

Design and Tuning of a Real-time Capable MPC for Automotive Fuel Cell Systems

Reglerentwurf und -tuning einer echtzeitfähigen MPR für automobiler Brennstoffzellensysteme

Von der Fakultät für Maschinenwesen der Rheinisch-Westfälischen Technischen Hochschule Aachen zur Erlangung des akademischen Grades einer Doktorin der Ingenieurwissenschaften genehmigte Dissertation

vorgelegt von

Thuc Anh Nguyen

Berichter: Univ.-Prof. Dr.-Ing. Dirk Abel
Univ.-Prof. Dr.-Ing. Bassam Alrifaae

Tag der mündlichen Prüfung: 27.01.2025

Diese Dissertation ist auf den Internetseiten der Universitätsbibliothek online verfügbar.

For my parents
Danh cho Ba va Me

Acknowledgements

This dissertation was conducted as part of my research at the Institute of Automatic Control at RWTH Aachen University within the Research Training Group (GRK) 1856 *Integrated Energy Supply Modules for On-Road Electromobility*, funded by the German Research Foundation (DFG).

First and foremost, I would like to express my sincere gratitude to Prof. Dirk Abel for giving me the opportunity to pursue my doctoral research under his supervision. His emphasis on applying control methods to real-world systems established a clear and meaningful framework for my research. At the same time, I greatly valued the freedom and trust to shape my own research topic. I would also like to thank Prof. Bassam Alrifaae for co-supervising my work and for his valuable feedback, Prof. Heinrich Rake for serving on the examination committee, and Prof. Sebastian Trimpe for kindly serving as chair of the examination committee.

I feel fortunate to have been part of such an established institution as the Institute of Automatic Control, which allowed me to focus on my research and provided access to the resources, infrastructure, and academic opportunities that significantly supported my work. As part of my doctoral studies, I was given the opportunity to visit Prof. Ali Mesbah's laboratory at the University of California, Berkeley, for which I am deeply grateful. This research stay allowed me to engage with applied control in a different academic environment and has meaningfully contributed to this thesis. I would like to thank Prof. Ali Mesbah and his group for their warm welcome, as well as for their time and effort.

Furthermore, I would like to thank Prof. Oliver Nelles for mentoring me within the TandemDok program, as well as Mrs. Gitta Döbert for her wholehearted support throughout this program. I would also like to thank Verena Neisen for welcoming me as her successor in this research topic. Her passion for fuel cell systems was contagious. My learning during this time extended beyond research activities and was also shaped by teaching responsibilities and the supervision of students. I am grateful for their trust and engagement, in particular Yannik Jünke and the TPR student assistants. I also very much valued the supportive and friendly atmosphere in our office. Andreas, Paul, and Xu – I truly appreciated having such kind colleagues around.

During my PhD, I once received valuable advice on the importance of having a stable social environment – a safety net to fall back on – during a conversation with Prof. Katharina Kohse-Höinghaus. This thought stayed with me, and it proved to be true. Looking back, I am deeply grateful to the many people who accompanied me during different phases of this journey –

those who checked on me when I did not reply, who encouraged me during difficult times, who celebrated with me in good times, and who quietly rooted for me in the background. This support has meant everything to me. In particular, I would like to mention Qiaoxiao, Sonali, Ulvi, Mohammed, Achraf, Aqsa and Ibrahim, Molly, Christina, Arash, Lika, Lara, Richard, Kimberly, Sirine, Nathalie, Philipp, and Marc.

I would especially like to thank my uncle Nam for his endless patience and for always being there for me, as well as my aunt Lan Phuong for her constant encouragement. Finally, I am forever grateful to my parents and my brother, who have always stood firmly by my side. My parents are my greatest role models. Through the way they live their lives, they have shown me perseverance in the face of adversity, the value of honest and hard work, and, most importantly, the importance of keeping a kind and open heart. They taught me to see the good beyond what is visible, to remain humble, and to act with compassion rather than judgment. These values have given me great strength throughout my doctoral journey and continue to guide me in everything I do.

Aachen, January 2026

Abstract

Polymer electrolyte membrane fuel cells are emerging as a promising technology for achieving carbon-free mobility, particularly in the automotive sector. In fuel cell-dominant vehicles, highly dynamic operation of the fuel cell system (FCS) is required, while simultaneously maintaining high efficiency and safe operation. This entails avoiding oxygen starvation, compressor surge and choke, as well as ensuring proper membrane hydration, all of which are also critical for system longevity.

This thesis presents a model predictive control (MPC) framework designed, implemented, and evaluated for an automotive FCS in a fuel cell-dominant hybrid electric vehicle (FCHEV). The control structure is arranged as a hierarchical MPC scheme, composed of a high-level and a low-level control layer. The high-level control layer includes a power-split nonlinear MPC (NMPC) and a target selector. The former optimizes the power distribution between the FCS and the hybrid battery to maximize FCS efficiency and improve battery charge sustainability. Given this power distribution, the target selector ensures a statically optimal allocation of the compressor operating point, preventing overactuation of the FCS. The high-level control layer, which considers a dynamic model of the battery and a static model of the FCS, operates on a one-second sampling time in line with the power demand dynamics. The low-level control layer comprises a tracking MPC and a state and disturbance estimator. Considering the transient behavior of the FCS, the tracking MPC computes the optimal inputs for FCHEV actuation, with the objective to dynamically track the references provided by the high-level control layer while respecting the system's operational bounds. The extended Kalman filter estimates internal states of the fuel cell stack and provides a static model error for zero steady-state offset tracking. To capture the transient system behavior, the tracking MPC operates at a faster rate, which poses challenges for real-time feasibility on embedded hardware. In this context, a comparison between linear time-varying MPC and NMPC shows that nonlinear prediction is advantageous under highly dynamic operation, particularly when rapid power changes interact with safety-critical constraints. This motivates a careful design of the NMPC, including the selection of sampling time, numerical integration scheme, and prediction horizon to ensure real-time feasibility. The real-time feasibility of the resulting controller is demonstrated on embedded hardware. The overall MPC framework, including the real-time capable tracking MPC, is tuned using multi-objective Bayesian optimization to balance dynamic capability, system efficiency, and constraint satisfaction. A further contribution of this work is the incorporation of humidity into the tracking MPC, which is essential for accurately capturing the dynamic behavior of the FCS and enabling proper water management.

Kurzfassung

Polymerelektrolytmembran-Brennstoffzellen gelten als vielversprechende Technologie zur Realisierung einer kohlenstofffreien Mobilität, insbesondere im Automobilbereich. In Fahrzeugen mit dominantem Brennstoffzellenantrieb ist ein hochdynamischer Betrieb des Brennstoffzellensystems (BZS) erforderlich. Gleichzeitig sind eine hohe Effizienz sowie ein sicherer Betrieb zu gewährleisten. Dazu zählen die Vermeidung von Sauerstoffmangel, der Betrieb des Kompressors innerhalb der Schluck- und Pumpgrenze sowie das Aufrechterhalten einer angemessenen Membranfeuchte. Diese Betriebsgrenzen sind zudem relevant für die Lebensdauer des Systems.

Die vorliegende Dissertation stellt ein modellprädiktives Regelungskonzept für ein BZS in einem hybriden Elektrofahrzeug mit dominantem Brennstoffzellenantrieb vor. Das Regelungskonzept ist als hierarchische Struktur ausgelegt und besteht aus einer überlagerten und unterlagerten Regelungsebene. Die überlagerte Ebene umfasst eine nichtlineare modellprädiktive Regelung (MPR) sowie einen Target-Selector. Die nichtlineare MPR (NMPR) optimiert die Leistungsaufteilung zwischen BZS und Hybridbatterie mit dem Ziel, die Effizienz des BZS zu maximieren und eine nachhaltige Batterieladung sicherzustellen. Auf Basis dieser Leistungsaufteilung bestimmt der Target-Selector statisch optimale Betriebspunkte für den Luftpfad und verhindert so eine Überaktuierung des BZS. Die überlagerte Regelungsebene berücksichtigt dabei ein dynamisches Batteriemodell sowie ein statisches BZS-Modell und arbeitet mit einer Abtastzeit von einer Sekunde, die der Dynamik der Leistungsanforderung entspricht. Die unterlagerte Regelungsebene besteht aus einer Tracking-MPR sowie einem Zustands- und Störgrößenschätzer. Die Tracking-MPR berechnet unter Berücksichtigung des transienten Verhaltens des BZS die optimalen Stellgrößen zur Aktuierung des Fahrzeugs, um den von der überlagerten Ebene vorgegebenen Referenzen unter Einhaltung der Betriebsgrenzen dynamisch zu folgen. Um das Systemverhalten abzubilden, wird die Tracking-MPR mit einer höheren Abtastrate ausgeführt, was die Echtzeitfähigkeit auf eingebetteten Steuergeräten zu einer Herausforderung macht. Ein Vergleich zwischen linearer und nichtlinearer MPR zeigt, dass eine nichtlineare Prädiktion unter hochdynamischen Betriebsbedingungen vorteilhaft ist. Vor diesem Hintergrund wird die NMPR gezielt entworfen und ihre Echtzeitfähigkeit durch eine Analyse zentraler Entwurfparameter sowie eine Implementierung auf eingebetteter Hardware nachgewiesen. In der überlagerten Regelung werden ausgewählte Parameter mithilfe mehrzielliger Bayes'scher Optimierung getunt, um den Zielkonflikt zwischen Dynamik, Effizienz und Einhaltung der Betriebsgrenzen systematisch zu adressieren. Darüber hinaus wird das Regelungskonzept um eine explizite Berücksichtigung der Feuchtedynamik erweitert.

