCR, one must 'undo' the stick dynamics using an inverse-stick block (H_{stick}^{-1}) as applied in Figure 2.7. In doing so, the underlying forces are revealed.

In Span's original work the inverse-stick model is given by Equation 2.6 with the following settings:

$$k = 3.58 \text{ Nm/rad} \quad K_{stick} = 10 \text{ inch/rad} \quad b = 0.22 \text{ rad/s} \quad I = 0.01 \text{ kg/m}^2$$

$$H_{stick}^{-1} = \left(\frac{K_{stick}}{I(j\omega)^2 + bj\omega + k} \right)^{-1} \quad (2.6)$$

Figure 2.8 shows the structure of the FDC-HSC controller. Since the controller was designed to be used in a preview tracking task, the human compatible reference is setup to mimic human behaviour corresponding

¹If no interaction is assumed between a HSC and HO, forces may not be required since summing u_{HSC} and u_{HO} is equivalent to first adding the (inverted) forces and passing the sum through a side-stick model.

to a preview task (settings given in Table 2.1). This is achieved by using Van der El's preview model to generate a guidance force F_R , as well as a reference state x_R . The same inverse-stick model that is used to determine T_{HO} is also used to determine F_R .

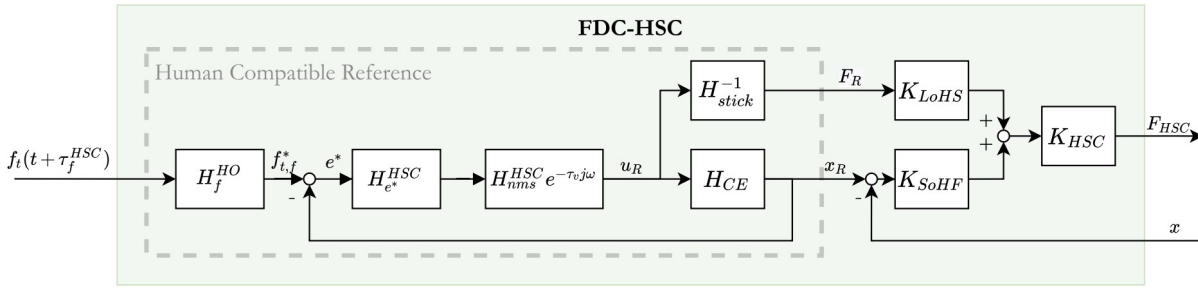


Figure 2.8: Four Design Choices Haptic Shared Controller [6]

Figure 2.8 shows that a CE block is used to produce a reference state x_R , which is compared with the true state x to produce a corrective force through K_{SoHF} . The reference state is also fed back within the HCR to determine an internal error e^* . This error represents the difference between the low-passed phase-shifted target $f_{t,f}^*$ and the reference state x_r .

The gain K_{HSC} represents the level of haptic authority, since the larger this gain, the more guidance torque will be applied (both corrective and feed-forward). K_{LoHS} determines the relative strength of the open-loop guidance channel, whereas K_{SoHF} indicates how strongly the system will 'pull' the current state x back towards the HCR reference state x_R . In this case, K_{LoHS} and K_{SoHF} are tuned such that $K_{HSC} = 1$, meaning that K_{LoHS} and K_{SoHF} implicitly determine the LoHA, which is set to 1.

Table 2.3: FDC and HO gains

K_{LoHS}	K_{SoHF}	K_{HSC}	K_{HO}	
0.6	0.8	1.0	1.0	[-]

In line with Section 2.1, the HCR model in Figure 2.8 steers towards a point on a low-pass filtered trajectory τ_f seconds into the future. Note that a remnant is lacking, which effectively makes the HCR an idealized HO adapted to preview information. As a result, the HCR reference state x_R represents the expected output of the system in the case of an *idealized, linear, time-invariant human operator*. This has two crucial benefits:

1. Guidance torques have a high likelihood of being accepted by the human, since the controller is trying to replicate the average operator's behaviour and control strategy, including corner cutting and delay cancellation. One could also point out that the system is, to an extent, subject to the same physical limitations as its human counterpart (neuromuscular filter and time delay), meaning that it essentially 'acts' as if it had a body!
2. Since the HCR lacks a remnant which accounts for the stochastic and non-linear element of human behaviour, it is expected to out-perform its human counterpart (assuming that the non-linear component of human behaviour does not contribute to superior control performance).

Main finding of Section 2.3

By choosing to implement Van der El's cybernetic model as a FDC human compatible reference, performance is expected to increase and HO-HSC conflicts are expected to be minimal. The HSC settings are taken to be the same as the 'average' HO, allowing the FDC-HSC to emulate the behaviour of a linear, time-invariant, idealized (no remnant) human being.

2.4. Effects of Controller Look-Ahead Time

2.4.1. Simulation and Experiment Results

Work by Span looked at the effects of varying look-ahead times using a FDC-HSC with a HCR configured according to Van der El's model, revealing a region of optimal look-ahead time's for $\tau_{HSC} = [0.7s, 0.8s]$ where conflict is minimal [6]. This corroborates with Van der El's work which proved that a critical look-ahead time exists for values of τ_f around the same value, beyond which additional preview information no longer influences the HO's behaviour.

Furthermore, these results also make intuitive sense: by adjusting the settings of the controller to match the expected behaviour of the average human, conflicts are evidently reduced. In the case of Span's FDC-HSC experiment the preview time was set to $\tau_p = 1.0s$, meaning that the HO was likely employing a look-ahead time of around $\tau_f = 0.6s$. When the HSC's look-ahead time was set to approximately the same value, $\tau_{HSC} \approx 0.6$, conflicts were found to be reduced. Therefore, by aligning the HSC's (HCR) settings with the parameters that describe HO behaviour (i.e. same look-ahead time), conflicts between the HO and HSC are minimal.

Using the parameters defined in Table 2.1 and Table 2.3, simulation results were then averaged over 1000 realizations yielding the results in Figure 2.9, Figure 2.10, Figure 2.11, and Figure 2.12. These simulations were initialised with the settings given in Section 2.3. Figure 2.9 shows the experiment tracking error and control activity relative to simulations, clearly suggesting a region of optimal look-ahead times around $\tau_{HSC} = 0.6s$ for both the simulations and the experiment. Figure 2.10 suggests that on average, the simulations *underestimate* the amount of torque applied by both the HO and HSC. This can be observed even more clearly in Figure 2.12, which shows that there was more conflict between the HO and HSC during the experiment than predicted for extreme values of τ_{HSC} . Figure 2.11 again supports the notion that the HSC and HO are optimally aligned around $\tau_{HSC} = 0.7s$.

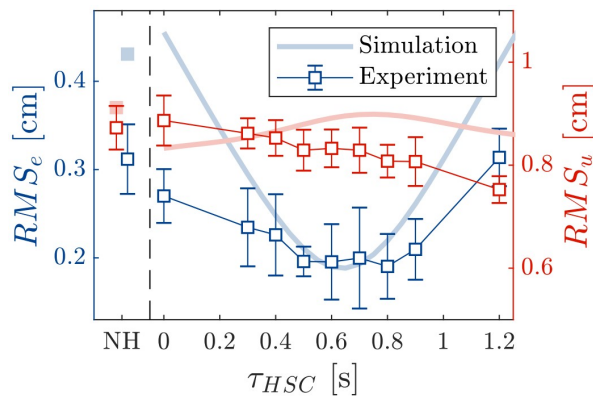


Figure 2.9: Tracking error and control activity [6]




Figure 2.10: Torque magnitudes [6]




Figure 2.11: Conflict time [6]




Figure 2.12: Conflict torque [6]

The main takeaway from Figure 2.9 and Figure 2.12 is that the simulations underestimate *both* tracking performance and conflicts for extreme values of τ_{HSC} . This suggests that HOs are willing to endure conflicts to achieve superior tracking performance. Because the model's main limitation is that it cannot account for direct physical interaction between the HSC and HO through the side-stick, this is likely the channel through which HO instigate conflicts in real-life to improve tracking performance.

2.4.2. Comparing Experiment Data with Simulation Results

The difference between experimental and simulation results around extreme values for τ_{HSC} is hypothesised to be due to human adaptation to the controller [6]. It is important to realise that in the simulations performed by Span, the HO is assumed not to react *directly* to torques from the controller, which is the primary benefit that HSCs have to offer in real-world scenarios.

However, in Span's simulations the HO *does* act upon the state x , which is influenced by torques stemming from the HSC. In simulations without explicit human adaptation to force-feedback, HSC torques are assumed to directly control the side-stick in parallel with the HO without impediment. This entails that the only method of interaction between the HO and HSC is through the state x , which is inherently slower than direct force-feedback due to the delays introduced by the side-stick and controlled element. This 'slower' path is illustrated by the green line in Figure 2.13.

Another way to consider this fact is by realising that an equivalent simulation would be to model two side-sticks through two independent channels, with one for the HSC and one for the HO. By summing the output of each side-stick as shown in Figure 2.13 the same results would be achieved as in the single side-stick model without adaptation. This makes intuitive sense, since the model assumes no direct physical interaction between the two parties through the side-stick.

Furthermore, even once the (simulated) HO has 'recognised' the delayed effects of the HSC on the state x , it must still process these inputs (adding more delays such as τ_v) and pass them through its own side-stick. In conclusion: due to the inherent structure of the baseline FDC-HSC simulation the HO is severely limited in its capacity to correct the HSC in the case of poor performance. This explains the discrepancy in tracking performance and conflicts between the simulation and experiment data for extreme values of τ_{HSC} .

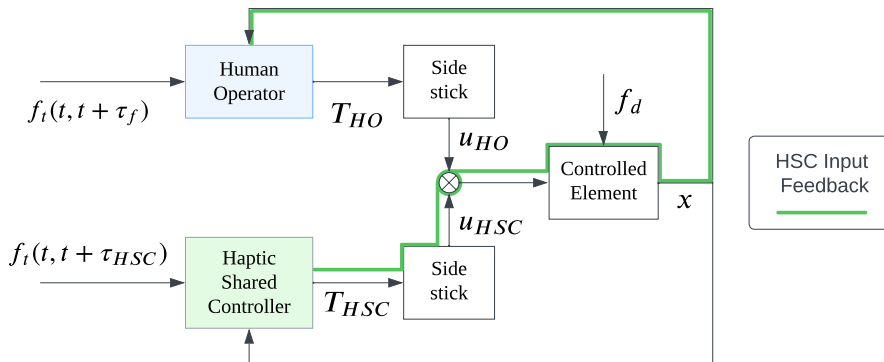


Figure 2.13: Equivalent FDC-HSC model without HO adaptation

2.4.3. HSC Performance for Varying Look-Ahead Times

The reason that conflicts occur for extreme values of τ_{HSC} is likely because the HO is compensating for the poor performance of the controller in these configurations. Evidently, the HO is aware that the HSC is not aiding the situation and is therefore willing to undergo conflicts to rectify the sub-optimal inputs from the HSC.

The hypothesis that the HSC is performing poorly for extreme values of τ_{HSC} is further supported by Figure 2.14, which shows the closed-loop behaviour of the HSC *without* the HO. Note that for $\tau_{HSC} = 0.8$ performance is optimal with $RMS_e = 0.2 \text{ cm}$, which is very similar to the level of performance achieved when the HO and HSC work together (Figure 2.9).




Figure 2.14: HSC error and control activity without HO

Figure 2.9 shows that *without* the HSC, the HO achieves a tracking performance of $RMS_e = 0.32\text{ cm}$ in the non-haptic (NH) condition. It is therefore to be expected that if the HSC is significantly under-performing this benchmark, then the HO will be dissatisfied with the performance of the controller (since the HO could do a better job if the controller was disabled). Figure 2.14 shows that for look-ahead times outside of $\tau_{HSC} = [0.6, 1.0]\text{ s}$ the controller achieves *worse* performance than the HO would if it performed the task alone. This is because for these values, the reference state either leads or lags the target signal, substantially decreasing tracking performance.

Figure 2.15 shows the effect of HSC look-ahead time on the state x in the case of a *closed-loop* HSC without HO. For $\tau_{HSC} = 0.3\text{ s}$ both the reference state and actual state tend to lead the target f_t . The reverse is true for $\tau = 1.0\text{ s}$, where x and x_r lag the target instead. The reason for this behaviour is due to the fact that the SoHF channel will always steer the state towards the reference x_r . For extreme look-ahead times the HSC-HCR itself will be either be leading or lagging the target, hence the state will be steered towards a sub-optimal reference.




Figure 2.15: Reference state phase shift, HSC without HO

Since the reference state is simply what the controller 'thinks' would happen if a human would be behind the wheel, there is no amount of gain-tuning or tweaking that could ever resolve the fundamental mismatch between x_r and f_t . The only remedy is to either design an additional feedback loop (furthermore convoluting an already complex controller), or simply correct the settings of the controller such that it produces a reference state that *is* aligned with the target.

The reason why the controller is not aligned with the target (apart from $\tau_{HSC} = 0.65$) can be attributed to the implementation of Van Der El's model as a HCR. Since this model describes how HO behaviour has evolved to compensate for inherent human delays (e.g. τ_v), its parameters are derived in such a way² that the look-ahead optimally compensates the model's lags. If this was *not* the case, then the model would not prove successful in describing HO behaviour, since in reality humans have evolved to use preview information (lead) to optimally compensate for our own biological limitations (lag). Therefore, one cannot expect to tweak model settings such as look-ahead time without voiding the HCR's accuracy in describing HO behaviour, hereby negatively impacting tracking performance and 'breaking' a well-tuned model.

²Not accounting for individual differences between HOs.

Figure 2.16 shows the situation in which the look-ahead time of the controller is tuned to a more realistic value of $\tau_{HSC} = 0.65$. In this case the target and reference state are optimally aligned in terms of phase shift. Since the SoHF channel minimises the error between x and x_r , if x_r happens to align with the target, then the state will also *seemingly* track the target (it is actually following the reference state, see Figure 2.8).




Figure 2.16: Reference state alignment, HSC without HO

Although correct cancellation of model delays (e.g. τ_v) through look-ahead time lead (τ) is essential, in theory, there is one way the look-ahead time could be reduced without a significant loss of performance. Figure 2.17 shows the effects of setting $\tau_v = 0.0$ in combination with reducing the look-ahead time to $\tau_{HSC} = 0.4$. By reducing both delay *and* lead equally, the final reference state happens to align with the target. This finding could prove useful for the case in which the HSC only has a limited preview time (e.g. $0.4s$), but still retains Van der El's model characteristics.




Figure 2.17: Reference state alignment, HSC without HO, $\tau_v = 0$

Main finding of Section 2.4

Extreme values of τ_{HSC} lead to conflicts between the HO and HSC, as well as a reduction in performance. For high look-ahead times the reference state tends to lead the target, for low look-ahead times the reference state tends to lag the target. The degree to which the reference state and target are aligned plays a pivotal role in both tracking performance and conflicts. If the HCR is not properly tuned, i.e. if the lead generated by look-ahead anticipation does not cancel out the model delays, then there will be a guaranteed mismatch between the reference state and target.