Contents

1	Introduction	1
1.1	Motivation for MPC in automotive fuel cell systems	2
1.2	Literature review	3
1.3	Research questions	7
1.4	Thesis outline and contribution	9
2	Preliminaries on fuel cell technology	11
2.1	PEM fuel cells	11
2.1.1	Structure and operating principle	11
2.1.2	Efficiency and performance	13
2.1.3	Degradation	18
2.1.4	Fuel cell stack	18
2.2	PEM fuel cell systems	19
2.3	Fuel cell vehicle	22
2.4	Control objectives	26
3	Control preliminaries	31
3.1	Optimal control	31
3.1.1	Solution methods for optimal control	32
3.1.2	Single shooting	33
3.1.3	Multiple shooting	33
3.1.4	Numerical simulation	34
3.2	Numerical optimization	36
3.2.1	Interior point methods	38
3.2.2	Sequential quadratic programming	39
3.2.3	Quadratic programming	41
3.3	Model predictive control	43
3.3.1	Linear model predictive control	44
3.3.2	Nonlinear model predictive control	49
3.3.3	State estimation	51
3.3.4	Offset-free reference tracking	53
3.4	Bayesian optimization	60

4	Design of model predictive control framework	67
4.1	Control task and control structure	67
4.2	Control-oriented modeling	68
4.2.1	Compressor model	69
4.2.2	Inlet manifold model	70
4.2.3	Fuel cell stack model	71
4.2.4	Outlet manifold model	75
4.2.5	Battery model	76
4.2.6	Power equations	77
4.3	High-level control	78
4.3.1	Target selector	78
4.3.2	Power-split NMPC	82
4.4	Low-level control	86
4.4.1	Prediction model formulation	86
4.4.2	Optimal control problem	87
4.4.3	Transcription to nonlinear program	89
4.4.4	Sequential quadratic programming formulation	92
4.5	State and disturbance estimation	96
5	Expansion to include humidity	101
5.1	Control-oriented modeling	101
5.1.1	Cathode flow model	102
5.1.2	Membrane hydration model	103
5.1.3	Other static relations	105
5.2	High-level control	105
5.3	Low-level control	108
5.3.1	Prediction model formulation	108
5.3.2	Optimal control problem	110
5.3.3	Transcription to nonlinear program	110
6	Simulative studies and results	113
6.1	Comparison of tracking LTV MPC and tracking NMPC	113
6.2	Real-time feasibility of tracking NMPC	118
6.3	Simulative validation of the overall control concept	123
6.3.1	Closed-loop control results	123
6.3.2	Closed loop control results with MOBO tuning	126
6.4	Evaluation of humidity consideration	129
6.5	Summary of results	133
7	Summary and outlook	137

A Appendix **141**

- A.1 Equivalent hydrogen consumption 141
- A.2 Coefficients for stack voltage model 142
- A.3 Approximation of square root function 143
- A.4 Modeling parameters 144

Personal Publications

This thesis incorporates content from the following peer-reviewed publications:

- [N1] NGUYEN, T. A. ; NEISEN, V. ; ABEL, D.: Investigation of the real-time feasibility of NMPC for air-path control in automotive fuel cell systems. In: *2024 Automotive meets Electronics & Control (AmEC); 15th GMM Symposium*. VDE, 2024, pp. 67–72.
- [N2] NGUYEN, T. A. ; WEBER, N. ; ABEL, D.: Parameterization of nonlinear model predictive control for automotive fuel cell systems. In: *2024 European Control Conference (ECC)*. IEEE, 2024, pp. 3588–3593. DOI: 10.23919/ecc64448.2024.10590712.
- [N3] NGUYEN, T. A. ; NEISEN, V. ; KELLER, M. ; KONRAD, T. ; ABEL, D.: Brennstoffzellensysteme: die Regelung als zentraler Wegbereiter: Teil II: Das Potential modellprädiktiver Regelungen. In: *atp magazin* 64. no. 3 (2022), pp. 100–107. DOI: 10.17560/atp.v64i3.2610.
- [N4] NGUYEN, T. A. ; NEISEN, V. ; KELLER, M. ; KONRAD, T. ; ABEL, D.: Brennstoffzellensysteme: die Regelung als zentraler Wegbereiter: Teil I: Systembeschreibung und Herausforderungen. In: *atp magazin* 63. no. 11-12 (2021), pp. 82–89. DOI: 10.17560/atp.v63i11-12.2580.

During my time as a research associate, I contributed to the following publication, which is not part of this thesis:

- [N5] BASLER, M. ; NGUYEN, T. A. ; HRUSCHKA, F. ; JASSMANN, U. ; ABEL, D.: A Genetic Algorithm-based Robust Control Approach for Wind Turbine System Test Benches. In: *2020 European Control Conference (ECC)*. IEEE, 2020, pp. 394–401. DOI: 10.23919/ecc51009.2020.9143631.

List of Figures

2.1	Schematic of the basic structure of a PEMFC. Illustration adapted from [18].	12
2.2	Exemplary polarization curve (left) and corresponding efficiency and power characteristic (right) of a FC. Illustration based on [49].	15
2.3	Exemplary illustration of a FCS for automotive applications, adapted from [85].	20
2.4	Exemplary illustration of a fuel-cell-dominant power train, adapted from [62].	22
3.1	The principle of model predictive control. In the general case, the control horizon N_u can differ from the prediction horizon N . In this work, we use identical control and prediction horizons with the lower limit $N_1 = 0$	44
3.2	Two-layered control structure: MPC regulator, state estimator, and target selector.	55
3.3	Two-layered control structure: MPC regulator, state and disturbance estimator, and target selector.	59
3.4	Illustration of the BO procedure over selected iterations. The left column plots show the estimated mean and confidence intervals of the objective function using a probabilistic model. While the actual objective function is displayed here for illustration, it remains unknown in practice. The right column plots depict the acquisition functions. The acquisition function peaks in regions where the model predicts a high objective value (exploitation) and where prediction uncertainty is high (exploration). [100]	62
3.5	Illustration of the steps in BO, adapted from [29].	63
3.6	Illustration of MOO concepts with $m_{\text{obj}} = 2$ objectives, adapted from [40]: The non-dominated objective vectors are depicted as blue circles. The hypervolume indicator of this non-dominated set \mathcal{A} with respect to the red reference point is the surface of the blue shaded region. The pink rectangle is the hypervolume improvement brought by $\Phi(\theta')$	64
3.7	Illustration of the MOBO framework with $m_{\text{obj}} = 2$ objectives, adapted from [61]. At iteration n , the black-box system is evaluated at θ_n , yielding a noisy observation \mathbf{z}_n of the objective vector. Based on the accumulated data, surrogate models are learned for each objective and an acquisition function, here based on the expected hypervolume improvement, is optimized to select the next query point θ_{n+1}	66
4.1	Simplified schematic of the control structure.	68

List of Figures

4.2	Schematic overview of the model structure used for control-oriented modeling the FCHEV, focusing on air-path dynamics without humidity considerations.	69
4.3	Compressor map.	71
4.4	Compressor efficiency data published in [85].	72
4.5	Fit of compressor efficiency η_{cp}	73
4.6	Static equivalent circuit model.	76
4.7	Resulting FCS efficiency curves η_{fcs} from the nominal steady-state target problem, plotted over corresponding FCS output powers $P_{fcs,bus}$ for different stack temperatures T_{st}	81
4.8	FCS power reference trajectory and operating points, at which the system dynamics are linearized.	90
4.9	Real parts of eigenvalues for each system mode and operating point.	90
4.10	(a) Relative discretization error vs. number of integration steps for tested integration methods, (b) relative discretization error vs. mean CPU time for tested integration schemes.	92
4.11	(a) Relative cost difference and CPU time vs. prediction horizon, (b) CPU time vs. relative cost difference.	96
4.12	Illustration of chronological sequence of computation steps for implementation of an extended Kalman filter (EKF) for a system without feedthrough (a) and with feedthrough (b).	98
4.13	Simplified schematic of the control structure with state and disturbance estimation.	99
5.1	Schematic overview of the model structure used for control-oriented modeling the FCHEV, focusing on air-path dynamics with humidity considerations.	101
5.2	Resulting FCS efficiency curves η_{fcs} from the nominal steady-state target problem with consideration of humidity, plotted over corresponding FCS output powers $P_{fcs,bus}$ for different stack temperatures T_{st}	108
5.3	Real parts of eigenvalues for each system mode and operating point, with humidity consideration.	111
5.4	(a) Relative discretization error vs. number of integration steps for tested integration methods, (b) relative discretization error vs. mean CPU time for tested integration schemes.	112
5.5	(a) Relative cost difference and CPU time vs. prediction horizon, (b) CPU time vs. relative cost difference.	112
6.1	Closed-loop tracking results of total power P_{drive} and FCS power $P_{fcs,bus}$ for SQP NMPC with one iteration and LTV MPC in the step evaluation scenario.	115
6.2	Constraint trajectories for OER and compressor choke and surge, resulting from closed-loop control with SQP NMPC using one iteration and LTV MPC in the step evaluation scenario.	116