Controller Design and Simulations

This chapter aims to describe the design process of the novel adaptive HSC, starting with the motivation behind the controller and the design methodology in Section 3.1, followed by Section 3.2 which replicates and validates results by Span for a FDC-HSC. Thereafter, Section 3.3 evaluates the usability of system identification results in capturing individual HO differences. Section 3.4 investigates the occurrence of conflicts and their relationship to the reference state of the HCR. This is followed by Section 3.5 which details how HO rejection of HSC torques is modelled.

The second half of the chapter describes the design of the novel controllers. This starts with Section 3.6, which lays out the test scenario for time-varying HO behaviour. Section 3.7, Section 3.8, and Section 3.9 describe the design and results for three adaptive HSCs: inertial, model-fitting, and cross-correlation, respectively. The performance of the controllers is compared in Section 3.10, after which additional findings are presented in Section 3.11.

3.1. Motivation and Methodology

The primary goal of the project is to develop an adaptive HSC. Making the controller adaptive has numerous benefits relative to previous designs:

1. There is no longer a need for manual tuning, a process that was often done heuristically and therefore subject to individual preferences of the designer.
2. Performance is expected to increase and conflicts are expected to reduce relative to a fixed controller, since by adapting itself the controller will 'find' the optimum settings. This entails that the controller will in theory be able to adapt to individual operator differences, assuming that these differences are significant enough to cause measurable anomalies in conflict and/or tracking performance.
3. By adapting solely based on performance and conflicts, the system is not limited to current hypotheses about the structure of the 'optimal' controller. In this way, the hope is that the controller will empirically lead to new insights on operator preferences in a haptic control setting. This is particularly useful in the event of time-variant operator behaviour as a consequence of limited preview information (e.g. due to weather).

The proposed methodology for designing an adaptive haptic shared controller (A-HSC) is summarised in the following steps:

1. Validation of FDC-HSC Model
2. Evaluation of System Identification Results
3. Conflict and HCR Analysis
4. Modelling HO Rejection of HSC Torques
5. Adaptive Haptic Shared Controller Design

The cornerstone of this project is a set of MATLAB simulations which are used to test and evaluate prospective adaptive controllers. By building an understanding of their strengths and limitations and extending their validity, the models are optimally leveraged in the design process of the novel adaptive controller. For this reason the chapter starts by validating the FDC-HSC model used by Span [6].

Thereafter previous experiment data and system identification results are used to build personalised cybernetic models of individual subjects. These models are then tested to evaluate if the new methods lead to a more accurate representation of subjects' behaviour, possibly leading to improvements over the current one-size-fits-all HCR model.

Focus will then shift towards the design of an adaptive controller by first performing a conflict analysis on both LoHS and SoHF channels, as well as developing a model to capture HO rejection of HSC torques.

The final step is to propose and test a structure for an adaptive controller with varying look-ahead time. Since the effects of look-ahead time on performance and conflicts is known from Span's research [6], this provides a reference that can be used to evaluate the novel adaptive controller. Testing the new controller and verifying that it leads to reduced conflicts and increased performance is the last step before evaluating it during a real experiment.

3.2. Validating FDC-HSC Model

The replica-simulation model parameters are set according to Table 2.1 and Table 2.3. The only difference between the replica-simulation and Span's original work is the use of a different inverse-stick model to reveal the underlying torques. For convenience a simplified inverse-stick model was used in the replica model, as given by Equation 3.1.

$$H_{stick}^{-1} = \frac{k}{K_{stick}} \quad (3.1)$$

Whereas in Span's original work Equation 3.2 was used as an inverse-stick model, consisting of the inverse of the transfer function describing the full stick dynamics. Note that this model is improper and non-casual. To circumvent this issue during simulation, Span combined this transfer function in series with its adjacent blocks as shown in Figure 2.8, such as H_{nms} . Due to the presence of poles in the transfer functions of the adjacent blocks the system becomes casual again and therefore realizable.

$$H_{stick}^{-1} = \frac{I(j\omega)^2 + bj\omega + k}{K_{stick}} \quad (3.2)$$

Implementing the simplified stick model as in Equation 3.1 yielded negligible difference in results compared to the full model as shown in Equation 3.2. This was validated by comparing actual LoHS torque from experiment data with simulations using the simplified stick-model. To measure the fit between the two signals, the Variance Accounted For was calculated using Equation 3.3, where \hat{T}_i is the simulated LoHS torque determined using the simplified model (Equation 3.1), and T_i the experiment LoHS torque.

$$VAF_i = \left(1 - \frac{\text{var}(T_i - \hat{T}_i)}{\text{var}(T_i)} \right) \cdot 100\% \quad (3.3)$$

Table 3.1: Simplified inverse stick VAF

Fofu nr.	1	2	3
VAF	94.93%	94.97%	94.98%

The LoHS channel is ideally suited for comparing the simplified and full inverse-stick model, since it is always independent of the state and HO input. Therefore, it can be said with certainty that the only reason the VAF values in Equation 3.9 are not 100% must be due to a difference in HCR structure or parameters. And indeed, the only difference in HCR between the two is that the experiment uses the *full* inverse-stick model, whereas the simulation uses a *simplified* inverse-stick model.

Figure 3.1 shows the LoHS channel for both full and simplified inverse stick-models. The minor phase shift between the experiment (full) and simulation (simplified) occurs due to the neglected dynamics in the simplified model. To go from Equation 3.2 to Equation 3.1 one can imagine multiplying with a

second-order filter, effectively dampening high frequency components and shifting the signal in time as shown in Figure 3.1. Because of the multiple low-pass filters in the HCR model (H_{nms} , H_{of}) most of the high frequency components are already removed, hence why the simplified model results in a negligible difference in results.




Figure 3.1: Full and simplified inverse-stick model T_{LoHS}

Using the experiment data from [6], each experiment run was simulated by initialising a remnant with a random seed, using the same forcing function as from the experiment run. The results are shown in Figure 3.2 and Figure 3.3, agreeing well with Figure 2.9 and Figure 2.12. Note that the replica-simulations also underestimate the error for extreme values of τ_{HSC} as evident in Span's work.




Figure 3.2: Validation tracking error

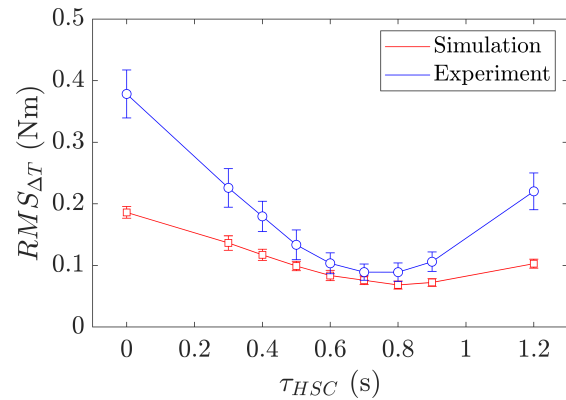


Figure 3.3: Validation conflict torque

Main finding of Section 3.2

Simulations were performed with a FDC-HSC controller to validate Span's work. It was found that a simplified inverse-stick model can be used without significantly affecting the accuracy of results. The simplified model essentially applies a low-pass filter, hereby introducing a small amount of phase lag (due to the presence of other filters, high-frequency components are already removed). To validate the use of the simplified model, Variance Accounted For (VAF) was computed between experiment T_{LoHS} and simulated T_{LoHS} , with all VAF values exceeding 94%.

3.3. Evaluation of System Identification Results

Span's work included fitting a model to individual subjects for the non-haptic condition, in the hope that a correlation could be identified between the identified look-ahead and the optimal look-ahead time found during the experiment [6]. However, only a weak correlation between the two was found, with an insignificant p-value of $p = 0.174 > 0.05$. This is a key finding that suggests the following hypothesis:

Individual differences between HO's are not large enough to produce noticeable changes in optimal controller settings (e.g. τ_{HSC})

The above implies that a one-size-fits-all HCR works equally well for the average participant compared to fitting a personalised HCR for each subject. To test this, system identification was performed on each subject in order to build a 'personalised' HCR that is expected to more accurately reflect each HO's individual behaviour. This hypothesis was then tested by simulating each of the subject runs using models that incorporate the system identification results *per subject*.

To model individual subjects' behaviour the same methodology was used as laid out by Span [6]. This method estimates parameters in the frequency domain according to the structure laid out in Figure 3.4. The HO model consists of two channels, with H_{o_t} describing the response to the forcing function as described by Equation 3.5 and H_{o_x} incorporating the response to the state x as given by Equation 3.4.

$$H_{o_x}(j\omega) = K_e e^{-\tau_v j\omega} \quad (3.4)$$

Note that Equation 3.5 captures the corner cutting behaviour and look-ahead time, whereas Equation 3.4 captures the inherent time delays of the HO.

$$H_{o_t}(j\omega) = H_{o_x} \frac{K_f}{T_{l,f} j\omega + 1} e^{\tau_{HO} j\omega} \quad (3.5)$$

As this structure does not assume any stick or neuromuscular dynamics, it is considered a *reduced model* (RM). Span proved that it is capable of capturing human behaviour for the non-haptic condition, with variance-accounted-for VAF values exceeding 90% for *all* subjects [6].




Figure 3.4: Reduced model (RM) for parameter estimation

The high VAF values initially suggest that neuromuscular dynamics can be omitted in the HCR, as long as the rest of the HCR parameters are updated according to the parameter estimation results yielded by the structure given in Figure 3.4. To evaluate the effects of omitting neuromuscular and stick dynamics, a full-model (FM) is tested as shown in Figure 3.5. To prevent increased complexity during parameter estimation, H_{nms} is assumed to be equal for all subjects according to Equation 3.6.

$$H_{nms} = \frac{\omega_{nms}^2}{(j\omega)^2 + \zeta_{nms}\omega_{nms}(j\omega) + \omega_{nms}^2} = \frac{10.5^2}{(j\omega)^2 + 0.35 \cdot 10.5 \cdot (j\omega) + 10.5^2} \quad (3.6)$$




Figure 3.5: Full model for parameter estimation

The results of the FM compared to the RM for all 10 subjects are given in Figure 3.6. For each NH run the parameter vector $\Theta = [\tau_{HO}, K_f, T_{l,f}, K_e, \tau_v]$ [6] is estimated by minimising the squared magnitude of a cost function, shown in Equation 3.7 [6].

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \sum_{i=1}^{N_f} |\epsilon(j\omega_i | \Theta)|^2 \quad (3.7)$$

The cost function itself is defined by Equation 3.8 for a given parameter vector Θ , with X as the Fourier transform of the state from experimental data, F_t being the Fourier transform of the target and U being the Fourier transform of the input.

$$\epsilon(j\omega_i | \Theta) = U(j\omega_i) - \left(\hat{H}_{o_t}(j\omega_i | \Theta) F_t(j\omega_i) - \hat{H}_{o_x}(j\omega_i | \Theta) X(j\omega_i) \right) H_{nms} \quad (3.8)$$

Figure 3.6 shows the effect of neglecting the neuromuscular system and side-stick dynamics on parameter-estimation results. For $T_{l,f}$ and τ_{HO} little difference is to be found, which makes intuitive sense since neuromuscular and stick dynamics are not necessarily expected to influence the amount of corner cutting or look-ahead time.

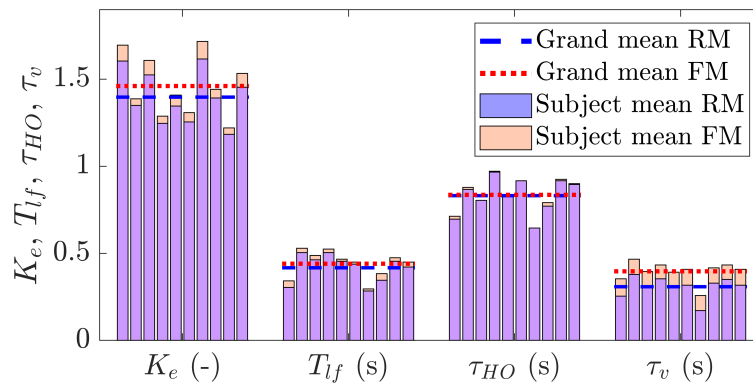


Figure 3.6: System identification results


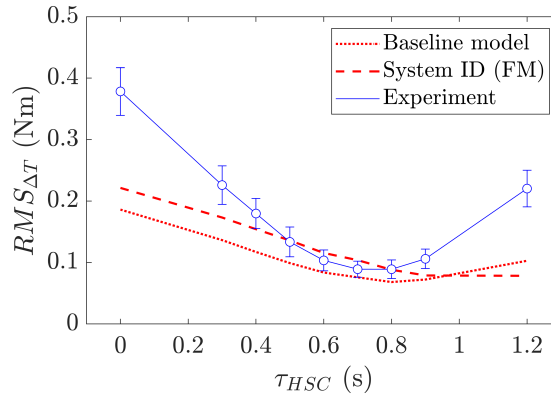
Table 3.2 shows that K_e decreases marginally for the FM and that τ_v decreases by 29%. Therefore, the main conclusion is that neuromuscular and side-stick dynamics are implicitly accounted for and 'lumped' into τ_v in the case of the RM. This is important to consider when designing a HCR according to parameter identification results. For example, when using the RM parameters in a HCR it is important to *omit* the NMS in the HCR, as apparently this has a significant impact on τ_v .

Table 3.2: Grand means RM versus FM

	K_e [-]	T_{lf} [s]	τ_{HO} [s]	τ_v [s]
RM	1.46	0.44	0.84	0.40
FM	1.40	0.42	0.83	0.31
Δ	4.3%	4.8%	1.2%	29.0%

Figure 3.7 and Figure 3.8 show the effects of using each subject's personalised parameter estimation results to model their individual behaviour *and* adjust the HCR parameters accordingly. Note that in this case the FM was chosen to perform system identification, meaning that the original HSC-HCR structure could be used (i.e. not omitting the NMS).

It is evident from Figure 3.7 and Figure 3.8 that by adapting the HCR and HO model, on average little benefit is gained in terms of accuracy relative to the experiment results. This supports the hypothesis that individual differences as shown in Figure 3.6 are not significant enough to yield an improvement in simulation accuracy relative to the one-size-fits-all baseline model as described in Section 2.4. Furthermore, human adaptation seems to have a far more significant effect on model accuracy than individual differences between HO's.

**Figure 3.7:** Personalised simulations tracking error**Figure 3.8:** Personalised simulations conflict

Main finding of Section 3.3

Parameter estimation results by Span are used to model individual HO behaviour, resulting in personalised HCRs. Using this method to simulate Span's experimental data did not result in increased performance (in terms of tracking error) relative to the baseline model, suggesting a one-size-fits-all model is adequate to model HO behaviour in an FDC-HSC controller. Furthermore, (the lack of) human adaptation seems to have a substantially stronger effect on model accuracy than individual differences between HOs.

3.4. Conflict and HCR Analysis

3.4.1. Root Causes of Conflicts

One of the driving factors behind the design of haptic control systems is the minimisation of conflict forces, which occur when the operator disagrees with the control inputs of the automation and decides to counteract them. In the worst case the user may even decide to disengage the automation altogether, negating the possible benefits of the guidance system.

Designers have two primary tools at hand to reduce conflicts. Firstly, the design architecture of the controller itself plays a large role in ensuring that the guidance system's control strategy is aligned with the human. To this end, cybernetic models are used to capture and replicate human control behaviour.

Secondly, a range of parameters within the controller can be tuned to further ensure satisfactory behaviour and alignment with the user. Some of these parameters pertain to the HO model (for example, look-ahead time), whereas others are a matter of preference or design choice (for example, the strength of guidance input relative to the user's input).

The details of how humans adapt to haptic controllers are yet unknown, although it can be expected that operators will counteract the guidance forces in the case of one or more of the following occurrences:

1. The operator expects that the system's inputs will have a negative impact on control performance (i.e. increase the error between the target and state)
2. The operator's intended or present control inputs run counter to the system's inputs
3. The operator lacks trust in the system

The expected result of one or more of the above is that the operator will seek to neutralise the guidance forces, negating their effect on the controlled element (CE) state x .

Note that (1) and (2) may coincide in practice, assuming that the human's control strategy is to minimise the error with the least amount of effort. Therefore, if the user anticipates that the system's inputs will lead to an increase in error, the chances are high that these inputs will also be in opposition to the human's control strategy. However, in theory it is possible that the human's intended control inputs do not align with the system's, despite the fact that the system's inputs would lead to better performance.

Consider the case in which the guidance system has a faster and more ambitious control strategy than the human. In this example conflicts may arise from the fact that the human is unable (or unwilling) to keep up with the automation, despite both parties being able to accomplish the task (albeit with varying performance).

If situations (1) or (2) occur, the implicit effect on the operator will likely be a lack of trust (3) and subsequent neutralisation of the system's control inputs through conflict torques. Although it may be possible that due to stress or inherent lack of confidence in the automation, the operator is reluctant to accept guidance forces *regardless* of their potential benefit.

In practice it is very difficult to distinguish between (1), (2) and (3) since it is highly unlikely that they occur independently of each other. This is the reason why modelling human rejection of guidance forces poses a significant challenge. However, it is still useful to contemplate the root causes of conflicts to shed light on the importance of the following design options:

1. **Human compatible reference:** To reduce the chances of the operator suspecting that the system will have a negative impact on control performance, a simple constraint can be implemented: the HCR reference state x_R should be designed to be as close to the target as possible, subject to the HO model constraints (such as the neuromuscular system or limited preview time). This way the HSC is guaranteed to result in a satisfactory tracking error, increasing the chances of acceptance.
2. **HO identification:** In an idealized simulation with cybernetic models, the guidance system will perfectly replicate the human's inputs in the case of zero remnant (i.e. the human is modelled as a linear time-invariant system with optimal parameters). In practice, it is inevitable that the guidance system will not be able to perfectly capture the human's behaviour due to a) the remnant, and b) sub-optimal model parameters. To account for this, system identification methods can in theory be used to capture individual differences between operators¹.

¹Albeit in the case of this project, individual differences between HO subjects are evidently not significant enough.

3. **Display:** For optimal trust and transparency, it may be beneficial to show the user a portion of the HCR, or a graphical representation of the current guidance inputs. In the case that the guidance is sub-optimal, this would allow the user to anticipate conflicts similar to how he or she anticipates the target signal. Furthermore, the operator would also have the possibility of building an intuition about the cause of conflicts by relating them to the system's HCR signal (for example, observing if the reference signal consistently lags or leads the target).

Finally, although it has been proven that conflict torques can be reduced by adopting a controller *structure* that mimics the human's behaviour, it is yet unknown exactly what set of *parameters* lead to optimal control performance and minimisation and conflicts. For example, it may be advantageous to implement a controller that reacts *faster* than the HO (e.g. less time delay), but is still subject to a cybernetic structure that describes the user's overall control behaviour. In this case the set of parameters that best *describe* the operator's behaviour may not correspond to the set of HSC parameters that best *support* the operator's behaviour.

3.4.2. LoHS and SoHF Conflicts

One of the benefits of the FDC architecture is that designers are able to separately tune the open-loop support channel (LoHS) and the feedback channel (SoHF). Because these channels both have a physically interpretable meaning, it is useful to consider how they individually contribute to conflicts and tracking.

Figure 3.9 shows the forcing function and torque taken from a single run of Span's experimental data, with $\tau_{HSC} = 0.0s$ (both HO and HSC active). From Span's work it is known that this condition is prone to inducing conflicts, since the controller is not producing any phase lead to compensate for its HCR delays (e.g. $\tau_{v_{HSC}}$). The effect of this is clearly seen by the fact that the reference signal x_r significantly lags the target, meaning that without the HO the HSC would ensure that the state would lag the target by tracking x_r instead.




Figure 3.9: Experiment target and torque for $\tau_{HSC} = 0.0s$

Figure 3.9 also demonstrates the effect of HO adaptation, showing that for the majority of the time, the HO is giving contradictory inputs to the HSC. It is these inputs by the HO that compensate for the phase lag generated by the HSC continuously trying to 'pull' the state x back towards a lagging reference state x_r .

Figure 3.10 shows the occurrences of conflict for the same run, analysed per HSC channel with opposite inputs to the HO marked in red. It can be clearly seen that the vast majority of conflicts are due to the SoHF channel. This makes sense, as this channel is by nature a corrective force that will react to deviations of the state x from the reference x_r . Therefore, the more successful the HO is in correcting the lagging x_r to improve tracking performance, the more feedback the SoHF channel will exert to 'correct' the HO! For this reason, the SoHF channel and HO tend to *amplify* each-other's control activity, but only when the reference x_r deviates substantially from the target.




Figure 3.10: Experiment T_{SoHF} and T_{LoHS} for $\tau_{HSC} = 0.0s$

Furthermore, Figure 3.10 shows that when x_r crosses f_t (see blue in top figure) there is a short reduction in conflicts for T_{SoHF} . This is because for a moment, the difference between x_r and x is small. Conversely, T_{LoHS} exhibits the opposite of this behaviour as it is prone to conflict at exactly these locations in time (compare the red regions in the second and third row of Figure 3.10). This can be explained as follows.

It is known that changing the look-ahead time of the controller has the effect of phase shifting the reference signal x_r , resulting in a pure phase shift of the output torque of the LoHS channel. In the grand scheme of things this phase shift has little effect on conflicts since the majority of the time, the LoHS channel is still 'on the right side of the fence': when the target is primarily positive, the LoHS channel is also positive and vice versa. This is illustrated by Figure 3.11, which shows the effect of phase shifting two pure sine waves, highlighting the regions where they are opposite in sign. It can be seen that in Figure 3.10 and Figure 3.11 most conflicts occur around $T_{LoHS} = 0$, with the duration of conflicts being longer for a larger phase shift.




Figure 3.11: Conflict regions between two sine waves with a phase shift

The right hand-side of Figure 3.12 shows the same data for a run with minimal conflicts ($\tau_{HSC} = 0.8s$). As expected, in this case the reference state x_r almost perfectly follows the target, hereby ensuring that the SoHF will always 'pull' the state towards the target signal since x_r is now aligned with f_t . The effect of this is a drastic reduction in conflict torques relative to $\tau_{HSC} = 0.0s$, with the remaining conflict that does still occur being significantly smaller in magnitude.

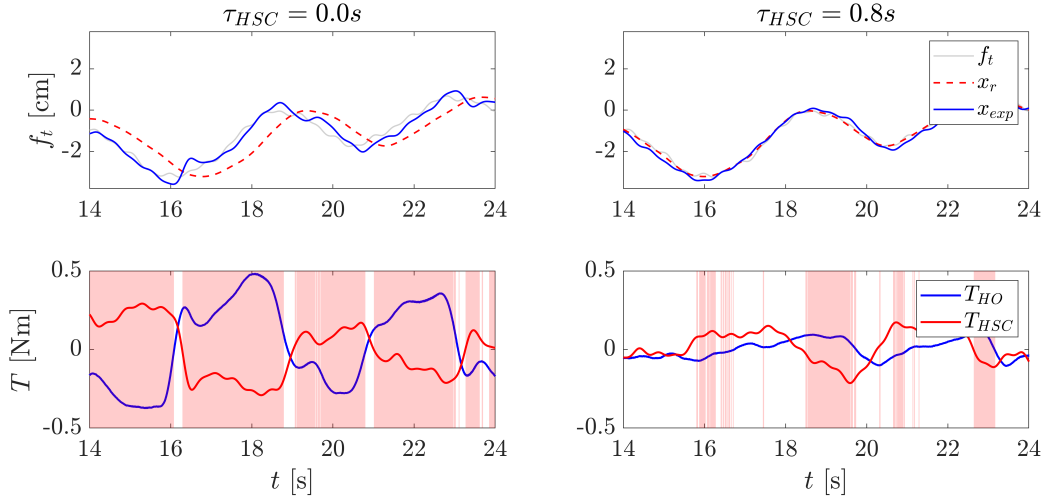


Figure 3.12: Experiment T_{SoHF} and T_{LoHS} for $\tau_{HSC} = 0.0s, 0.8s$

Main finding of Section 3.4

HO-HSC conflicts are mostly due to the SoHF channel, which becomes active in the event that $x \neq x_r$. If the HSC is using extreme values for τ_{HSC} , then $x \neq x_r$ will be the predominant case, hence leading to significant conflicts. The more successful the HO is in rectifying the offset of the reference state by 'pulling' the state back to the target, the larger the difference between x and x_r , further activating the SoHF channel. For this reason, HO and HSC torques tend to amplify each-other in the event of sustained SoHF conflict.

3.5. Modelling HO Rejection of HSC Torques

This section proposes a model to capture the behaviour of HO's when τ_{HSC} deviates substantially from the look-ahead time of the HO. Modelling this aspect of human behaviour will help to improve the accuracy of the simulations used to develop the adaptive HSC, hereby aiding the design process. Figure 3.13 illustrates the proposed amendment to Van der El's existing structure, incorporating a *rejection* loop that supplements the original torque modelled by Van der El, T_{HO^*} , with an additional *rejection torque* T_{Δ} .

The amount of rejection torque T_{Δ} applied is dependent on two factors:

1. The difference between the HO's expected applied torque T_{HO^*} and the current felt torque of the HSC through force feedback, T_{HSC} .
2. The magnitude of the HSC torque itself, as this will determine to the strength of rejection in the event that it occurs.

Note that this model assumes the force feedback from the stick as felt by the HO to be 100% accurate with zero time delay.

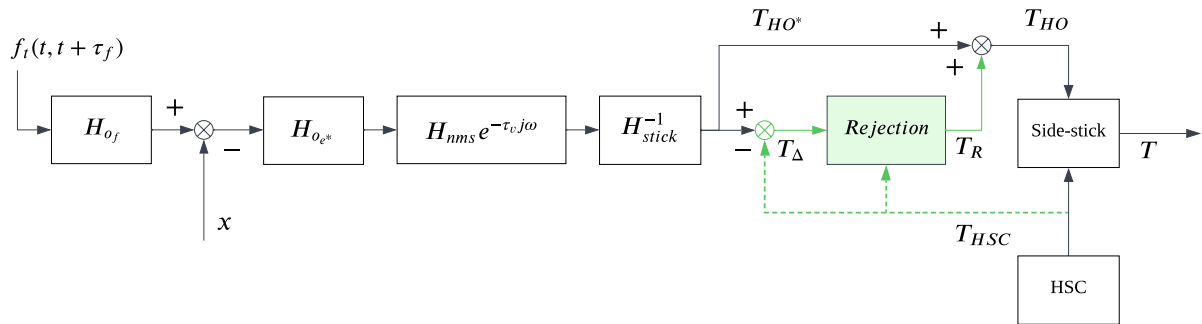


Figure 3.13: Adaptive HO model

Figure 3.14 shows the inner workings of the rejection block. This consists of taking the root-mean-squared (RMS) value of T_{Δ} as an indication for how much the haptic controller is currently operating counter to the intentions of the HO. Alternatively, this can also be a running RMS to simulate a buildup (or lack) of trust over time. Thereafter, the RMS block output is multiplied with a gain before it passes through a saturation-limiter with limits $[0, -1]$, after which it is multiplied with the current HSC torque. Finally, the product is passed through a low-pass NMS filter taken to be the same as H_{nms} in Van der El's preview model.

The effect of this structure is to reverse the effects of the HSC torque by directly counter-acting them (negative gain), *if* they run counter to the intended torques of the HO for the current state x . To prevent this loop from being active in the case that the HSC and HO inputs are of the same sign, a saturation limiter is used, also ensuring that the rejection forces never exceed the original HSC torque as felt through force-feedback. By adapting the rejection gain K_r the amount of counter-torque (rejection) for a given T_{Δ} can be adjusted. This allows for modelling varying degrees of rejection, with $K_r = 0$ resulting in no rejection forces (no HO adaptation).




Figure 3.14: Rejection block structure

Effectively, this structure makes the following assumptions:

1. Human rejection of conflict torques only occurs when HSC torques run opposite to the HO's intentions.
2. Humans are able to feel the magnitude of applied HSC torque without delay and noise.
3. In the event of conflicting inputs, humans will attempt to cancel out the HSC's inputs by applying torques in the opposite direction, with a magnitude that is at most equal to the original HSC torque causing the conflict.

The effects of this change to the model are best illustrated by considering how the original baseline simulation performed in the case of high conflict. This situation is shown in Figure 3.15 where simulated conflict torques and compared with experiment data.