6.3	Control input trajectories resulting from closed-loop control with SQP NMPC using one iteration and LTV MPC in the step evaluation scenario.	117
6.4	Selected output and constraint trajectories and their prediction trajectories based on control results from LTV MPC. Top row: Predictions are calculated with nonlinear prediction model; bottom row: Predictions are calculated with linearized prediction model.	118
6.5	Closed-loop output tracking results for all tested NMPC algorithms.	120
6.6	Closed-loop constraint trajectories for all NMPC algorithms.	121
6.7	Closed-loop constraint trajectories of control concept with plant model. . .	124
6.8	Closed-loop constraint trajectories resulting from the overall control concept with plant model.	125
6.9	Left: All observed points and observed Pareto front from MOBO optimization, right: Hypervolume evolution over MOBO iterations.	126
6.10	Closed-loop output tracking results of the overall control concept with plant model and MOBO parameterization.	128
6.11	Closed-loop constraint trajectories of control concept with plant model and MOBO parameterization.	129
6.12	Closed-loop output tracking results for low-level control with humidity consideration.	130
6.13	Closed-loop constraint trajectories for low-level control with humidity consideration.	131
6.14	Closed-loop input trajectories for low-level control with humidity consideration.	132

List of Tables

6.1	Performance metrics for dynamic power delivery and constraint satisfaction, evaluated for SQP NMPC with one iteration and LTV MPC in the step evaluation scenario.	115
6.2	Maximum turnaround time and performance metrics for dynamic power delivery and constraint satisfaction, evaluated for all tested NMPC algorithms in the step evaluation scenario.	122
6.3	Performance metrics for evaluation of the overall control concept with plant model.	126
6.4	Performance metrics for evaluation of the overall control concept with plant model and MOBO tuning.	127
6.5	Performance metrics for dynamic power delivery and constraint satisfaction, evaluated for the tracking NMPC with humidity consideration in the step evaluation scenario.	132
A.1	Thermodynamic constants used for modelling the FCS	144
A.2	Parameters used for modelling the FCS	145

List of Abbreviations

ACL	anode catalyst layer
BEV	battery electric vehicle
BFGS	Broyden-Fletcher-Goldfarb-Shanno
BO	Bayesian optimization
BoP	balance-of-plant
BPP	bipolar plate
BVP	boundary value problem
CCL	cathode catalyst layer
CL	catalyst layer
DM	diffusion media
DP	dynamic programming
ECMS	equivalent consumption minimization strategy
ECSA	electrochemical surface area
EHVI	expected hypervolume improvement
EKF	extended Kalman filter
EMS	energy management strategy
ERK	explicit Runge-Kutta
FC	fuel cell
FCHEV	fuel cell hybrid electric vehicle
FCS	fuel cell system
GDL	gas diffusion layer
GGN	Generalized Gauss-Newton
GHG	greenhouse gas
GL	Gauss-Legendre
GN	Gauss-Newton

List of Abbreviations

GP	Gaussian process
GPR	Gaussian process regression
HEV	hybrid electric vehicle
HHV	higher heating value
HOR	hydrogen oxidation reaction
ICE	internal combustion engine
IRK	implicit Runge-Kutta
IVP	initial value problem
KF	Kalman filter
KKT	Karush-Kuhn-Tucker
LHV	lower heating value
LICQ	linear independence constraint qualification
LMPC	linear MPC
LTI	linear time-invariant
LTV	linear time-varying
MEA	membrane electrode assembly
MIMO	multiple inputs and outputs
MOBO	multi-objective Bayesian optimization
MOO	multi-objective optimization
MPC	model predictive control
MPL	microporous layer
NLP	nonlinear program
NMPC	nonlinear MPC
OCP	optimal control problem
OCV	open-circuit voltage
ODE	ordinary differential equation
OER	oxygen excess ratio
OP	operating point
ORR	oxygen reduction reaction
PEM	polymer electrolyte membrane

PEMFC	polymer electrolyte membrane fuel cell
QP	quadratic programming
RDIIa	Radau IIa
RK	Runge-Kutta
RTI	real-time iteration
SoC	state of charge
SQP	sequential quadratic programming
STP	standard temperature and pressure
USDoE	United States Department of Energy

List of Symbols

Symbol	Description	Unit
Greek Symbols		
α	Step size used in the update of primal and dual variables in SQP methods	—
η	Efficiency	—
$\eta_{fc,max}$	Maximum theoretical FC efficiency	—
Λ	Water content	—
λ	Stoichiometric ratio	—
μ	Lagrange multipliers associated with the inequality constraints	—
ν	Lagrange multipliers associated with the equality constraints of a nonlinear optimization problem	—
ω	Angular velocity	rad s ⁻¹
$\Phi(\cdot)$	Unknown black-box performance function mapping the design parameters θ to a scalar closed-loop performance measure	—
$\phi(\cdot)$	Static mapping function used to describe algebraic input–output relations	—
$\Phi(\cdot)$	Vector-valued objective function in multi-objective optimization, collecting m competing objective functions $\Phi_i(\cdot)$ into a single mapping $\Phi : \Theta \rightarrow \mathbb{R}^m$	—
$\hat{\Phi}(\cdot)$	Surrogate model of the unknown performance function $\Phi(\cdot)$ used in Bayesian optimization, typically obtained from a Gaussian process posterior	—
$\hat{\Phi}(\cdot)$	Vector-valued surrogate model of the objective function, typically modeled using independent Gaussian processes for each objective	—
$\phi_m(\cdot)$	Nonlinear measurement function mapping the system state and measurement noise to the measured outputs	—
Π	Pressure ratio	—

List of Symbols

σ_{mem}	Protonic conductivity of the membrane	S cm^{-1}
ε	Slack variables	—
τ	Barrier parameter in interior point methods, controlling the smooth approximation of the complementarity conditions	—
Θ	Design space of the Bayesian optimization problem, defining the admissible domain of the design parameters θ	—
θ	Design parameter (hyperparameter) vector optimized by Bayesian optimization	—
Θ^*	Pareto set	—
ξ	Primal–dual iterate comprising the decision variables and Lagrange multipliers, $\xi := (w, v, \mu)$, used in Newton-type optimization methods	—
ζ	Stacked linear stage cost vector, $\zeta_k := \begin{bmatrix} \zeta_{x,k}^\top & \zeta_{u,k}^\top \end{bmatrix}^\top$	—
$\zeta_{u,k}$	Linear cost coefficient vector for the input term at stage k	—
$\zeta_{x,k}$	Linear cost coefficient vector for the state term at stage k	—
$\zeta_{x,N}$	Linear terminal state cost vector in the quadratic objective of the sparse linear MPC formulation	—

Latin Symbols

A	Continuous-time state matrix of the linear state-space model	—
$a(\cdot)$	Acquisition function guiding the selection of the next query point in Bayesian optimization	—
A_{eq}	Equality constraint matrix of a quadratic program	—
$A_{\text{eq},S}$	Stacked equality constraint matrix of the sparse linear MPC quadratic program	—
A_{fc}	Active FC area	cm^2
A_{ineq}	Inequality constraint matrix of a quadratic program	—
$A_{\text{ineq},S}$	Stacked inequality constraint matrix of the sparse linear MPC quadratic program	—
B	Continuous-time input matrix of the linear state-space model	—
b	Right-hand side vector of the affine inequality constraints in linear model predictive control	—
B_d	Disturbance-to-state coupling matrix in the augmented system model	—
b_{eq}	Equality constraint vector of a quadratic program	—

$b_{\text{eq,S}}$	Right-hand side vector of the stacked equality constraints in the sparse linear MPC formulation	—
b_{ineq}	Inequality constraint vector of a quadratic program	—
$b_{\text{ineq,S}}$	Right-hand side vector of the stacked inequality constraints in the sparse linear MPC formulation	—
b_N	Right-hand side vector of the affine terminal inequality constraint in linear model predictive control	—
a_{ir}, b_i, c_i	Coefficients $\{a_{ir}, b_i, c_i\}$ defining a Runge–Kutta integration scheme, where a_{ir} are the stage coupling coefficients, b_i are the weights for combining stage derivatives, and c_i specify the relative time locations of the stages	—
C	Continuous-time output matrix of the linear state-space model	—
$c(\cdot)$	Equality constraint function of the nonlinear optimization problem	—
C_D	Inequality constraint matrix of the dense quadratic program	—
c_D	Offset vector of the dense inequality constraints	—
C_m	Measurement matrix mapping the system state to measured outputs	
C_d	Disturbance-to-output coupling matrix in the measured output equation	—
$C_{u,k}$	Input constraint matrix at time step k in linear model predictive control	—
$C_{x,k}$	State constraint matrix at time step k in linear model predictive control	—
$C_{x,N}$	Terminal state constraint matrix in linear model predictive control	—
D	Feedthrough matrix	—
d	Constant disturbance vector acting on the measured output	—
\hat{d}	Estimated disturbance vector obtained from the augmented state observer	—
\mathcal{D}_n	Set of observed data points up to iteration n , consisting of input–output pairs (θ_i, y_i)	—
$E_{\text{lhv,eq}}$	Total equivalent energy consumption based on the LHV	J
E_{rev}	Reversible OCV	V
E_{rev}^0	Standard reversible OCV	1.229 V