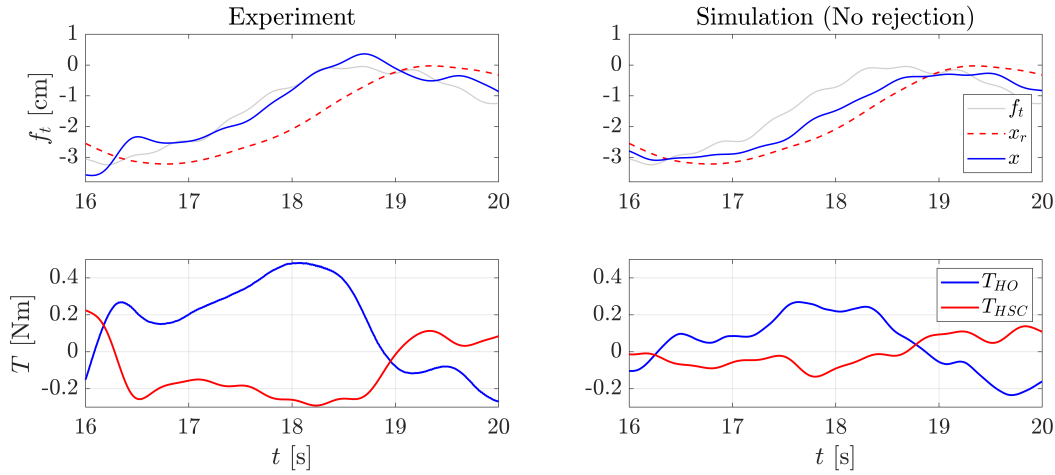


Figure 3.15: Simulated versus experiment torques for $\tau_{HSC} = 0.0s$

Note that due to the effects shown by Figure 2.13, even without modelling rejection, the simplified HO model is capable of correcting the poor performance of the HSC to an extent. This is evident from the fact that the simulated state x is situated between the target and reference x_r , indicating the the simulated HO is 'pulling' the HSC towards the target and introducing lead. Although the simplified simulation captures the fact that HO and HSC torques run counter to each other in terms of sign, Figure 3.15 shows that the magnitude of both HO- and HSC torques are underestimated by the simplified model. This supports the hypothesis made in the previous section that the SoHF channel tends to amplify HO torques and vice versa, which is exactly what is lacking in this case.

Figure 3.16 shows the same run but with the HO torque-rejection model, which results in increased simulated magnitudes of both T_{HO} and T_{HSC} . Furthermore, the error is reduced from $RMS_e = 0.47\text{ cm}$ for the non-adaptive HO, to $RMS_e = 0.30\text{ cm}$ for the adaptive HO shown in the right-hand side of Figure 3.16, indicating that the adaptive HO model *achieves better tracking performance through conflicts*. In this run the experiment error was $RMS_e = 0.35\text{ cm}$, meaning that the simulated adaptive HO even out-performs the subject for this run. This may be due to the idealized assumptions governing the adaptive model.

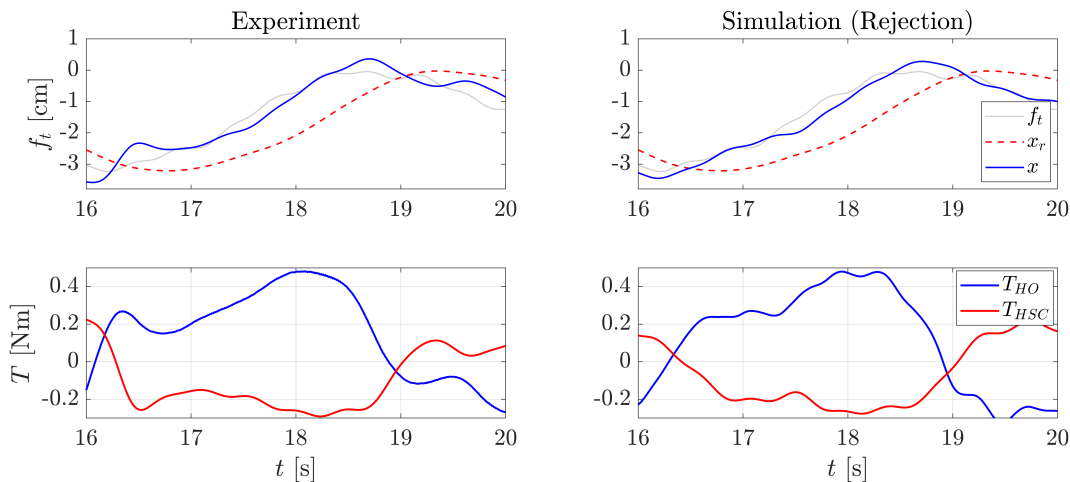


Figure 3.16: Simulated versus experiment torque for an adaptive HO model, $\tau_{HSC} = 0.0s$

Figure 3.17 shows the magnitude of both HO and HSC torques as a function of rejection gain, K_r . For no rejection ($K_r = 0$) both the HO and HSC torques are relatively small in magnitude. As rejection increases, the amplification effect of the SoHF can be clearly seen as both torques increase in magnitude (but opposite in direction). Further increasing the magnitude ($|K_r| > 8$) yielded negligible differences, indicating that

there is a limit to the amount of rejection/adaptation that can be done. This makes intuitive sense due to the saturation limiter, which ensures that rejection torques will, at most, fully cancel out the HSC torques.




Figure 3.17: Conflict as a function of varying rejection, $\tau = 0.0s$

Figure 3.18 shows the results of setting $K_r = -4$ and computing the VAF between the simulated values of x and experimental data for Span’s 10 subjects. For the majority of values of τ_{HSC} , the rejection model yields better or similar VAF scores as the Span’s baseline simulations. The exception to this rule is $\tau_{HSC} = 1.2s$, where using the rejection model results in a worse fit than the baseline (no rejection). **This result suggests that HO’s use a different strategy when adapting to a controller with high-look times relative to a controller with low look-ahead times.** Finally, Figure 3.19 shows that the adaptive, conflict-rejecting HO model fits the experiment data better for extreme values of τ_{HSC} than the baseline, validating its use in designing an adaptive HSC.




Figure 3.18: VAF for rejection model

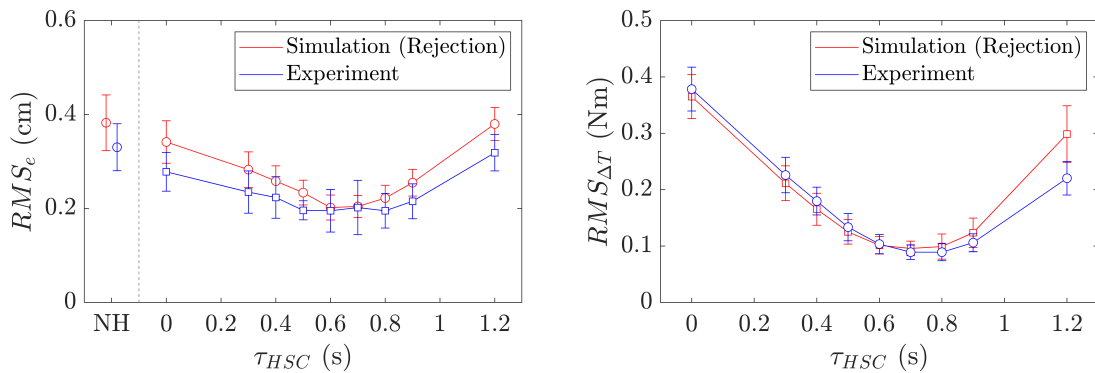


Figure 3.19: Error and conflict for rejection model and varying τ_{HSC}

Main finding of Section 3.5

HO rejection forces are modelled based on the assumption that if a HO notices that the current HSC input runs counter to its own, then it will undergo additional effort to reject this input. This is modelled by adding a 'rejection' loop that runs parallel to Van der El's model, which outputs a rejection torque T_R by comparing the HO's current nominal torque (T_{HO^*}) and the current HSC torque (T_{HSC}). If the HSC torque deviates too much from the HO's expectation, the HO will neutralise it by superimposing the rejection torque on top of its nominal output (T_{HO^*}), hereby cancelling the effects of the HSC.

3.6. Modelling Time-varying HO Behaviour

The final and ultimate goal of this report is to propose and test a structure for an adaptive HSC. To this end a test scenario is created in which the preview time is varied during a simulation run, resulting in time-variant HO behaviour. Figure 3.20 shows a timeline of the proposed scenario, starting off with a run-in time of $8s$. This is followed by the first condition, which will be either full preview ($1.0s$) or limited preview time ($0.1s$). At $t = 38s$ the condition will switch, going from either $1.0s$ to $0.1s$ or vice versa. A period of transient behaviour is expected, after which another $30s$ period allows for system identification during the second condition for real-life experiment runs.

The HO model used during simulations is identical to the one previously described in Section 2.1, with the exception of look-ahead time. To model time-varying HO behaviour in light of the varying preview-time, the HO is expected to change its look-ahead time from $\tau_{HO} = 0.6s$ to $\tau_{HO} = 0.1s$ when the preview time is reduced to $\tau_p = 0.1s$ (see Section 2.1.3). Initially this change will be modelled to be instantaneous, hereby offering the novel controller designs the best chance of recognising and adapting to (abrupt) changes in HO behaviour. Later on in the chapter the effects of transient behaviour will be evaluated. All simulations assume SI CE dynamics.



Figure 3.20: Test case

The remainder of this chapter uses the above test case and the target signal given by Table 2.2 to evaluate three distinct adaptive haptic shared controllers. Note that simulations in this chapter assume HO rejection of HSC torques, modelled according to Section 3.5 with $K_r = -4$ and $K_n = 8$ (remnant gain). The effects of varying K_r on A-HSC behaviour are evaluated at the end of the chapter.

3.7. Controller Design: Inertial A-HSC

3.7.1. Inertial A-HSC Structure

The structure of the first of three proposed A-HSC's is discussed in this section: the *inertial A-HSC*, shown in Figure 3.21. Note that the HO given in Figure 3.21 is modelled to be time-variant according to the test case given in Section 3.6. The inertial A-HSC consists of three core parts:

1. Adaptive controller: determines the optimal look-ahead time for the HSC
2. Phase shifter: modulates the target function depending on the adaptive look-ahead time
3. FDC-HSC: a conventional Four-Design-Choices HSC as described in Section 2.2

The adaptive part of the controller ('Adaptive Controller' in Figure 3.21) determines the look-ahead time τ_a for the traditional FDC-HSC controller ('HSC' in Figure 3.21), taking as an input the difference between the HO and HSC torques (ΔT) in order to measure conflict, also leveraging the target f_t and state x . Thereafter, the calculated adaptive look ahead time τ_a is determined and used to phase shift the target signal before it is passed to the FDC-HSC.

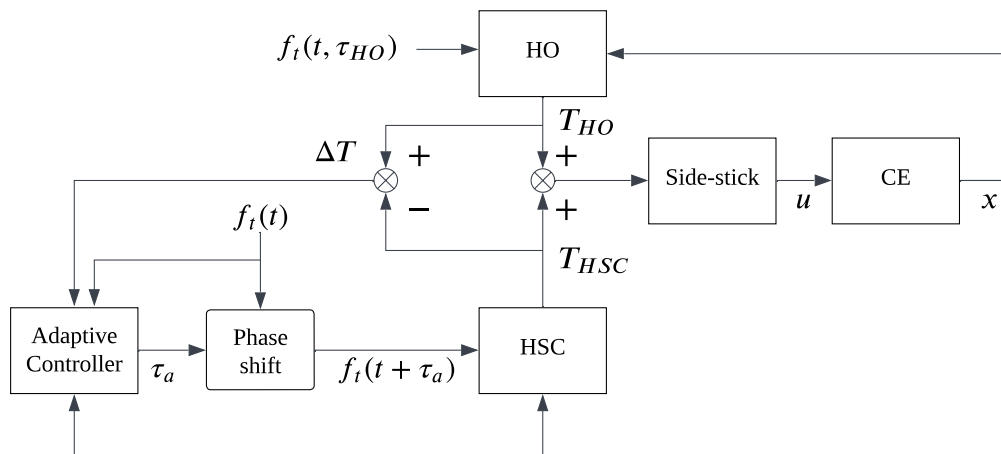


Figure 3.21: Adaptive controller structure

The most difficult part of the design process is determining how to determine the optimal look-ahead time τ_a (the inside of the 'Adaptive Controller' block), which relates to research questions 1 and 2:

Research Question 1

Which metrics should be used to trigger or drive adaptation of the haptic shared controller?

Research Question 2

How should the system adapt itself in relationship to the chosen metrics?

To answer these questions, Figure 3.22 shows the inner workings of the 'Adaptive Controller' block in Figure 3.21. It consists of two main parts: 1) the calculation of a cost value c which is to be minimised and 2) a memory buffer and decision policy to decide how τ_a should be adapted.




Figure 3.22: Adaptive controller cost and policy

Exactly how the cost is best defined is yet to be fully investigated, since it raises a deeper question about if the controller should adapt itself to reduce conflicts *or* error. For example, for limited preview times it is known that the HO's tracking performance will deteriorate as it is not able to generate the same amount of phase lead as before, hence lagging the target. In this case the controller can be configured to *also* lag the target by aiming to reduce the maximum amount of conflict torques ΔT . Alternatively, the HSC can be set to (additionally) reduce the error. For the purposes of this project, the cost value is taken to only be dependent on ΔT , meaning that the error channel is inactive in Figure 3.22.

The memory buffer takes the average cost of 5-second chunks of time to average out some of the inherent stochastic conflict that is always present due to the HO remnant and disturbance. It was found that a lower chunk length resulted in the controller not being able to distinguish random conflicts from its own adaptive influence. The last two cost-chunks C_{i-1}, C_i are stored before the decision policy uses them to evaluate in which direction τ_a should adapt.

The decision policy itself simply evaluates whether the last cost-chunk C_i performs better than the previous chunk C_{i-1} . If $C_{i-1} - C_i < 0$, then the value of the look-ahead time for the next chunk will evolve in the same direction as previous look-ahead times were going in, because performance is evidently increasing. The next iteration of look-ahead time τ_{i+1} will then experience a change as described by Equation 3.9.

$$\tau_{i+1} = \tau_i + 0.68 \cdot |C_{i-1} - C_i| \cdot \text{sign}(\tau_i - \tau_{i-1}) \quad (3.9)$$

Note that the proportionality constant 0.68 was tuned heuristically to strike a balance between adaptation speed and stability. This constant can be seen as a meta-gain: the larger it is, the larger steps the controller will take during adaptation. However, similar to conventional controllers, high gains come at the cost of reduced stability and increased volatility (sensitivity to stochastic signals).

If the performance of the latest chunk is *not* sufficiently better than the one before it, i.e. $C_{i-1} - C_i > 0$, then the next look-ahead time will adapt in the *opposite* direction as before, shown in Equation 3.10 by a subtraction of the step value instead of addition.

$$\tau_{i+1} = \tau_i - 0.68 \cdot |C_{i-1} - C_i| \cdot \text{sign}(\tau_i - \tau_{i-1}) \quad (3.10)$$

The controller exhibits a certain 'momentum' or inertia when it is adapting in a direction that reduces the cost value (it keeps going until it encounters an increase in conflict), hence its name: ***inertial*** HSC.

Figure 3.23 shows the cost and chunks over time for a typical run starting at $\tau_{a0} = 0.6s$. Note how the raw cost is volatile due to the stochastic nature of the disturbances and HO remnant, in turn influencing HO-HSC conflicts which are measured by the cost value. For this reason 5-second chunks are used to accumulate an average across time, so that the influence of the adaptive controller can be inferred reliably.




Figure 3.23: Inertial A-HSC example run 1

After the fog 'hits' at $t = 38$ the simulated HO reduces its look-ahead time to $\tau = 0.1s$. The effect of this is a sharp increase in raw cost, as there are now more conflicts. In turn, this increase is measured by the subsequent chunks, driving the adaptation policy to change τ_a .