List of Symbols

$E_{\text{rev,h}}^0$	Standard reversible OCV based on LHV	1.25 V
E_{rev}°	Reversible OCV at standard pressure	V
F	Faraday constant	96485 C mol ⁻¹
$\mathbf{f}(\mathbf{x})$	Objective vector (BO context)	—
$f(\cdot)$	System dynamics function	—
f_d	discrete-time system dynamics	—
G_f	Gibbs free energy of formation	J
$g(\cdot)$	Inequality constraint function of the finite-dimensional nonlinear optimization problem	—
\bar{g}_f	Molar specific Gibbs free energy of formation	J mol ⁻¹
H	Hessian or Hessian approximation of the Lagrangian used in sequential quadratic programming	—
$h(\cdot)$	Path constraint function in the optimal control problem	—
H_D	Hessian matrix of the dense quadratic program resulting from state elimination	—
$h_f(\cdot)$	Terminal constraint function in the optimal control problem	—
\bar{h}_f	Molar specific enthalpy of formation	J mol ⁻¹
$\Delta\bar{h}_f^0$	Standard change in molar specific enthalpy of formation	-241.83 kJ mol ⁻¹
H_{GN}	Gauss-Newton Hessian approximation used in SQP-based nonlinear model predictive control	—
H_N	Terminal Hessian (or terminal cost Hessian) used in nonlinear model predictive control	—
H_S	block-diagonal Hessian matrix of the sparse quadratic program in linear MPC	—
I	Electric current	A
i	Current density	A cm ⁻²
I_H	Hypervolume indicator	—
$J(\cdot)$	Objective (cost) function of the nonlinear optimization problem	—
K	Kalman gain	—
k	Correction coefficient for indirect hydrogen consumption	—
$\mathcal{L}(\cdot)$	Lagrangian function of the nonlinear optimization problem	—
$\ell(\cdot)$	Stage cost function in the optimal control problem	—
$\ell_f(\cdot)$	Terminal cost function in the optimal control problem	—

L	Jacobian of the discrete-time system dynamics f_d with respect to the process noise	—
M	Molar mass	kg mol^{-1}
m	Mass	kg
m_{obj}	Number of objective functions in a multi-objective optimization problem	—
\dot{m}	Mass flow rate	kg s^{-1}
n_{θ}	Dimension of the design parameter vector θ	—
n_d	Number of constant disturbances	—
\dot{n}	Molar flow rate	mol s^{-1}
n_{eq}	Number of equality constraints in the nonlinear optimization problem	—
n_{fc}	Number of individual fuel cells (FCs) connected in series	—
n_{ineq}	Number of inequality constraints in the optimal control problem and in the corresponding finite-dimensional nonlinear optimization problem	—
$n_{\text{ineq,f}}$	Number of terminal inequality constraints in the optimal control problem	—
N_{int}	Number of integration steps per sampling interval	—
N	Prediction horizon	—
n_u	Dimension of the control input vector	—
n_w	Dimension of the finite-dimensional decision variable vector of the discretized optimal control problem	—
n_x	Dimension of the state vector	—
n_{y_m}	Number of measured output signals	—
n_y	Number of controlled (tracked) output signals	—
p	Pressure	Pa
p_0	Standard pressure	101 325 Pa
P	Electrical power	W
	A-posteriori estimation error covariance matrix	—
\mathcal{P}_{Θ}	Pareto set, i.e. the set of non-dominated decision vectors $\theta \in \Theta$ in a multi-objective optimization problem	—
\mathcal{P}_y	Pareto front	—
p_{bo}	Vector of back-off parameters used to tighten safety constraints	—

List of Symbols

P^-	A-priori estimation error covariance matrix	—
\hat{P}_y	Approximated Pareto front in the performance space	—
Q	Positive semi-definite state weighting matrix in the stage cost of model predictive control formulations	—
q	Linear term in the objective function of a quadratic program	—
q_D	Linear cost vector of the dense quadratic program	—
Q_e	Process noise covariance matrix used in state estimation	—
Q_f	Terminal state weighting matrix used in the terminal cost of model predictive control formulations	—
Q_N	Terminal state weighting matrix used in the terminal cost of linear model predictive control	—
q_S	Stacked linear cost vector of the sparse quadratic program	—
\mathbf{r}	Reference point used in the definition of the hypervolume indicator	—
R	Universal gas constant	8.314 J mol ⁻¹ K ⁻¹
	Positive definite input weighting matrix	—
$r(\cdot)$	Residual function defining the nonlinear least-squares objective in Gauss–Newton and generalized Gauss–Newton methods	—
R_e	Measurement noise covariance matrix used in state estimation	—
RH	Relative humidity	—
\bar{s}	Molar specific entropy	J mol ⁻¹ K ⁻¹
s	Number of stages of the Runge-Kutta method	—
$\Delta\bar{s}^\circ$	Change in molar specific entropy at standard pressure	J mol ⁻¹ K ⁻¹
S_k	State–input cross-weighting matrix	—
S_l	Linear penalty weight matrix for slack variables	—
S_q	Quadratic penalty weight matrix for slack variables	—
T	Temperature	K
T_0	Standard temperature (corresponds to 25 °C)	298.15 K
T_s	Sampling time	s
T_{int}	Numerical integration time step	s
T_y	Output mapping matrix relating the measured output y_m to the controlled output y	—
t	Time	s

U	Terminal voltage	V
u	Control input vector	—
u_{OP}	Input vector at the operating point used for linearization	—
u_{rel}	Relative input variable defined as deviation from the operating point	—
$u_{0:N-1}$	Control input sequence over the prediction horizon $u_{0:N-1}^{\top} = [u_0^{\top}, u_1^{\top}, \dots, u_{N-1}^{\top}]$	—
u_{aug}	Input rate vector in delta formulation	—
u_{∞}	Asymptotic closed-loop input	—
u_{-1}	Input applied to the system at the previous sampling instant	—
u_s	Steady-state target input	—
V	Jacobian of the measurement function ϕ_m with respect to the measurement noise	—
v_m	Measurement noise	—
w	Decision variable vector of the nonlinear optimization problem	—
w_p	Process noise	—
$x_{0:N}$	State sequence over the prediction horizon $x_{0:N}^{\top} = [x_0^{\top}, \dots, x_N^{\top}]$	—
x	State vector	—
x_{OP}	State vector at the operating point used for linearization	—
x_{rel}	Relative state variable defined as deviation from the operating point	—
x_{aug}	Augmented state vector in delta formulation	—
\hat{x}	A-posteriori state estimate	—
\hat{x}^-	A-priori state estimate	—
x_{∞}	Asymptotic closed-loop state	—
x_s	Steady-state target state	—
x_{SoC}	Battery state of charge	—
\tilde{x}	Intermediate state used during numerical integration	—
y	System output vector	—
\tilde{y}_{ref}	Modified output reference that compensates for the estimated constant disturbance	—
\mathcal{Y}	Objective (performance) space induced by the vector-valued objective function in multi-objective optimization	—

List of Symbols

y_∞	Asymptotic closed-loop output	—
y^{llc}	Output vector of the low-level controller	—
y_m	Measured system output vector used in state estimation	—
y_{ref}	Output reference	—
y_s	Steady-state target output	—
z	Noisy observation of the Bayesian optimization objective function	—
Sets		
\mathbb{A}	Index set of active inequality constraints	—
\mathbb{I}	Index set of inactive inequality constraints	—
Indices		
a	dry air	
act	activation; associated with charge transfer overpotential	
an	anode	
aux	auxiliary	
bat	battery	
bat,bus	battery DC bus quantity	
BoP	balance-of-plant; refers to all auxiliary system components	
ca	cathode	
cm	compressor motor	
conc	concentration; associated with mass transport limitations	
cp	compressor	
cr	compressor corrected	
DCDC	DC–DC converter	
drive	refers to the total power for vehicle movement	
eq	equivalent; indicates a value converted to a corresponding hydrogen basis	
fc	fuel cell (FC)	
fchev	fuel cell hybrid electric vehicle (FCHEV)	
fcs	fuel cell system (FCS)	
fcs,bus	FCS measured at DC bus	
H_2	hydrogen	

H_2O	water vapor
hum	humidifier
im	inlet manifold
in	inflow
mem	membrane
N_2	nitrogen
net	refers to the effective value after subtracting parasitic losses
O_2	oxygen
ohm	ohmic; related to ionic and electronic conduction resistances
om	outlet manifold
out	outflow
react	reacted
ref	reference; indicates a requested setpoint or target value
rev	reversible
<i>sat</i>	saturation condition
st	FC stack
sto	stoichiometric

1 Introduction

Global warming is one of the greatest challenges facing mankind. Among its many contributors, the transport sector accounts for a significant share, responsible for 23 % of global energy-related CO₂ emissions. Notably, 70 % of these direct transport emissions originate from road vehicles [81]. The ongoing rise in greenhouse gas (GHG) emissions [102], exacerbating global warming and its consequences, calls for an urgent transformation of our existing mobility systems.

A significant portion of these emissions stems from the use of fossil fuels in internal combustion engines (ICEs). Additionally, the operation of these engines contributes to local noise and pollutant emissions, impairing the quality of life in densely populated urban areas with heavy traffic volumes. Therefore, introducing new technologies in the mass market to drastically reduce the carbon footprint of road vehicles is essential. In recent years, battery electric vehicles (BEVs) have gained market share due to significant cost reductions [102], national subsidy programs [22], and increasing restrictions on ICE vehicles. BEVs, particularly those with higher ranges, require large batteries with capacities of up to 100kWh or more. This results in vehicles that are more expensive and heavier compared to their ICE-driven counterparts. Moreover, the electricity for charging must be supplied by the grid exactly at the time of charging, which poses a growing challenge due to the increased share of volatile renewable energy sources [18].

Fuel cell hybrid electric vehicles (FCHEVs) offer a promising alternative to complement BEVs. When fueled with hydrogen, the FCHEV operates without emitting GHGs. Particularly when the hydrogen is sourced from renewable sources, this method of propulsion has the potential to achieve a mobility solution that is both pollutant-free and CO₂-neutral. A notable advantage of these vehicles is their short refueling times, comparable to those of conventional vehicles. For this reason, this technology holds significant promise for applications requiring high daily mileage.

Despite multiple manufacturers offering FCHEVs, their adoption remains limited, with an estimated 87,600 FCHEVs in operation in 2023, according to the International Energy Agency [57]. The volatility of renewable energy sources necessitates large-scale energy storage solutions, with hydrogen storage - produced from water via electrolysis - being a viable option. However, the scarcity of refueling stations presents a significant challenge, creating a chicken-and-egg dilemma between the availability of vehicles and the supporting infrastructure.

In addition to this infrastructural challenge, further advancements in FC technology for vehicle applications are essential. The United States Department of Energy (USDoE) has established detailed technical targets related to power density, efficiency, cost, and durability. While significant progress has been made in achieving power densities and efficiencies near the USDoE's ultimate targets, reducing system costs and significantly enhancing durability remain critical goals [80]. Additionally, Klell, Eichlseder, and Trattner [62] identify improving dynamic performance as another major challenge that requires further research. [18]

Advanced control algorithms, particularly those in the field of optimal control, can complement other efforts to address the challenges associated with the application of FC technology in vehicles, serving as a crucial lever for improvement. By formulating control tasks as optimization problems, control objectives can be encoded in a cost function optimized within the constraints of desired or safe operation. To leverage feedback from online measurements for disturbance rejection, plant-model mismatch compensation, and system stabilization or reference tracking, these control algorithms are designed for closed-loop, online operation. This requires their deployment on embedded hardware with limited memory and computing power. Specialized algorithms have been developed to implement the optimal control strategies for execution in real time, within strict time constraints [31–34]. This thesis aims to apply these tailored optimal control algorithms to the control of automotive FCS within FCHEVs.

1.1 Motivation for MPC in automotive fuel cell systems

In a FCHEV, the FCS is hybridized with a battery. This work focuses on vehicle setups where the battery cannot supply the full electric power required by the drivetrain, referred to as fuel-cell-dominant setups. The additional degree of freedom introduced by hybridization necessitates an energy management system for optimal power distribution between the FCS and the battery, particularly to maximize steady-state efficiency. Furthermore, and of particular importance in this work, automotive applications demand challenging dynamic capabilities, such that fuel-cell-dominant vehicle setups result in the most demanding dynamic requirements for the FCS.

To understand this better, we briefly explain how FCs operate in vehicles. FCs convert hydrogen and oxygen into electricity, heat, and water via an electrochemical reaction. In the automotive context, low-temperature polymer electrolyte membrane fuel cells (PEMFCs) are predominantly used, as is the case in this work. The reliable and efficient operation of the FC stack, which consists of multiple FCs in series, requires integration into a comprehensive FCS. The FCS incorporates peripheral components, known as balance-of-plant (BoP) components, responsible for supplying reactants and managing other operating conditions such as temperature and humidity. This results in a complex, nonlinear system with time constants of different scales and various interconnected control loops with multiple inputs and outputs (MIMO). During operation, these outputs must be controlled to follow a dynamic trajectory and stay within

operational boundaries for safe operation. These operational limits, associated with both the FC and the BoP components, make efficient and especially dynamic power delivery at the vehicle level challenging. Specifically, for the FCS to achieve highly dynamic power delivery, it is often required to operate close to or at the boundaries without violating them to avoid instabilities or damage to the system.

Among the family of optimal control algorithms, MPC is particularly suited to address these challenges. In each sampling interval, MPC predicts the future evolution of the system based on a model within a specified time horizon and optimizes the control inputs according to objectives and constraints specified in an optimization problem, which is solved online. The first portion of the optimized control sequence is commanded to the plant before the optimization process is repeated with new feedback information. The ability of MPC algorithms to consider coupled and nonlinear MIMO systems, and explicitly consider operational and safety limits in the control design, makes them an attractive choice for the control task in automotive FCS.

In conclusion, the motivation to consider MPC for controlling automotive FCS stems from the need to deal with control objectives and system constraints that are opposing, especially with respect to the target of dynamic power delivery. At the same time, designing an effective and computationally efficient MPC algorithm for practical deployment, particularly in the presence of nonlinear predictions, is not trivial and requires careful consideration in the control design process. This challenge is addressed as part of this work.

1.2 Literature review

The control of FCSs encompasses various critical tasks, including regulating air and fuel supply and managing humidity and temperature to ensure efficient power delivery and longevity of the system. Numerous control approaches have been employed to address these tasks, either individually or in combination. This literature review focuses on the application of MPC to automotive FCSs. The primary objective is to evaluate significant contributions in the field and to assess the current state of research on MPC applied to automotive FCSs.

Early pioneering work in this field has been significantly influenced by Professor Anna Stefanopoulou and her team at the University of Michigan. One of their major contributions, led by Pukrushpan et al., is the development of a control-oriented, semi-physical dynamic model of an automotive FCS [83, 85]. This model uses a zero-dimensional approach, assuming no spatial variation within the FC, as well as no cell-to-cell variation within the FC stack. A similar modeling approach is presented in [48]. However, the model in [83, 85] is implemented in MATLAB/SIMULINK[®] and made open source [86], facilitating controller development not only within her group, but also broadly in the FC control community, particularly for model-based control strategies.

The model particularly provides a comprehensive depiction of the air-path dynamics. The emphasis on the air-supply subsystem makes sense due to its significant impact on dynamic power delivery, efficiency and safe operation of the system. The stack temperature is considered constant, as its dynamics evolve much more slowly. The dynamics of the hydrogen-supply subsystem are assumed to be fast and therefore negligible, since the corresponding mass flow can be quickly adjusted when supplied from a high-pressure tank. Modeling humidity in the FC was attempted but found to be not representative and was therefore not considered for further analysis. Instead, a constant membrane hydration is assumed. The effect of flooding is not modeled at all. In general, humidity and membrane hydration are complex to model and difficult to validate. Additionally, membrane hydration was not a primary focus in the FCS control community for a long time. Although the overall model is not fully validated, it qualitatively represents the air-path transient behavior well.

As a result, the authors focus their control efforts on the air-supply subsystem to replenish and regulate the oxygen depleted from the cathode during step changes in current drawn from the stack. This task needs to be achieved fast and efficiently to avoid oxygen starvation, a condition that occurs when the partial pressure of oxygen drops below a critical level at any location within the cathode, leading to a rapid decrease in cell voltage, which in severe cases can cause short circuits and hot spots on the surface of the cell's membrane. Although spatially variant, it is widely believed that oxygen starvation can be mitigated by maintaining an appropriate oxygen excess ratio (OER) in the cathode, a lumped (spatially invariant) variable. Maintaining a certain OER-level means ensuring a certain surplus of oxygen supplied to the cathode compared to that necessary for the reaction. In [83], this is done by actuating the compressor motor voltage. The analyzed control strategies are static feedforward, dynamic feedforward, and a linear-quadratic regulator in combination with a state observer. The latter two are based on linearizations of the nonlinear model around a nominal operating point. Analyses of the air path reveal that minimizing parasitic losses and providing fast air-flow regulation are conflicting objectives. The conflict arises because the compressor uses part of the stack power to accelerate. Therefore, a compromise needs to be made between preventing oxygen starvation and providing FCS net power for fast transients.

Based on the model in [83], Suh et al. elaborate the inherent limitations of autonomous operation of PEMFCs with compressor-driven air supply. They emphasize the challenges in controlling the air supply for dynamic net FCS power delivery when the compressor is powered by the FC stack. Specifically, their work quantifies the coupling between the electric and the flow paths in the system. This coupling imposes control limitations through nonminimum-phase zeros between the compressor motor voltage, the air mass flow and the OER, resulting in the aforementioned trade-off. [108, 109]

The model presented in [83] is further utilized by Vahidi et al. in a hybrid system incorporating a capacitor [114]. The resulting model is linearized once, and the linearized model employed in an MPC. The idea is to enable high dynamic power delivery while maintaining an OER of 2 through hybridization. The control inputs are the compressor motor voltage and the current

drawn from the stack, which was previously considered a disturbance. The same approach is extended in [115] to include constraints for compressor choke and surge. In [117], the same idea is realized via a hierarchical, decentralized MPC framework. In [113, 116], the authors employ load governors for constraint management, where oxygen starvation is treated as a constraint, ensuring it remains above a minimum value rather than being tracked. Overall, the results show visible violations of the surge boundary in closed-loop simulations with the nonlinear plant, where the MPC handles the constraints more effectively at the cost of potentially higher computational burden. Generally, the authors established that considerable differences exist between the linear model from linearization around one operating point and the nonlinear model, leading to compromised performance of the proposed algorithms.