The primary disadvantage of the inertial A-HSC is the fact that at the time of fog, the controller does not yet know which direction is 'correct'. In the case of Figure 3.23, the adaptive look-ahead time first increases

(between $t = 45 - 55$) before the controller realises that it should have gone the other way. This can be seen by the increase in cost/conflict at $t = 50 - 55$, causing the controller to revert its adaptation direction and reduce the look-ahead time, this time adapting in the correct direction. Finally, the controller converges to $\tau_a = 0.09s$.

Figure 3.24 shows the same data for another run with a different remnant seed. In this case it can be seen that the same initial spike in cost occurs after the fog hits. However, for this particular run the controller happens to immediately adapt in the correct direction (lowering τ_a). Therefore, only one spike in cost value is present in Figure 3.24, as opposed to two in Figure 3.23 (the second one is due to the controller's mistake).




Figure 3.24: Inertial A-HSC example run 2

Finally, notice how the chunk value always lags behind the raw cost in both Figure 3.23 and Figure 3.24. This is because the chunk buffer needs time to fill-up and average the previous 5 seconds of data, hereby negating the stochastic effects of the raw cost/conflict. The fact that each chunk lags reality severely limits the speed of adaptation and essentially causes the A-HSC to be 5 seconds late at all times. Especially if the controller happens to be going in the 'wrong' direction, this poses a serious limitation.

3.7.2. Inertial A-HSC Results

The results of the inertial A-HSC structure are given in Figure 3.25, showing the averaged results of 100 runs for SI CE dynamics with randomly selected target phases and remnant seeds, initialised with $\tau_a = 0.6$. The target signal is given by Table 2.2. Preview time is suddenly limited at $t = 38$, the corresponding change in HO behaviour is simulated by changing τ_{HO} from $0.6s$ to $0.1s$.

Figure 3.26 shows that initially all runs converge around a mean of $\tau_a = 0.68$ with a standard deviation of $\sigma_{38} = 0.13s$. After the fog hits the distribution diverges, increasing the standard deviation to $\sigma_{54} = 0.30s$. Finally, at $t = 199s$ the runs converge to a new mean of $\tau_a = 0.26s$ with $\sigma_{199} = 0.14s$. The divergence of simulation runs at $54s$ is indicative of the controller being driven to adapt due to an increase in conflict. When HO behaviour changes the controller is initially unaware of which direction it should adapt τ_a , hence it will take a relatively large step in whatever direction τ_a happened to be evolving in before the fog. After this period of divergence, at $t = 199s$ the simulation runs converge to an average look-ahead time of $\tau_a = 0.26s$, which is closer to the modelled HO's 'fogged' look-ahead time of $\tau_f = 0.1s$.

Table 3.3: Inertial A-HSC mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.68	0.61	0.26
Standard dev. τ_a (s)	0.13	0.30	0.14




Figure 3.25: Inertial A-HSC look-ahead time, N=100




Figure 3.26: Inertial A-HSC histogram, N=100

Main finding of Section 3.7

The inertial A-HSC uses a memory buffer and decision policy to adapt the look-ahead time of an FDC-HSC controller. The memory buffer averages the RMS value of 5-second chunks of conflict (cost) to separate the influence of the controller from random disturbance functions or remnant. If the values of subsequent chunks indicate that conflict is reducing, then the controller continues to adapt in whatever direction it was going in. Conversely, if subsequent chunks indicate an increase in conflict, then the policy is to reverse the direction of adaptation.

3.8. Model-Fitting Adaptive HSC

3.8.1. Model-Fitting A-HSC Structure

Figure 3.27 shows the structure of the second A-HSC named **Model-Fitting** (MF). The philosophy behind this controller is to run multiple HSC's in the background and select the one that fits best.

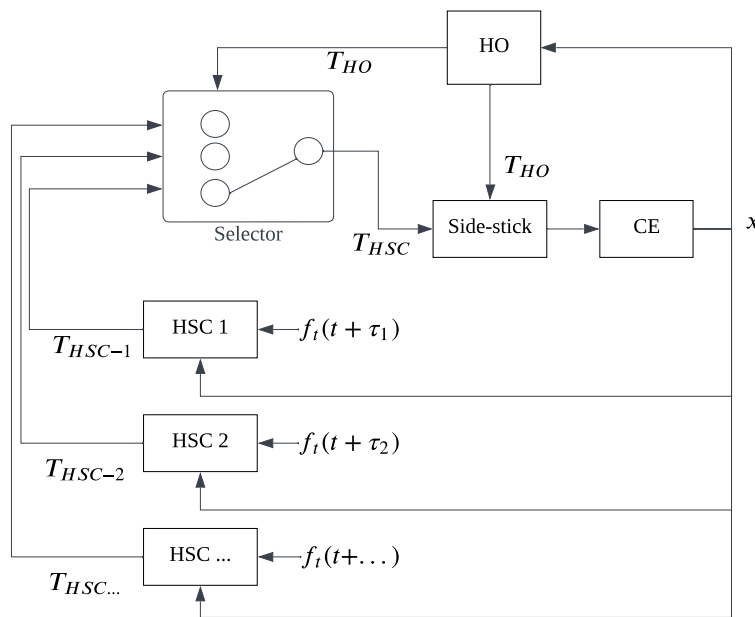


Figure 3.27: Model-fitting controller structure

Even though only one of the HSC's can be in the loop together with the HO at any given time, it is still possible to determine what the other controllers proposed torques *would have been* to an extent. Each controller receives two inputs:

1. The target function shifted according to the controller's unique look-ahead time
2. The current state x , influenced by the HO and the current active HSC

Every 5 seconds the selector compares the last 5 seconds of T_{HSC-i} with T_{HO} for each HSC model. Based on the amount of hypothetical conflict between T_{HO} and each controller, the selector then chooses the HSC that would have been closest to the HO in terms of least amount of conflict torques.

The limitation of this design is that it does not account for the fact that the HO is likely adapting its torque depending on the HSC that is currently active, which is something the inactive HSC's cannot 'see' or account for. This makes selecting a HSC based on past resemblance to the HO's torque a questionable method. Furthermore, each latent HSC is dependent upon the current state x , which is influenced by the HO's (adaptive) behaviour, drawing into question the validity the inactive HSC's latent torques (since they are a function of x). This may be the reason that during initial testing of the controller, the selector occasionally got 'confused' and switched between HSCs with extreme look-ahead times, as shown in Figure 3.28.




Figure 3.28: Model-fitting without 0.1s restriction

After implementing a restriction that only allows the selector to change to a HSC with a deviation of $\Delta\tau_{HSC} = 0.1s$ at a time, the results improved as shown in Figure 3.29.




Figure 3.29: Model-fitting with 0.1s restriction

Another possible remedy for the aforementioned limitations is to simply select the HSC whose proposed torque is smallest in terms of RMS value for the past 5 seconds. The reasoning behind this is best explained with an example. Take HSC-1 to be the active HSC with $\tau_1 = 0.0s$ and a HO with $\tau_{HO} = 0.6s$. Due to the low look-ahead time of the active HSC, the HO will likely use conflict forces to rectify the poor performance of the active HSC and ensure that x follows the target, despite HSC-1 having a sub-optimal look-ahead time of 0.1s.

Due to the SoHF path, inactive HSC's with extreme look-ahead times (e.g. $\tau_{HSC} = 1.2s$) will propose to 'correct' the current state (which is following the target due to HO rejection of HSC-1) and 'pull' it towards their own reference state. The hypothesis is that the latent HSC with $\tau_{HSC} = 0.6$ will propose a *minimum* level of adjustment, since the current state is *already* aligned with its own reference. Therefore, this HSC ($\tau_{HSC} = 0.6$) is expected to be chosen as the prime candidate. Analogous reasoning can be applied to the situation where the HO is 'fogged' and forced to use a lower look-ahead time.

3.8.2. Model-Fitting Results

Figure 3.30, Figure 3.31, and Table 3.4 show the results for a model-fitting A-HSC that selects based on a minimal difference in torque between latent HSCs and the HO (MF-1).




Figure 3.30: MF-1 look-ahead time, N=100




Figure 3.31: MF-1 histogram, N=100

Table 3.4: MF-1 mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.74	0.63	0.52
Standard dev. τ_a (s)	0.11	0.13	0.14

Figure 3.32, Figure 3.33, and Table 3.5 show the results for a model-fitting A-HSC structure that selects based on which HSC proposes the minimum RMS of torque, *without* explicitly comparing with HO torques. This method is denoted MF-2. The results are notably different from MF-1, with the means of the MF-2 distribution being significantly lower than those of MF-1.



Figure 3.32: MF-2 look-ahead time, N=100




Figure 3.33: MF-2 histogram, N=100

Table 3.5: MF-2 mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.38	0.22	0.17
Standard dev. τ_a (s)	0.11	0.11	0.09

Main finding of Section 3.8

The model-fitting A-HSC uses multiple HSCs running in the background together with a selection algorithm. Every 5-seconds the selection algorithm may choose to 'activate' a latent HSCs based on one of the two following criteria:

- **MF-1:** The difference between the (proposed) HSC and HO torques of the latent HSC is smaller than the current difference
- **MF-2:** The RMS value of the (proposed) HSC torque of the latent HSC is smaller than the current RMS value of the active HSC torque

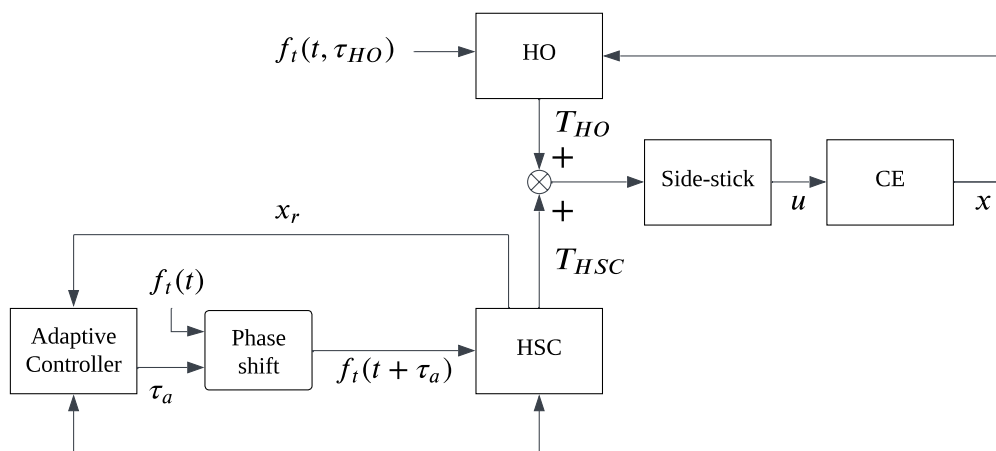
The disadvantage of model-fitting lies in the fact that only one model can be active at any given time. It is therefore difficult to predict which inactive model will outperform the current one, given the fact that inactive models cannot truly be evaluated unless they are made active together with the HO.

3.9. Cross-Correlation Adaptive HSC

3.9.1. Cross-Correlation A-HSC Structure

The structure of the final and most promising A-HSC design is given by Figure 3.34. The cross-correlation haptic shared controller (XC-AHSC) is similar in design to the inertial A-HSC as it also seeks to determine the optimal look-ahead time before passing this information to a conventional FDC-HSC. However, the method it uses to determine the optimal look-ahead time is fundamentally different from the inertial A-HSC.

Figure 3.35 shows the contents of the 'Adaptive Controller' block in Figure 3.34 and reveals the mechanism by which the A-HSC adjusts the look-ahead time of the FDC-HSC. The adaptation policy is based on the relationship between the reference state and HO-HSC conflicts. Section 3.4 concluded that if the actual state and the reference state are not aligned then it is highly likely that the HSC is not in agreement with the HO's actions. To measure sustained mismatch between x_r and x in terms of phase shift, the cross-correlation between the two is computed over a buffer as shown in Figure 3.35.

**Figure 3.34:** Cross-correlation A-HSC structure

By evaluating at which index the cross-correlation function reaches its maximum, the amount of lead or lag between the state and reference can be determined. How much the reference state is leading/lagging the actual state (measured by i_{max}) determines to what extent the current look-ahead time τ_a should be adapted. To control the speed of adaptation, the gain K_a can be tuned. If the reference is consistently lagging the actual state, by implication this must mean that the HO is 'pulling' the state and causing lead (and vice versa), resulting in conflict. By measuring this misalignment through cross-correlation and using it to drive adaptation, the root cause of conflicts is addressed head-on without circumvention. The buffer size is set to $t = 9.00s$ with an overlap of $t = 8.99s$, K_a is set to $1.25 \cdot 10^{-5}$.




Figure 3.35: Cross-correlation adaptive controller block

3.9.2. Cross-Correlation Results

The results of the XC-AHSC are given by Figure 3.36, Figure 3.37, and Table 3.6. Out of the three A-HSC designs the XC-AHSC exhibits the most robust behaviour, indicated by the low spread of data. This is likely due to the simplicity of the design and its capacity to effectively address the root cause of conflicts. Furthermore, the XC-AHSC is the only adaptive controller that does not require discontinuous jumps between distinct look-ahead times. By selecting a sufficiently large buffer overlap the optimal look-ahead time can be updated at arbitrarily small intervals (in this case, every $0.01s$). The ability to smoothly shift between look-ahead times may also increase the likelihood that the controller will be accepted by HO's during an experiment (compared to a discontinuous A-HSC).




Figure 3.36: X-corr look-ahead time, N=100




Figure 3.37: X-corr histogram, N=100

Table 3.6: Cross-correlation mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.78	0.29	0.30
Standard dev. τ_a (s)	0.03	0.13	0.04

Main finding of Section 3.9

The XC-AHSC's decision policy is based on the fact that almost all conflicts are rooted in a misaligned reference state. Using the index of the maximum value of the cross-correlation function of x and x_r , the XC-HSC updates the current look-ahead time as to reduce the phase lead or lag between the reference state and actual state. Preliminary results show the XC-AHSC to be robust against disturbances and more stable than the other controllers.

3.10. Comparing A-HSC Designs

Table 3.7 and Figure 3.38 show the combined means and standard deviations for all A-HSCs. In terms of spread of look-ahead data, X-corr has a similar narrowness of distribution as MF-2. However, X-corr has the added benefit of being able to adapt to a wider range of look-ahead times with simulations showing adaptation from an initial convergence of $\mu_\tau = 0.78s$ to $\mu_\tau = 0.30$ ($\Delta_\tau = 0.48s$). This is a significant improvement over MF-2's adaptation range, which is $\Delta_\tau = 0.38s - 0.17s = 0.21s$. The second-best performer in terms of adaptation range is the inertial-AHSC, going from $0.68s$ to $0.26s$ ($\Delta_\tau = 0.42s$). However, the inertial-AHSC's volatile nature (large spread) renders it a risky choice for an experiment.