Though originating from the same group, Goshtasbi et al. expanded the research scope beyond the work previously established at the University of Michigan. Moving away from the zero-dimensional approach, the authors developed a spatially variant model of the FC [46]. This effort is motivated by the need to model internal states of the FC more precisely, which are critical to control and/or to monitor during operation. Integrating this model into an automotive FCS, they applied MPC based on successive linearization, also known as linear time-varying (LTV) MPC [44, 45]. The control task addresses not only the air-path subsystem, but also membrane hydration and FC temperature. Thus, the constraints incorporated are related to the corresponding subsystems and the critical internal states of the FC. The control objective is to track a dynamic power demand while maximizing efficiency. This holistic control approach addresses important degradation issues, such as flooding and drying of the membrane, offering a more comprehensive solution for maintaining the health and efficient operation of the automotive FCS.

Related to the work by Goshtasbi et al., which was further developed at FORD[®] in Michigan, is the research published by Bacher-Chong et al. from the University of Vermont. They use a linear time-invariant (LTI) model obtained from linearizing the nonlinear model proposed in [83] to implement a reference governor-based control scheme for the air path of an automotive FCS. The goal is to achieve efficient power delivery without causing oxygen starvation [7, 8]. One important contribution is the introduction of a third actuator, the backpressure valve at the outlet of the air path. This additional degree of freedom, in conjunction with the stack current and the compressor motor voltage, enables more efficient power delivery and better management of the OER constraint without significantly increasing parasitic losses.

Another notable group working on control-oriented modeling and control of automotive FCSs is from the Vienna University of Technology. Daniel Ritzberger et al. developed a zero-dimensional, dynamic model for the FC stack, which provides a more comprehensive representation of the anode and the humidity in the stack. This model is evaluated using experimental data from a FC research vehicle during on-road operation [94, 95]. Based on this model, Vrlić et al. developed an LTV MPC for automotive FCSs. The control problem is designed to track dynamic power demands and minimize losses, specifically voltage losses of the FC stack. Constraints are imposed on the anode and cathode inlet pressure, the pressure difference across

the membrane, and variables related to oxygen and hydrogen starvation. The control inputs include the air and hydrogen volume flow, the stack current, and the cathode backpressure valve position. Note that these actuated variables are related to the FC stack, indicating that the BoP components are not considered in the control problem. Notably, the control problem extends beyond the cathode side of the FC stack [121]. In a subsequent publication, the authors incorporate a compressor power calculation into the model, integrating it into the net system power and system efficiency calculations. Thus, the compressor power consumption is considered in the objective function designed for power tracking and efficiency maximization, alongside minimizing the pressure difference across the membrane [120]. In another study, the authors implement an MPC based on successive linearization, with further adaptations to the control model and problem. This includes adding the purge valve at the outlet of the hydrogen path as an actuator and expanding the objective function to track an optimal line in the compressor map [122].

Another important research hub in the area of automotive FCSs is at BOSCH® in Renningen. Notable work in the field of control is the work by Hahn et al. The authors apply LTV MPC based on a lumped model of an automotive FCS, which mainly comprises the air path but also considers the heat generation during FC operation. The control problem aims to minimize the net power tracking error and hydrogen consumption. Control inputs include the inlet and outlet air path valve openings, the torque provided by the electric machine of the compressor, and the oxygen stoichiometry. Notable constraints, which extend beyond the air path and power generation of the FCS, include bounds enforced on the cathode outlet activity and the heat flow rate [50].

The MPC approaches presented so far each rely on a linear prediction model derived from linearizing a nonlinear system representation. Depending on the study, the linearization remains constant or is re-evaluated at each sampling instant. Research on the application of NMPC to automotive FCSs is relatively scarce. Hähnel et al. from Helmut Schmidt University in Hamburg implemented various NMPC versions with different structures and slightly different control formulations [51–55]. The version in [53] features an offset-free NMPC implementation for dynamic tracking of a stack power reference. Control inputs include the references for the anode and cathode pressure, stack temperature and cathode inlet mass flow. Bounds are enforced on the stack current, and the controller-internal model is a lumped representation of the FC stack. All proposed algorithms are tested on a test bench with PM200 FC stack from PROTON MOTORS®, consisting of 48 cells and producing up to 4.4 kW of electrical power. Due to the relatively small size of the stack, the system and its requirements are quite different from those in automotive applications.

Luna et al. from the Polytechnic University of Catalonia implemented several NMPC algorithms using a more involved, spatially distributed model of an automotive FCS. The control inputs include the compressor current, cathode relative humidity, stack temperature and hydrogen inlet flow. The primary control objective is to operate the peripheral subsystems with maximum efficiency, which is encoded in the cost function, while power tracking is not part of the control

problem formulation. Constraints are enforced on the spatially distributed concentrations of hydrogen, oxygen and water within their respective channels. The sampling time of the NMPC is chosen to be 100ms [69]. In a subsequent publication, the authors adjusted the stack temperature input to the coolant mass flow. Furthermore, the control objective shifted to maximizing the electrochemical surface area (ECSA) at the cathode catalyst layer, enhancing the available reaction area of the stack and preventing starvation at the catalyst sites [70]. Overall, the authors present a holistic control approach with a focus on efficiency maximization and degradation mitigation.

Neisen et al. from the Institute of Automatic Control at RWTH Aachen University apply NMPC to the air path of an automotive FCS based on a prediction model derived from [83]. Their objective function combines an economic cost function term and a reference tracking cost function term to maximize FCS efficiency while following dynamic power demands. Key constraints include avoiding compressor surge and choke boundaries as well as preventing oxygen starvation. The optimal control problem (OCP) is solved using the GRAMPC toolbox, which employs the indirect method for optimal control. [74]

In none of the above-mentioned publications is the real-time capability on embedded hardware discussed or analyzed for solving the optimization problem arising in the control of automotive FCSs. Several researchers have addressed the challenge of real-time capability in NMPC methods by developing numerical techniques to enhance their computational efficiency, which have been applied across various applications [33, 88]. Despite these methods, achieving fast, reliable, real-time NMPC solutions remains a recurring challenge due to the intricacies of each application. Research addressing this gap for automotive FCSs remains limited. Schmitt et al. have notably demonstrated the real-time capability of NMPC algorithms for air-path control on a small-scale FCS test rig. The FC stack, consisting of 70 cells with a nominal operating point at 200 A, is simulated in the loop during experimental evaluation. In their studies, [97] uses a Wiener model as the control model, while [98] employs a physics-based control model comparable to that in [83]. Efficient numerical integration schemes and condensing methods are explored in [98] and [99], respectively.

1.3 Research questions

Based on the literature review in Section 1.2, we can conclude that:

- Most studies employ a zero-dimensional control model and utilize system-level control levers instead of stack- or cell-specific control levers.
- The majority of the research focuses on the air-path dynamics and dynamic power delivery.

- Many studies apply MPC with a linear prediction model and do not investigate the real-time feasibility of their proposed control algorithm.

To set the stage for further investigations, we first define the requirements for the control of automotive FCS.

Research question 1:

What are the requirements for control systems in automotive FCSs?

To remain optimization-friendly, we also adopt a system-level, zero-dimensional modeling approach and primarily focus on the air path and dynamic power delivery problem in this work. Though this limited perspective does not cover all control requirements in automotive FCSs control, it addresses the most relevant ones. On this basis, we analyze whether a linear prediction is sufficient for the control performance required in automotive FCSs.

Research question 2:

Is a linearized prediction sufficient to meet the control requirements for automotive FCSs when applying MPC?

Based on the findings from research question 2, we evaluate the real-time feasibility of the linear or nonlinear MPC control algorithm.

Research question 3:

Is it possible to design an effective and real-time feasible MPC algorithm for automotive FCSs?

Subsequently, we present a hierarchical MPC control scheme in closed loop with a plant model of a FCHEV. This hierarchical MPC control scheme has been adapted based on earlier research [73] and incorporates our developed MPC algorithm. The closed-loop control system is tuned with multi-objective Bayesian optimization (MOBO) to achieve an optimal trade-off between relevant control objectives.

Research question 4:

How can the hierarchical MPC scheme achieve an optimal trade-off between dynamic power delivery, constraint satisfaction, and efficiency?

Lastly, we expand the automotive FCS model to consider humidity and membrane hydration in the control task and re-evaluate the performance of our MPC algorithm.

Research question 5:

How does the consideration of humidity affect control performance?

1.4 Thesis outline and contribution

The thesis is structured into seven chapters. Chapter 1 provides an introduction to the topic of optimal control of automotive FCSs and its relevance. A literature review is conducted in Section 1.2, which is focused on the application of MPC to automotive FCSs. Subsequently, research questions for this work are derived in Section 1.3. We conclude the chapter with an outline and state the contributions of this thesis. The fundamentals of the application, automotive FCSs in fuel-cell dominant vehicles, are described in Chapter 2, which provides the basis for answering research question 1. This is followed by the relevant basics on control theory presented in Chapter 3. Chapter 4 then dives into the controller design of the hierarchical MPC framework with a focus on the low-level tracking MPC. The latter's design is expanded for humidity consideration in Chapter 5. All simulative studies and results to answer the research questions 2 to 5 are presented in Chapter 6. Finally, a conclusion to the thesis and potential directions for future studies is given in Chapter 7.

The main contributions of this thesis are as follows:

- (i) We investigate whether a linear prediction is sufficient for the purposes of automotive FCS control, and design a real-time feasible MPC algorithm.
- (ii) The MPC algorithm is expanded to incorporate membrane hydration.
- (iii) We demonstrate how MOBO can be applied to closed-loop control tuning.