Table 3.7: Adaptive look-ahead time mean and standard distribution per controller

	X-corr			MF-2			MF-1			Inertial		
	38s	54s	199s	38s	54s	199s	38s	54s	199s	38s	54s	199s
μ_τ	0.78	0.29	0.30	0.38	0.22	0.17	0.74	0.63	0.52	0.68	0.61	0.26
σ_τ	0.03	0.13	0.04	0.11	0.11	0.09	0.11	0.13	0.14	0.13	0.30	0.14




Figure 3.38: Adaptive look-ahead time mean and standard distribution per controller

Figure 3.39 shows the averaged values of all runs for three A-HSC designs ($N=3 \times 100$) combined in one graph. MF-2 is not considered due to its tendency to converge to short look-ahead times (even in situations with high preview time), which would likely result in conflicts [6]. Figure 3.40 shows the same data with an emphasis on convergence, showing how the look-ahead time of each controller remains bounded.




Figure 3.39: Comparing averaged look-ahead times for A-HSCs ($N=3 \times 100$)




Figure 3.40: Convergence of A-HSCs ($N=3 \times 100$)

Table 3.8 presents the data shown in Figure 3.40 in tabulated format. The settling time refers to the start of the convergence regions in Figure 3.40, i.e. the time beyond which τ_a remains bounded between the upper and lower limits specified in Table 3.8. The inertial A-HSC performs the best with respect to convergence range, however, its settling time (115s) is more than double that of the XC-AHSC, which has the fastest settling time of 53.27s. Note that the *convergence range* of the mean values is not to be confused with the *standard deviation* of aggregated realizations, as given by σ_τ in Table 3.7.

Table 3.8: Convergence of A-HSCs ($N=3 \times 100$)

Convergence metric	Inertial	MF-1	XCorr
τ_a lower bound (s)	0.2559	0.4750	0.2608
τ_a upper bound (s)	0.2937	0.5750	0.3036
τ_a range (Δ) (s)	0.038	0.100	0.043
Settling time (s)	115.00	65.00	53.27

Main finding of Section 3.10

The XC-AHSC performs better or equal to other AHSCs according to the following metrics:

1. **Adaptation speed:** how quickly the controller is capable of converging after fog.
2. **Robustness:** the spread of data, i.e. how resilient is the controller against false adaptations (remnant or disturbance).
3. **Adaptation range:** the range of look-ahead times that the controller can adapt to.

For these reasons, only the XC-AHSC will be considered going forward in the project.

3.11. Cross-Correlation A-HSC: Additional Findings

3.11.1. Effects of Transient Behaviour on XC-AHSC Adaptation

This section evaluates the effects of using various methods to model the HO's transient behaviour. Thus far simulations have been performed under the assumption that the HO discontinuously changes its look-ahead time from $0.6s$ to $0.1s$. In reality this change will be continuous rather than instantaneous, since it is likely the HO will remember certain parts of the previewed signal just before they disappear due to the limited preview time.

To model time-varying behaviour of the HO more in a more realistic manner, two sigmoid functions of varying speeds are used. These functions are used to shift between the original HO model ($\tau_{HO} = 0.6$) and the adapted HO model ($\tau_{HO} = 0.1$) by changing each model's contribution to the output. This is given by Equation 3.11, where S_1 and S_2 are complementary sigmoid functions such that $S_1(t) + S_2(t) = 1 \quad \forall t$. Table 3.9 shows the parameters used to generate both fast and slow sigmoid functions, with *duration* measuring the time it takes the function to go from 0.1 to 0.9. The results are given in Figure 3.41, showing the value of $S_2(t)$ in addition to how the average adaptive look-ahead time is influenced by the three types of transient functions ($N=3 \times 100$).

$$T_{ho} = T_{0.6} \cdot S_1(t) + T_{0.1} \cdot S_2(t) = T_{0.6} \left(\frac{1}{1 + e^{a(t-c)}} \right) + T_{0.1} \left(\frac{1}{1 + e^{-a(t-c)}} \right) \quad (3.11)$$

Table 3.9: Sigmoid parameters

	c (s)	a (-)	Duration (s)
Fast	40	3.3	1.32
Slow	43	1.0	4.39

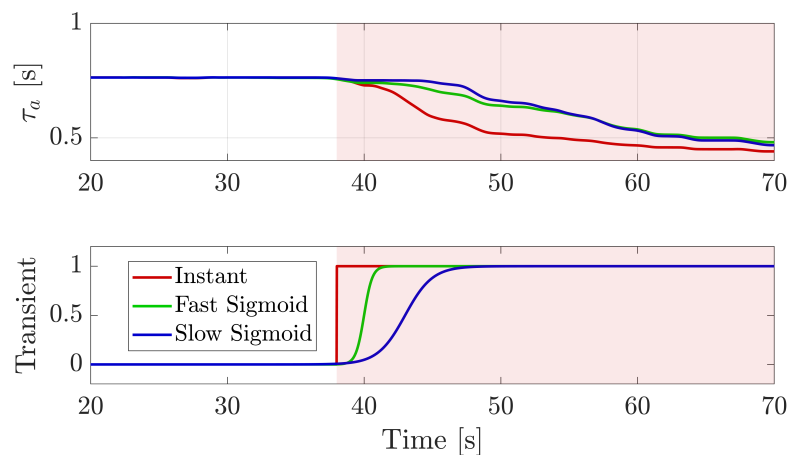


Figure 3.41: Effects of varying transient behaviour, $N=3 \times 100$

As expected, it takes slow sigmoid the longest time to adapt. For an instant transient function it takes on average $6.7s$ for the look-ahead time to adapt from $0.75s$ to $0.6s$. For the fast and slow sigmoids it takes $17.2s$ and $17.5s$, respectively.

Figure 3.42, Figure 3.43, and Table 3.10 show more detailed results for the slow sigmoid compared to instantaneous HO adaptation. It can be seen that the initial rate of adaptation for the slow sigmoid is more shallow than for the instantaneous case, with the sigmoid mean look-ahead time taking on average an additional 48 seconds to converge to its final value (after some stagnation at $\tau = 0.4s$).

Table 3.10: XC-AHSC Slow sigmoid mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.78	0.59	0.30
Standard dev. τ_a (s)	0.03	0.09	0.04

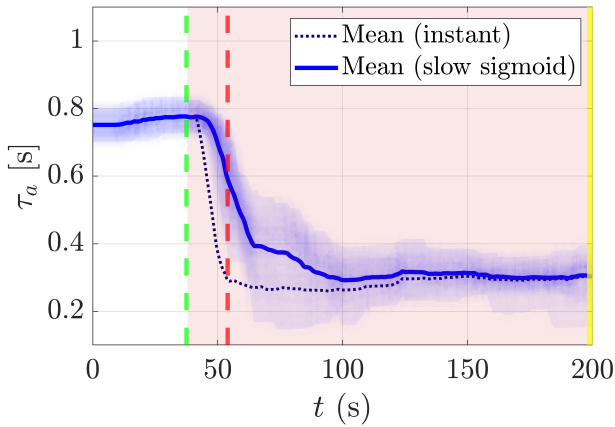


Figure 3.42: XC-AHSC look-ahead time for slow sigmoid, N=100




Figure 3.43: XC-AHSC histogram for slow sigmoid, N=100

3.11.2. XC-AHSC Adaptation for Varying HO Levels of HO Rejection

This section briefly evaluates the effects of varying levels of HO adaptation on the average look-ahead times. Figure 3.44 shows the results for 5x100 realizations performed for varying values of K_r between $[0, -8]$ with increments of -2 . It can be seen that when HO rejection of HSC is disabled ($K_r = 0$), look-ahead time adaptation is almost non-existent. However, **when the HO model is configured to reject HSC forces that are sufficiently different from their own intended actions, look-ahead time adaptation increases significantly.** This finding demonstrates the necessity of modelling HO rejection of HSC forces and raises the following question: to what extent will HO's reject HSC forces under limited preview time in real life? Answering this question is one of the goals of this project's HO-in-the-loop experiment.

The reason the controller is more prone to adaptation for higher levels of rejection is because the HO exerts an increasing level of influence on the state x for higher values of K_r . When configured to reject HSC forces, a fogged HO will not accept a reference state that follows the target, because in this case the HSC's torques will not be in line with the HO's inputs. Therefore, adaptation is required to align x_r with the actions of the HO. However, if the HO is configured *not* to reject HSC forces, then evidently it has no issue with accepting a look-ahead time that is significantly higher than its own. Whether or not HO's choose to reject HSC forces may be dependent on the task instructions, trust, and personal preference.




Figure 3.44: XC-AHSC look-ahead time for varying K_r (N=5x100)

3.11.3. Effects of Fog Direction on XC-AHSC Adaptation

Figure 3.45 shows the same experiment performed in the opposite direction, going from limited preview time to full preview after $t = 60s$. Contrary to Figure 3.44, rejection seems to have a negligible effect on adaptation if the look-ahead time is increasing. Depending on the amount of rejection (as shown in Figure 3.44), **this finding suggests that the controller and HO may be more prone to adaptation in the direction of small-to-large look-ahead times than vice versa.**




Figure 3.45: XC-AHSC look-ahead time for varying K_r , reversed (N=5x100)

Figure 3.46, Figure 3.47, and Table 3.11 show the results of 100 realizations for when the HO initially starts with a limited preview time ($\tau_{HO} = 0.1s$) and instantaneously shifts to a higher look-ahead time at $t = 38s$, essentially reversing the previous simulations. Table 3.11 shows that the adaptation range ($\Delta\tau = 0.80 - 0.33 = 0.47s$) is similar to the original case (Table 3.6: $\Delta\tau = 0.48s$). Furthermore, the spread of the distributions is narrower, having an average standard deviation of $\sigma_\tau = 0.04s$, compared to the original case which has an average of $\sigma_\tau = 0.07s$. These findings may be related to results of Section 3.5, which suggest that HO's may use a different strategy to reject HSC forces depending on if the look-ahead time is higher or lower than their own.

Table 3.11: XC-AHSC reversed mean and standard deviation

Time	38s	54s	199s
Mean τ_a (s)	0.33	0.49	0.80
Standard dev. τ_a (s)	0.03	0.06	0.04




Figure 3.46: XC-AHSC reversed adaptive look-ahead time, N=100




Figure 3.47: XC-AHSC reversed histogram, N=100

3.11.4. Effects of Buffer Size on XC-AHSC Adaptation

The top row of Figure 3.48 shows the evolution of adaptive look-ahead time for a single run without remnant, with a buffer length of 5 seconds and fog at $t = 48s$. The corresponding phase delay (ϕ) between the reference state and actual state is given by the second row, as determined through the index of the maximum value of the cross-correlation function. The last row shows the reference and actual state.

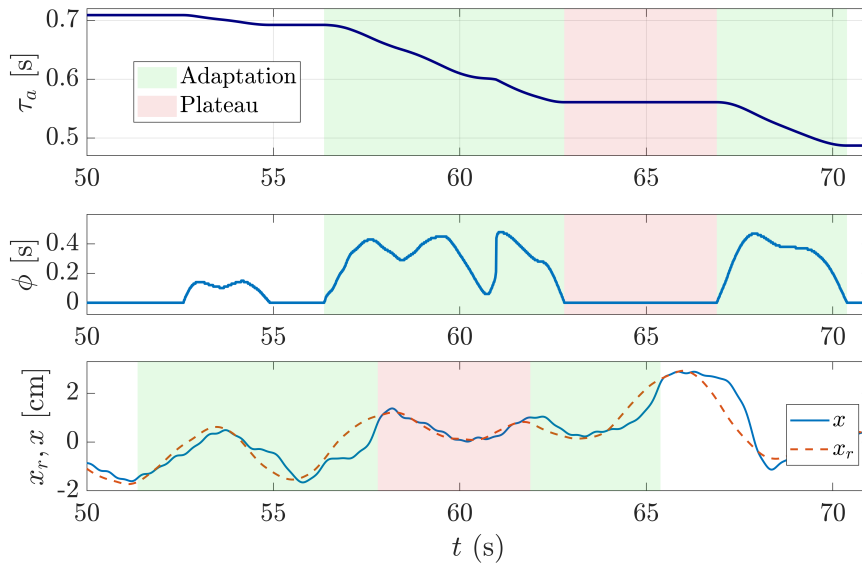


Figure 3.48: Reference state, phase lag and adaptive look-ahead time during plateau

Figure 3.48 shows that after initial excitement between $t = 56 - 63s$, adaptation briefly stops before continuing at $t = 67s$. This plateau coincides with zero phase delay, as adaptation is driven by a phase lead or phase lag. The reason for the plateau is due to the characteristics of the target signal. The adaptive algorithm is most effective during large excursions of the target signal, i.e. when the HSC is able to follow the target, but with a certain amount of measurable lag relative to the reference. This occurs in Figure 3.48 between $t = 63 - 66s$ (bottom row), for example, where the reference leads the state due to a relatively high look-ahead time ($\tau_{HO} = 0.1s$).

However, when the target signal is relatively constant (or 'flat'), the phase lead or phase lag between x and x_r is minimal. An example of this can be seen highlighted in red in the bottom row of Figure 3.48. As a result of overlap between x and x_r , adaptation stagnates. It is important to bear in mind that the cross-correlation at a given point in time refers to the last 5s chunk of data in the event of a buffer size equal to 5 seconds. This is why the measured phase lag or phase lead is delayed relative to the actual phase difference, hence why the coloured windows are shifted in the last row as seen in Figure 3.48.

Figure 3.49 shows how increasing the buffer length for the same run (no remnant, same target signal) results in smoother adaptive behaviour. This is because due to the fact that larger time-frames are being considered, it is less likely that the target will be 'flat' throughout the entirety of a chunk of data.




Figure 3.49: Effect of buffer size on adaptive look-ahead time

Main findings of Section 3.11

This section evaluated the effects of the following factors on the XC-AHSC controller:

1. **Transient functions:** two sigmoid curves were used to gradually adapt HO behaviour. It was found that by using a sigmoid function, the start of adaptation time was delayed by approximately the length of the sigmoid curve.
2. **Varying HO rejection:** a clear correlation was established between the amount of torque-rejection and look-ahead time adaptation. *In the case of high-to-low look-ahead adaption*, the more HSC torques are rejected by the HO, the more the controller is prone to adaptation. This effect was not found in the opposite direction (low-to-high look-ahead times), suggesting an inherent asymmetry in HO-HSC behaviour.
3. **Fog direction:** by configuring the HO to initially have a limited preview time that increased halfway through the runs, the effect of reversing the fog direction was evaluated. The spread of adaptive look-ahead times was found to be narrower for the reversed situation. A similar adaptation range was observed relative to the original case.
4. **Buffer size:** for small buffer sizes, the controller is prone to intermittent periods of stagnation (plateau-ing). This is due to characteristics of the target signal, where medium-to-large excursions of the state are required to measure a phase difference between the reference state. By increasing the buffer size adaptation becomes less prone to stagnation, since phase differences are evaluated over a longer time-span with more data.

4

Conclusion

The purpose of this report was to propose an adaptive haptic shared controller design to make progress towards the following research objective:

Research Objective

To design, simulate, and validate an adaptive haptic shared controller that adjusts itself in real-time to reduce conflicts between a human operator (HO) in a preview tracking task, without undermining control performance.

To this end, simulations are performed using three novel adaptive haptic shared controllers: inertial, model-fitting, and cross-correlation (X-corr). Each controller is tested in a simulated scenario with time-varying human behaviour to evaluate adaptation performance. Thereafter, the results for each controller are compared and the X-corr controller is selected based on robustness, adaptation range, and adaptation speed.

The first research question pertains to the cost metric:

Research Question 1

Which metrics should be used to trigger or drive adaptation of the haptic shared controller?

Answer to research question 1: the cross-correlation between the internal reference state of the HSC and the actual state x is used to drive adaptation of the look-ahead time.

The second research question pertains to the decision policy:

Research Question 2

How should the system adapt itself in relationship to the chosen metrics?

Answer to research question 2: A phase shift between x_r and x indicates that the HO and HSC are not in agreement, since in this case the HO is likely correcting for sub-optimal performance of the HSC. By evaluating the index of the maximum value of the cross-correlation function, the phase shift between the two signals can be determined. If x_r is found to be leading or lagging x , a step in look-ahead time $\Delta\tau$ is calculated and added to the current look-ahead time.

Simulation results averaged over 100 realizations show that for a instant change in the HO's look-ahead time ($0.6s - 0.1s$), the X-corr controller is capable of adapting from $\tau_a = 0.78s$ to $\tau_a = 0.29s$ during a period of $16s$. After $199s$, on average the controller adapts its look-ahead time from $\tau_a = 0.78s$ to $\tau_a = 0.30s$. The average standard deviation throughout simulation runs is $\sigma = 0.07s$.


References

- [1] S. M. Petermeijer, D. A. Abbink, M. Mulder, and J. C. F. de Winter. "The effect of haptic support systems on driver performance: A literature survey". In: *IEEE Transactions on Haptics* 8.4 (2015), pp. 467–479.
- [2] T.J. Gordon and M. Lidberg. "Automated driving and autonomous functions on road vehicles". In: *Vehicle System Dynamics* 53.7 (2015), pp. 958–994.
- [3] L. Bainbridge. "Ironies of automation". In: *Automatica* 19.6 (1983), pp. 775–779.
- [4] B. A. Forsyth and K. E. MacLean. "Predictive haptic guidance: Intelligent user assistance for the control of dynamic tasks". In: *IEEE Transactions on Visualization and Computer Graphics* 12.1 (2005), pp. 103–113.
- [5] F. Flemisch, J. Kelsch, C. Löper, A. Schieben, J. Schindler, and M. Heesen. *Cooperative Control and Active Interfaces for Vehicle Assistance and Automation*. 2008.
- [6] J. Span, D. M. Pool, M. M. van Paassen, and M. Mulder. *Effects of Look-Ahead Time in a Haptic Shared Controller for Preview Tracking*. Unpublished MSc thesis. 2021.
- [7] M. Mulder, D.M. Pool, D.A. Abbink, E.R. Boer, P.M. Zaal, F.M. Drop, K. Van Der El, and M.M. Van Paassen. "Manual Control Cybernetics: State-of-the-Art and Current Trends". In: *IEEE Transactions on Human-Machine Systems* 48.5 (Oct. 2018), pp. 468–485.
- [8] D. A. Abbink, M. Mulder, and E.R. Boer. "Haptic shared control: Smoothly Shifting Control Authority?". In: *Cognition, Technology and Work* 12.1 (Nov. 2011), pp. 1–10.
- [9] C. Taylor, J. Keller, R.O. Fanjoy, and F.C. Mendonca. "An exploratory study of automation errors in part 91 operations". In: *The Journal of Aviation/Aerospace Education & Research* 29.1 (2020), pp. 33, 48.
- [10] Y.P. Kwak, Y.C. Choi, and J. Choi. "Analysis between Aircraft Cockpit Automation and Human Error Related Accident Cases". In: *International Journal of Control and Automation* 11.3 (Mar. 2018), pp. 179–192.
- [11] K.J. Vicente and J. Rasmussen. "Ecological interface design: theoretical foundations". In: *IEEE Transactions on Systems, Man, and Cybernetics* 22.4 (1992), pp. 589–606.
- [12] R. Boink, M.M. Van Paassen, M. Mulder, and D. A. Abbink. "Understanding and reducing conflicts between driver and haptic shared control". In: *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (Oct. 2014), pp. 1510–1515.
- [13] K. Van Der El, S. Padmos, D.M. Pool, M. M. Van Paassen, and M. Mulder. "Effects of Preview Time in Manual Tracking Tasks". In: *IEEE Transactions on Human-Machine Systems* 48.5 (2018), pp. 486–495.
- [14] M.P. Mileusnic and G.E. Loeb. "Mathematical Models of Proprioceptors. II. Structure and Function of the Golgi Tendon Organ". In: *Journal of Neurophysiology* 96.4 (2006), pp. 1789–1802.
- [15] D. W. J. Van Der Wiel, M. M. van Paassen, M. Mulder, M. Mulder, and D. A. Abbink. "Driver Adaptation to Driving Speed and Road Width: Exploring Parameters for Designing Adaptive Haptic Shared Control". In: *2015 IEEE International Conference on Systems, Man, and Cybernetics* (2015), pp. 3060–3065.
- [16] S. Wietske, S. Barendswaard, D. Pool, R. Van Paassen, and D. Abbink. "A New Haptic Shared Controller Reducing Steering Conflicts". In: *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2018, pp. 2705–2710.

- [17] K. van der El, D.M. Pool, H. J. Damveld, M.M. van Paassen, and M. Mulder. "An Empirical Human Controller Model for Preview Tracking Tasks". In: *IEEE Transactions on Cybernetics* 46.11 (2016), pp. 2609–2621.
- [18] M. Tomizuka and D.E. Whitney. "The Preview Control Problem with Application to Man-Machine System Analysis". In: *Proceedings of the Ninth Annual Conference on Manual Control*. 1973, pp. 429–441.
- [19] D. T. McRuer and H. R. Jex. *A review of quasi-linear pilot models*. Vol. HFE-8. 3. Sept. 1967, pp. 231–249.

Part III

Appendices to Scientific Article



Computing Cross-correlation

The aim of this appendix is to describe the theoretical foundations behind the cross-correlation function ('X-corr'), as well as how it is implemented to estimate the time shift ζ_m between the state x and reference state x_r . Figure A.1 shows how the adaptive look-ahead time, τ_a , is updated by increments $\Delta\tau_a$, based on the estimated time shift ζ_m between x and x_r .




Figure A.1: Adaptive look-ahead time loop

A.1. Cross-product, Cross-covariance, and Cross-correlation

Before discussing how the cross-correlating function between x and x_r is computed, first a general definition of the cross-product function is given. For two stochastic processes \bar{f} and \bar{g} , the cross product function is defined according to Equation A.1, where $E\{\cdot\}$ refers to the expected value. The output of this function is a measure of the coherence between the two signals as a function of varying time shifts ζ . One can think of the second signal (in this case, \bar{g}) 'sliding' over the first signal according to a time shift ζ , with $R_{\bar{f}\bar{g}}(\zeta)$ measuring the similarity between the two signals depending on the time shift.

$$R_{\bar{f}\bar{g}}(\zeta) = E\{\bar{f}(t)\bar{g}(t + \zeta)\} \quad (\text{A.1})$$

If the means of both signals happen to be similarly large, then this will result in larger values of the cross product. Therefore, in order to have a more 'pure' measure of the two signals' coherence, it is useful to correct for the means of the two signals. This is achieved through the use of the cross-covariance function, as shown in Equation A.2, where $\mu_{\bar{f}}$ and $\mu_{\bar{g}}$ are the averages corresponding to \bar{f} and \bar{g} , respectively.

$$C_{\bar{f}\bar{g}}(\zeta) = E\{(\bar{f}(t) - \mu_{\bar{f}}) \cdot (\bar{g}(t + \zeta) - \mu_{\bar{g}})\} \quad (\text{A.2})$$

The final step is to normalize the covariance, such that the output of the function equals 1 when the signals are maximally correlated (and 0 if uncorrelated). To this end, the cross correlation function $K_{\bar{f}\bar{g}}$ is defined according to Equation A.3, where $\sigma_{\bar{f}}$ and $\sigma_{\bar{g}}$ are the standard deviations of \bar{f} and \bar{g} , respectively.

$$K_{\bar{f}\bar{g}}(\zeta) = E\left\{\frac{\bar{f}(t) - \mu_{\bar{f}}}{\sigma_{\bar{f}}} \cdot \frac{\bar{g}(t + \zeta) - \mu_{\bar{g}}}{\sigma_{\bar{g}}}\right\} \quad (\text{A.3})$$

Equation A.3 provides a useful tool for estimating the time delay or lead between two signals. If \bar{f} and \bar{g} are maximally correlated for a particular value of ζ , this means that the two signals share a common feature

which occurs both at t for \bar{f} , and at $t + \zeta$ for \bar{g} . If the maximum cross-correlation occurs for a positive time shift ζ , this means that the common feature occurs later in time (at $t + \zeta$) for \bar{g} than it does for \bar{f} (at t). If maximum cross-correlation occurs for a positive ζ , then \bar{g} is said to be *lagging* \bar{f} . Equivalently, one could also think of having to shift \bar{g} earlier by ζ seconds before maximum alignment between the two signals occurs. For a negative ζ , \bar{g} is said to be *leading* \bar{f} .

The same method for estimating time shift also applies for $R_{\bar{f}\bar{g}}$ and $C_{\bar{f}\bar{g}}$, as will now be explained. The cross-product and cross-covariance function are related by Equation A.4, which shows that the only difference between the two metrics lies in subtraction of a constant number (for a stationary process). Therefore, if a peak in $C_{\bar{f}\bar{g}}$ occurs at a given ζ , then $R_{\bar{f}\bar{g}}$ will also be maximised at the same value of ζ . However, the magnitude of each peak will be different (due to subtraction of the means).

$$C_{\bar{f}\bar{g}}(\zeta) = R_{\bar{f}\bar{g}}(\zeta) - \mu_{\bar{f}}\mu_{\bar{g}} \quad (\text{A.4})$$

In turn, the relationship between $C_{\bar{f}\bar{g}}$ and $K_{\bar{f}\bar{g}}$ is given by Equation A.5. This shows that $R_{\bar{f}\bar{g}}$ and $K_{\bar{f}\bar{g}}$ are related through scaling and subtraction. Therefore, any ζ resulting in the maximisation of $R_{\bar{f}\bar{g}}$ or $C_{\bar{f}\bar{g}}$ will also result in the maximisation of $K_{\bar{f}\bar{g}}$.

$$K_{\bar{f}\bar{g}}(\zeta) = \frac{C_{\bar{f}\bar{g}}(\zeta)}{\sigma_{\bar{f}}\sigma_{\bar{g}}} = \frac{R_{\bar{f}\bar{g}}(\zeta) - \mu_{\bar{f}}\mu_{\bar{g}}}{\sigma_{\bar{f}}\sigma_{\bar{g}}} = \frac{1}{\sigma_{\bar{f}}\sigma_{\bar{g}}}R_{\bar{f}\bar{g}} - \frac{\mu_{\bar{f}}\mu_{\bar{g}}}{\sigma_{\bar{f}}\sigma_{\bar{g}}} \quad (\text{A.5})$$

The conclusion is that each function (cross-product, cross-covariance, cross-correlation) will output the same measured time delay ζ_m when maximising the corresponding function. However, it is preferred to use $K_{\bar{f}\bar{g}}$, since it provides a normalized measure for cross-correlation between $[-1, 1]$.

A.2. Cross-correlation of a Finite Discrete Signal

Cross-product

To calculate an estimate for cross-correlation, first the cross-product is analysed as a first step. For real finite discrete signals F and G with N samples each, the cross-product is estimated as shown in Equation A.6, where m is the relative shift in terms of number of samples: $m = 1 - N, 2 - N, \dots, N - 2, N - 1$.

$$\hat{R}_{FG}[m] = \begin{cases} \sum_{n=1}^{N-m-1} F[n]G[n+m] & m \geq 0 \\ \sum_{n=1}^{N+m-1} F[n-m]G[n] & m < 0 \end{cases} \quad (\text{A.6})$$

In the following it will be shown that the cross-product is related to the convolution of the same two signals, which can in turn be calculated by multiplication in the frequency domain.

Convolution

The cross-product of two signals is closely related to the convolution of the same signals. The general definition of a convolution for signals \bar{f} and \bar{g} is:

$$(\bar{f} * \bar{g})(t) = \int_{-\infty}^{\infty} \bar{f}(\tau)\bar{g}(t - \tau)d\tau \quad (\text{A.7})$$

For real finite discrete signals F and G , the convolution is determined according to:

$$(F * G)[t] = \sum_{\tau} F[\tau]G[t - \tau] \quad (\text{A.8})$$

Where the summation is performed for all values of τ that result in valid indices for the given t (i.e., the indices for which both signals overlap). Swapping τ for n , for t with m results in:

$$(F * G)[m] = \sum_n F[n]G[m - n] \quad (\text{A.9})$$

Replacing n with $-n$ across the entire expression then results in Equation A.10.

$$(F * G)[m] = \sum_n F[-n]G[m+n] \quad (\text{A.10})$$

Equation A.10 shows that by flipping the discrete signal $F[n] \rightarrow F[-n]$, *before* computing the convolution, the output then becomes:

$$(F * G)[m] = \sum_n F[n]G[m+n] \quad (\text{A.11})$$

Which is the same algorithm as shown in Equation A.6, where the summation also only applies to the relevant indices n . This finding corresponds to the property that the cross-product of two signals, $f \star g$, is related to the convolution of the same signals, $f(-t) * g$, by flipping f in the time-domain before the convolution. This relationship is formalized in Equation A.12.

$$[f(t) \star g(t)](\zeta) = [f(-t) * g(t)](\zeta) \quad (\text{A.12})$$

Computing convolution in the frequency domain

To compute the convolution, use is made of the *convolution theorem*, which states that a convolution in the time domain corresponds to multiplication in the frequency domain:

$$\mathcal{F}\{f(-t) * g(t)\} = \mathcal{F}\{f(-t)\} \cdot \mathcal{F}\{g(t)\} \quad (\text{A.13})$$

By taking the inverse Fourier transform of the right-hand side of Equation A.13, an expression is found for the cross-product $f(t) \star g(t)$:

$$[f(t) \star g(t)](\zeta) = [f(-t) * g(t)](\zeta) = \mathcal{F}^{-1}(\mathcal{F}\{f(-t)\} \cdot \mathcal{F}\{g(t)\}) \quad (\text{A.14})$$

Finally, Equation A.15 shows the equivalent discrete-time algorithm for computing an estimate for the cross-product using Discrete Fourier Transforms (DFT). To account for the disjunction in overlap between two finite signals, $F[-n]$ and $G[n]$ are both zero-padded up to $2N - 1$ (corresponding to the combined length of the time-shifted signals, with only one sample of overlap). iDFT refers to the inverse DFT.

$$\hat{R}_{FG}[m] = F[n] \star G[n] = F[-n] * G[n] = \text{iDFT}\{\text{DFT}\{F[-n]\} \cdot \text{DFT}\{G[n]\}\} \quad (\text{A.15})$$

Equation A.15 shows that there are three equivalent methods for estimating the cross-product. To estimate the cross-correlation, the cross-product is normalized according to Equation A.16, such that the auto-correlations of both signals equal 1 at zero lag.

$$\hat{K}_{xx_r}[m] = \frac{\hat{R}_{xx_r}[m]}{\sqrt{\hat{R}_{xx}(0)\hat{R}_{x_r x_r}(0)}} \quad (\text{A.16})$$

Estimated time shift

Figure A.2 shows the cross-product and cross-correlation estimates of windows of the simulated state and reference state. The buffer length is $N = 700$ and the simulation time-step is $\Delta t = 0.01s$. All three methods for calculating the cross-product (given in Equation A.15) are used to estimate \hat{R}_{xx_r} , with each algorithm resulting in the same output, as expected. The last row shows the cross-correlation, with a peak occurring at index $m = +26$. This corresponds to x_r lagging x , as shown in the top row.