2 Preliminaries on fuel cell technology

This chapter provides an overview of the fundamentals of PEMFC operation and the corresponding FCS in the context of automotive applications. Understanding these basics is crucial for deriving the control objectives. The elaboration begins with the operation of the PEMFC itself, then expands to automotive FCSs, and finally to FC vehicles. The chapter concludes with the derivation of the control task relevant to this thesis, thereby addressing research question 1.

2.1 PEM fuel cells

2.1.1 Structure and operating principle

In a FC, an electro-chemical reaction converts chemical energy into electrical energy and heat. PEMFCs, a specific type of FC commonly used in automotive applications, utilize hydrogen stored in high pressure tanks as fuel and oxygen from the ambient air as the oxidant. The reaction between hydrogen and oxygen produces water, a non-toxic product:

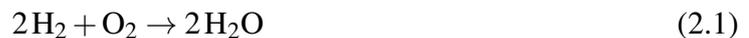
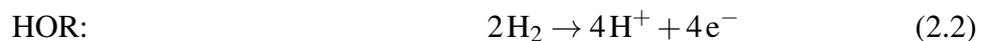


Fig. 2.1 illustrates the basic structure of a PEMFC. A single cell consists of a polymer electrolyte membrane (PEM), catalyst layers (CLs) and diffusion media (DM) on both the anode and cathode side, clamped together by two bipolar plates (BPPs). Each DM includes a gas diffusion layer (GDL) and an optional microporous layer (MPL), with only the GDL depicted here. The BPPs have machined grooves for reactant and coolant transport. [43]

The electrodes are the CLs, specifically the anode catalyst layer (ACL) on the anode side and the cathode catalyst layer (CCL) on the cathode side, separated by a proton-conductive electrolyte membrane. At each electrode, a sub-reaction of the overall FC reaction occurs. As expressed by (2.2), hydrogen is split into protons and electrons in the hydrogen oxidation reaction (HOR) at the ACL. The protons transfer from the anode to the cathode via the PEM, while the electrons travel through the other layers of the FC and the external circuit. At the CCL, they recombine during the oxygen reduction reaction (ORR) shown in (2.3), generating water and heat. [18]



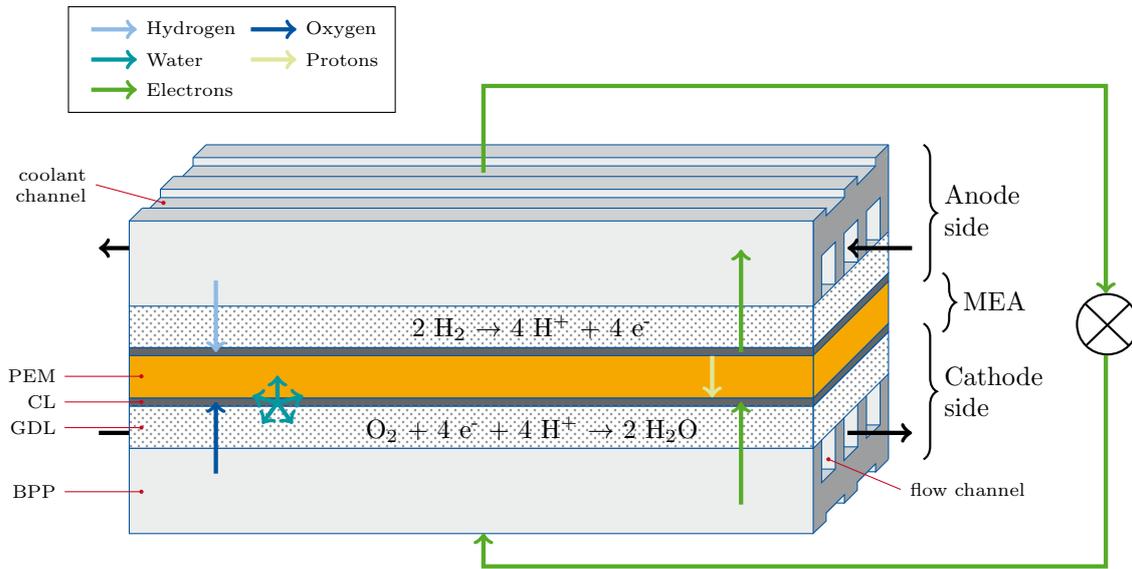
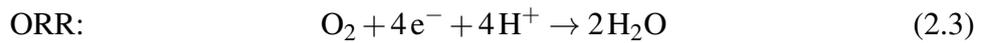


Fig. 2.1: Schematic of the basic structure of a PEMFC. Illustration adapted from [18].



The components of a FC are designed to facilitate these electrochemical reactions efficiently. The ACL and the CCL provide three-phase boundaries that allow for electrons, protons, and reactant gases to participate in the reaction. [43] They are made from an electron-conductive porous medium, like carbon fiber, to maximize surface area. To increase the reaction rate, a catalyst is dispersed in small particles on the electrode structure. The most common catalyst in PEMFCs is platinum, although alternative materials are being explored to substitute the relatively scarce and expensive material [25]. [18] These catalyst particles are partially enveloped by a proton-conducting ionomer, further enhancing the electrochemical reaction within the three-phase boundary. [49]

The CLs and the membrane are often referred to as the membrane electrode assembly (MEA), as the CLs are typically coated on the membrane early in production. The membrane material predominately used in low-temperature PEMFCs is Perfluorosulfonic Acid, also known as NAFION[®]. The PEM enables selective transport of species, creating the interface required for the electrochemical reaction. It conducts protons released through the HOR, while inhibiting electron and reactant transport. [18, 43]

Additional layers are introduced to distribute reactants across the active area and transport the produced water out of the cell. The GDLs facilitate reactant diffusion towards the active sites and conduct the electrons and the heat produced by the reaction. They are typically made from porous media like carbon paper or carbon cloth. [77] A hydrophobic agent such as Polytetrafluoroethylene is often added to reduce the liquid water accumulation within the GDLs, though this reduces porosity, which adversely affects oxygen transport. [79] Sometimes,

the GDL is combined with an MPL, to enhance liquid water transport and lower the contact resistance at its interface with the CL. [18, 68, 79]

Reactants are transported into the cell through flow fields, commonly parallel channel/land flow fields, though alternative channel geometries are also researched. [67, 96] Here, the flow field is considered part of the BPP. Apart from distributing reactants to the cell, the BPP also serves the critical function of distributing the coolant. It further provides the electrical connection to the current collector or adjacent cells. The BPP thus exhibits superior thermal and electrical conductivities and is either made of carbon-based material such as graphite [104] or metals [127]. [18, 49]

2.1.2 Efficiency and performance

To describe the energy being converted into electricity in a FC, we use the Gibbs free energy, which can be defined as the energy available to perform external work, excluding any work done by changes in pressure and/or volume. In a FC, the external work involves moving electrons around an external circuit. When working with chemical reactions, the zero energy point for the chemical energy is typically defined as pure elements in their normal state at standard temperature and pressure (STP) ($T_0 = 25^\circ\text{C}$, $p_0 = 101\,325\text{ Pa}$). For a PEMFC operating at STP, the Gibbs free energy of formation, G_f , of the inputs is zero. The change in Gibbs free energy of formation, ΔG_f , gives us the energy released and is calculated as the difference between the Gibbs free energy of the products and the Gibbs free energy of the inputs or reactants. Considering the reaction for the PEMFC in (2.1), the product is two moles of H_2O , and the reactants are two moles of H_2 and a mole of O_2 . Thus, we have:

$$\Delta \bar{g}_f = 2 \cdot (\bar{g}_f)_{\text{H}_2\text{O}} - 2 \cdot (\bar{g}_f)_{\text{H}_2} - (\bar{g}_f)_{\text{O}_2}$$

where \bar{g}_f is the molar specific Gibbs free energy of formation. This energy varies with temperature and state (liquid or gas). If there are no losses in the FC, i.e. if the process is reversible, all the Gibbs free energy is converted into electrical energy. We will use this to determine the reversible open-circuit voltage (OCV) of a FC.