Figure A.2: Computing estimated cross-product/correlation of simulated x , x_r ($N = 700$)

To estimate the time shift ζ_m between x and x_r , the index for maximum \hat{K}_{xx_r} is multiplied by the time-step:

$$\zeta_m = \Delta t \cdot \underset{m}{\operatorname{argmax}}(\hat{K}_{xx_r}[m]) \quad (\text{A.17})$$

For Figure A.2 this results in $\zeta_m = 0.01 \cdot 26 = 0.26s$. This concludes how the 'X-corr' block in Figure A.1 calculates ζ_m based on the two buffers x and x_r .

XC-AHSC: Robust Protection

This appendix explains the effects of implementing a safety measure to stop the controller adapting in situations where a time-shift cannot be measured with sufficient certainty. The safety consists of setting a minimum value for the peak cross-correlation, such that adaptation only occurs when a certain minimum cross-correlation between x and x_r is established. If the peak cross-correlation is not strong enough, then the look-ahead time increment is set to $\Delta\tau_a = 0$ to prevent instability.

The decision was made to implement this protection as a result of encountering strange behaviour during real life testing of the adaptive cross-correlation design. Occasionally, the adaptive look-ahead time would increase or decrease erratically, seemingly unprovoked. This only occurred during the non-haptic condition, when the user did not apply any input. To investigate this behaviour further, Figure B.1 shows a simulated run with no inputs to the stick, setting $K_{HSC} = K_{HO} = 0$ and running the cross-correlation algorithm. Remarkably, very similar results emerge as were observed in real life.

The reason behind the erratic behaviour is related to the fact that $x = 0$ (simulations), or the state is equal to noise around a mean of zero (real life). As a result, the cross-correlation between x and x_r is expected to be approximately equal for all phase shifts, since x does not contain any distinct features. Despite x_r containing periodic elements, since x contains no temporal information and is entirely unpredictable, the correlation between the two will never be dependent on their relative phase. Since adaptation is incumbent on finding a peak correlation for a given phase shift, this effect is precisely what causes erratic look-ahead adaptation as shown in Figure B.1.




Figure B.1: Normal behaviour




Figure B.2: Erratic behaviour




Figure B.3: Protected behaviour

The solution to this problem lies in setting a minimum threshold for \hat{K}_{xx_r} , such that adaptation will only take place *if* a reliable correlation has been found. The results of applying this threshold are shown in Figure B.3, where the minimum was set at $\hat{K}_{xx_r} > 0.14 = \hat{K}_{min}$.

Figure B.4 shows a typical simulation run from pursuit to preview with transition at $t = 50s$. The second and third row show the peak cross-correlation and corresponding phase shift, respectively. The dotted line shows the threshold that was used for protection during the experiment. The last row shows the phase shift between the actual state and reference state of the controller. It can be seen that around

$t = 68s$ look-ahead adaptation briefly stagnates, even though there is a clear phase shift between x and x_r , correctly measured by the algorithm to be approximately $t = 0.48s$. Just as the phase shift peaks, adaptation stops because the minimum threshold for peak cross-correlation is not met. This is an example of the protection inadvertently hindering adaptation.




Figure B.4: Inadvertent plateauing during adaptation because of \hat{K}_{min} protection

Figure B.5 shows the effect of this safety measure on the non-haptic look-ahead adaptation for $2 \times 5 \times 16$ realizations. For the pursuit to preview case, the minimum threshold has a significant impact on the speed of adaptation. For low values of \hat{K}_{min} the controller is prone to overshooting before converging to the final look-ahead time.

Results show that $\hat{K}_{xx_r} > 0.26$ strikes a balance between speed and robustness. For preview to pursuit there is little risk of overshooting, making protection redundant. This is likely the reason why look-ahead adaptation for preview to pursuit in the non-haptic case was less aggressive than expected during the experiment. To get the full picture on what the optimal settings for the controller would be, ideally \hat{K}_{min} should be considered in conjunction with varying adaptation gains K_a . That way, the optimal balance between performance and stability could be found.




Figure B.5: Effect of \hat{K}_{min} on average look-ahead time, $N = 2 \times 5 \times 16$, non-haptic

Finally, Figure B.6 shows the results of varying the threshold, together with an active HSC in the loop. The results show that there is no need for the protection as long as both the HSC and HO are active, since there is no risk of overshooting. In hindsight, ideally this could have been accounted for during the experiment.





Figure B.6: Effect of K_{min} on average look-ahead time, $N = 2 \times 5 \times 16$, $K_{HSC} = 1$



Experiment Design

This appendix briefly discusses the experiment design. Sixteen subjects were invited to perform a total of eight conditions, shown in Table C.1. During each condition a total of four runs were performed, where the first run was discarded to mitigate learning effects. Figure C.1 shows the segments of each run, where the first 8 seconds are discarded as run-in time. Thereafter, the display settings are switched at $t = 50s$, either going from pursuit to preview, or vice versa.





Figure C.1: Run segments

Table C.1: Experiment conditions

'NH' refers to 'non-haptic', 'PS' indicates pursuit to preview, 'PR' indicates preview to pursuit, 'F0' refers to a controller with pursuit settings, 'F6' refers a controller with preview settings, 'A' refers to the adaptive controller.

Subject	Training			Conditions							
	NH	F0	F6	F0-PS	NH-PR	F6-PS	A-PS	A-PR	NH-PS	F6-PR	F0-PR
1	NH	F0	F6	F0-PS	NH-PR	F6-PS	A-PS	A-PR	NH-PS	F6-PR	F0-PR
2	NH	F6	F0	NH-PR	A-PS	F0-PS	NH-PS	F6-PS	F0-PR	A-PR	F6-PR
3	NH	F0	F6	A-PS	NH-PS	NH-PR	F0-PR	F0-PS	F6-PR	F6-PS	A-PR
4	NH	F6	F0	NH-PS	F0-PR	A-PS	F6-PR	NH-PR	A-PR	F0-PS	F6-PS
5	NH	F0	F6	F0-PR	F6-PR	NH-PS	A-PR	A-PS	F6-PS	NH-PR	F0-PS
6	NH	F6	F0	F6-PR	A-PR	F0-PR	F6-PS	NH-PS	F0-PS	A-PS	NH-PR
7	NH	F0	F6	A-PR	F6-PS	F6-PR	F0-PS	F0-PR	NH-PR	NH-PS	A-PS
8	NH	F6	F0	F6-PS	F0-PS	A-PR	NH-PR	F6-PR	A-PS	F0-PR	NH-PS
9	NH	F0	F6	F6-PR	F0-PS	A-PR	F0-PR	NH-PS	F6-PS	A-PS	NH-PR
10	NH	F6	F0	F0-PS	F0-PR	F6-PR	F6-PS	A-PR	NH-PR	NH-PS	A-PS
11	NH	F0	F6	F0-PR	F6-PS	F0-PS	NH-PR	F6-PR	A-PS	A-PR	NH-PS
12	NH	F6	F0	F6-PS	NH-PR	F0-PR	A-PS	F0-PS	NH-PS	F6-PR	A-PR
13	NH	F0	F6	NH-PR	A-PS	F6-PS	NH-PS	F0-PR	A-PR	F0-PS	F6-PR
14	NH	F6	F0	A-PS	NH-PS	NH-PR	A-PR	F6-PS	F6-PR	F0-PR	F0-PS
15	NH	F0	F6	NH-PS	A-PR	A-PS	F6-PR	NH-PR	F0-PS	F6-PS	F0-PR
16	NH	F6	F0	A-PR	F6-PR	NH-PS	F0-PS	A-PS	F0-PR	NH-PR	F6-PS



Experiment Briefing and Consent Form

The following pages include the experiment briefing and consent form used for the experiment.

Experiment Briefing

Adaptive Haptic Shared Control

Thank you for volunteering to participate in this experiment!

This briefing will provide you with an overview of the goals of the project, the experiment setup, and what is expected from you as a participant. Please read this document carefully, and do not hesitate to ask the researcher any questions that you may have.

Research Context

You are currently situated in the human-machine interface (HMI) lab of the Control and Simulation section of TU Delft's Aerospace Engineering faculty. The main purpose of the laboratory is to allow researchers to perform experiments which investigate how humans control vehicles.

Today you will be asked to perform a **preview tracking task** as shown in Figure 1. This task requires you to place the circle (○) on the target cross (X). This can be done by moving the joystick either left or right as shown in the figure. The future of the target is given by the curved line up ahead.




Figure 1: A Preview Tracking Task

Experiment Objective

The goal of this research is to evaluate different types of support systems that will try and help you match the circle (○) with the cross (X). These systems will do so by applying forces to the joystick, similar to lane-assist systems in cars. For example, considering Figure 1, you may feel a force on the joystick that is pulling in the direction of the target.

In some situations you may feel forces that are counterintuitive or even unhelpful. In these cases it is your responsibility to keep the circle on the target as best as you can, **regardless of what the support system is doing**. The aim of the experiment is to investigate under which circumstances such conflicts occur, and how effective various support systems are in minimising them.

During runs you will see the amount of preview information (the curved line in Figure 1) change halfway throughout the run. It will either go from *full preview* (you see the upcoming route) to *pursuit* (you only see the X and O), or vice versa. **The end goal remains the same throughout this transition: to keep the circle on the target as best as possible at all times, whatever the circumstances.**

Experiment Procedure

Initially you will be offered the chance to perform some runs with and without haptic support to familiarise yourself with the task. Each run lasts approximately 2 minutes.

Once training is complete, a number of runs will again be performed to collect data. Prior to each run you will be told whether or not haptic support will be present. You may or may not feel differences in the behaviour of haptic support during runs. It is important to remember that the primary goal is to minimise the distance between the circle and target as much as possible.

A single run consists of the following steps:

1. The participants waits for the researcher to configure the next run.
2. The researcher checks that the participant is ready and initiates the run after counting down (3-2-1-go).
3. The participant performs the tracking task until completion.

After a set of runs **with** haptic support you will be asked to fill in a questionnaire, rating your experience of the support system. Once the questionnaire is completed, your total score for the set will be communicated.

The entire experiment will last at most 3 hours, including planned breaks taken between sets. More breaks are always available should you feel discomfort at any point during the experiment.

Your Rights & Consent

Participation is entirely voluntary, meaning that you can decide to withdraw from the experiment at any time without stating a reason. By participating in the experiment you agree to the collection of your data (e.g. tracking performance) and agree that it may be published anonymously. All data will be recorded in a way that makes it untraceable back to individual participants, except for the researcher. To ensure you understand and comply with the conditions of the experiment, you will be asked to sign an informed consent form.

Thank you for participating!

<u>Contact information researcher</u> Max Mckenzie	<u>Contact information research supervisor</u> Max Mulder
---	--

Experiment Consent Form

Adaptive Haptic Shared Control

I hereby confirm, by ticking the box, that:

1. I volunteer to participate in the experiment conducted by the researcher (**Max Mckenzie**), under supervision of **Prof. dr. ir. Max Mulder**, from the Faculty of Aerospace Engineering of TU Delft. I understand that my participation in this experiment is voluntary and that I may withdraw (“opt-out”) from the study at any time, for any reason.
2. I have read the briefing document and I understand the experiment instructions, and have had all remaining questions answered to my satisfaction.
3. I understand that my participation involves performing manual tracking tasks in the HMI-Lab simulator at TU Delft. I understand that only the pseudonymized recorded time traces of the tracking tasks are saved and used for data analysis.
4. I understand that the researcher will not identify me by name in any reports or publications that will result from this experiment, and that my confidentiality as a participant in this study will remain secure. Specifically, I understand that any demographic information I provide (gender, age range, *see next page*) will only be used for reference and always presented in aggregated form in scientific publications.
5. I understand that this research study has been reviewed and approved by an ethics commission. To report any problems regarding my participation in the experiment, I know I can contact the researchers using the contact information below.

Date

My Signature

Signature of researcher

Contact information researcher

Max Mckenzie

Contact information research supervisor

Prof. dr. ir. Max Mulder

Participant Demographic Information

Adaptive Haptic Shared Control

Age range:

- 18 – 19
- 20 – 24
- 25 – 29
- 30 – 34
- 35 – 39
- 40 – 44
- 45 – 49
- 50 – 55
- 55+

Gender: _____


Participant number: _____
(filled out by the researcher)

Contact information researcher

Max Mckenzie

Contact information research supervisor

Prof. dr. ir. Max Mulder



Van der Laan Questionnaire

This section includes the Van der Laan questionnaire which was used to evaluate participants' subjective experience.

Condition 1/8

My judgments of the haptic support system are ... (please tick a box on each of the 9 lines)

1	Useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Useless
2	Pleasant	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Unpleasant
3	Bad	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Good
4	Nice	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Annoying
5	Effective	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Superfluous
6	Irritating	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Likeable
7	Assisting	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Worthless
8	Undesirable	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Desirable
9	Raising Alertness	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Sleep-inducing

Condition 2/8

My judgments of the haptic support system are ... (please tick a box on each of the 9 lines)

1	Useful	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Useless
2	Pleasant	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Unpleasant
3	Bad	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Good
4	Nice	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Annoying
5	Effective	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Superfluous
6	Irritating	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Likeable
7	Assisting	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Worthless
8	Undesirable	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Desirable
9	Raising Alertness	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Sleep-inducing