For the PEMFC, two electrons pass through the external circuit for each water molecule produced and each molecule of hydrogen used. If the system is reversible and E_{rev} is the corresponding voltage of the FC, then the electrical work done moving the charge of two moles of electrons around the circuit is:

$$\begin{aligned} \text{electrical work done} &= \text{charge} \cdot \text{voltage} \\ \Leftrightarrow \Delta \bar{g}_f &= -2 \cdot F \cdot E_{\text{rev}} \end{aligned}$$

where the Faraday constant F is the charge of one mole of electrons. Thus, the reversible OCV of a PEMFC is given by [64]:

$$E_{\text{rev}} = \frac{-\Delta \bar{g}_f}{2 \cdot F} \quad (2.4)$$

At STP, it equals $E_{\text{rev}}^0 = 1.229 \text{ V}$ [62]. This figure assumes no irreversibilities. In practice, however, the voltage is lower. Some irreversibilities occur even when no current is drawn, so the actual OCV of a FC is usually lower than the value given by (2.4). [66]

Since a FC uses materials that are typically burnt to release their energy, it is useful to compare the electrical energy produced with the heat generated by burning the fuel. This heat is described by the change in enthalpy of formation $\Delta\bar{h}_f$. The efficiency of the FC is then defined as follows:

$$\frac{\text{electrical energy produced per mole of fuel}}{-\Delta\bar{h}_f} \quad (2.5)$$

There are two different values for $\Delta\bar{h}_f$: one for the burning of hydrogen with the product water in gaseous form and another for the product water condensed back to liquid. At STP, these are:



The difference of $44.01 \text{ kJ mol}^{-1}$ between these values represents the molar enthalpy of vapourisation of water. The higher value is the higher heating value (HHV), and the lower value is the lower heating value (LHV). In this thesis, efficiency is based on the LHV. We can see that there is a limit to the efficiency, if we define it as in (2.5). The maximum electrical energy available is equal to the change in Gibbs free energy. Thus, the maximum efficiency limit $\eta_{\text{fc,max}}$, sometimes referred to as the thermodynamic efficiency, is [66]:

$$\eta_{\text{fc,max}} = \frac{\Delta\bar{g}_f}{\Delta\bar{h}_f} \quad (2.6)$$

If all the energy from the hydrogen fuel were converted into electrical energy, the reversible OCV of a FC would be:

$$E_{\text{rev,h}}^0 = \frac{-\Delta\bar{h}_f^0}{2 \cdot F} = 1.25 \text{ V} \quad (2.7)$$

This voltage represents a 100% efficient system, based on the LHV. The actual efficiency of the cell is then the actual FC voltage U_{fc} divided by this ideal voltage [66]:

$$\eta_{\text{fc}} = \frac{U_{\text{fc}}}{E_{\text{rev,h}}^0} = \frac{U_{\text{fc}}}{1.25 \text{ V}} \quad (2.8)$$

The Gibbs free energy, and thus the reversible OCV, varies with temperature and reactant pressure and concentration. The variation with temperature is given by:

$$E_{\text{rev}}^\circ = E_{\text{rev}}^0 + (T_{\text{fc}} - T_0) \cdot \left(\frac{\Delta\bar{s}^\circ}{2 \cdot F} \right) \quad (2.9)$$

where T_{fc} is the FC temperature, and E_{rev}° and $\Delta\bar{s}^{\circ}$ are the reversible OCV and the change in entropy at standard pressure, respectively [4]. The Nernst equation expresses how the pressure and concentration of the reactants affect the Gibbs free energy, and thus the reversible OCV. One common form, with reactant and product pressures given in atm and the product water in the form of steam, is:

$$E_{rev} = E_{rev}^{\circ} + \frac{R \cdot T_{fc}}{2 \cdot F} \cdot \ln \left(\frac{p_{H_2} \cdot p_{O_2}^{\frac{1}{2}}}{p_{H_2O}} \right) \quad (2.10)$$

where R is the universal gas constant, p_{H_2} is the partial pressure of hydrogen on the anode side, and p_{O_2} and p_{H_2O} are the partial pressures of oxygen and vapour on the cathode side, respectively. This equation shows that the reversible OCV increases with higher reactant partial pressures and decreases with higher partial pressures of the product [4, 66]. Combining (2.9) and (2.10) yields the following expression for the reversible OCV:

$$E_{rev} = E_{rev}^0 + (T_{fc} - T_0) \cdot \left(\frac{\Delta\bar{s}^{\circ}}{2 \cdot F} \right) + \frac{R \cdot T_{fc}}{2 \cdot F} \cdot \ln \left(\frac{p_{H_2} \cdot p_{O_2}^{\frac{1}{2}}}{p_{H_2O}} \right) \quad (2.11)$$

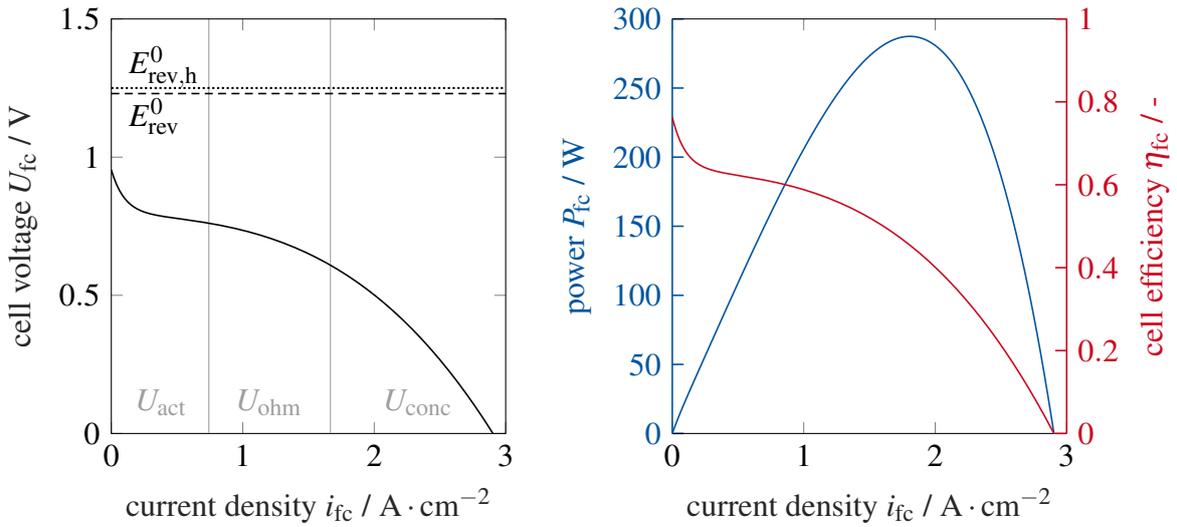


Fig. 2.2: Exemplary polarization curve (left) and corresponding efficiency and power characteristic (right) of a FC. Illustration based on [49].

In practice, the technically usable FC voltage U_{fc} is lower due to losses or irreversibilities, also commonly referred to as overpotentials. They reduce the reversible OCV E_{rev} , as expressed by the following equation:

$$U_{fc} = E_{rev} - U_{act} - U_{ohm} - U_{conc} \quad (2.12)$$

The multiple loss contributions can be divided into three groups: activation losses (U_{act}), ohmic losses (U_{ohm}), and concentration or mass transport losses (U_{conc}). These overpotentials depend

on the operating conditions within the cell and increase with higher current. The proportion of each overpotential varies depending on the FC operating point. This can be illustrated using the polarization curve of the FC, which relates the cell voltage to the area-specific current density. An exemplary polarization curve is depicted in Fig. 2.2 (left). In FC research and application, the polarization curve is a widespread representation of the FC's performance. To enable comparison of the performance of different cell sizes, the cell voltage is represented as a function of current density instead of the current. The current density i_{fc} relates the current I_{fc} to the active cell area A_{fc} as follows:

$$i_{fc} = \frac{I_{fc}}{A_{fc}} \quad (2.13)$$

The polarization curve can be broadly segmented into three characteristic regions, each dominated by a specific type of overpotential.

At lower current densities, the **activation overpotential** U_{act} is by far the largest loss contributor. It arises from the intrinsic resistance to electrochemical reactions at the electrode-electrolyte interface. Consequently, part of the available energy is lost driving the chemical reactions that facilitate electron and ion transfer to and from the electrodes. The extent of these losses also depends on the catalyst used. Due to the significantly slower reduction of O_2 compared to the oxidation of H_2 , the activation losses are primarily attributed to the ORR on the cathode side. Therefore, the voltage drop due to activation loss strongly depends on the oxygen partial pressure p_{O_2} , and additionally on the temperature T_{fc} . [18, 49, 66]

With rising current density, the **ohmic loss contributions** U_{ohm} grow almost linearly and dominate the typical operating range of a FC. Ohmic losses are caused by the finite electronic and protonic conductivity of the layers and the non-ideal contact between the layers. They primarily arise from the resistance of the polymer membrane to the transfer of protons. The membrane's protonic conductivity σ_{mem} is heavily dependent on its humidification state which is characterized by the membrane's water content Λ_{mem} , and the cell temperature T_{fc} . Additional factors include the quality of electrical contacts and the conductivity of the electrodes to transfer electrons. [18, 49]

At high current densities, the **concentration losses** U_{conc} cause a steep nonlinear cell voltage decline. These occur due to the depletion of reactant concentrations at the electrode surfaces as the fuel is used. As reactants are consumed at the electrodes, a continuous supply matching the consumption must be maintained. The achievable current density of the cell is constrained by the effective reactant mass transport for replenishment, which results in a reduction in voltage. Besides that, inadequate air circulation on the cathode side, more specifically nitrogen accumulation after oxygen is consumed, can cause a mass transport problem at high currents as it effectively blocks the oxygen supply. The lack of water removal can also be a cause of mass transport or concentration overvoltage. [18, 49, 66]

Additionally, the measurable OCV is significantly lower than the theoretical reversible OCV. This discrepancy at zero current density is attributed to fuel crossover through the electrolyte,

and, to a lesser extent, internal current flows from electron conduction through the electrolyte. The electrolyte should only transport ions through the cell. However, a certain amount of fuel diffusion and electron flow will always be possible. [66]

Based on these elaborations, we now understand that the FC voltage U_{fc} is not only a function of the load current density i_{fc} , but also dependent on various other operating conditions in the PEMFC, such as the pressure on the anode and cathode side, and the reactant partial pressures. These do not only impact the voltage through the reversible OCV (2.11) but also through the introduced overpotentials. Especially the oxygen partial pressure p_{O_2} in an air-fed PEMFC affects the activation and concentration losses. Additionally, the FC temperature T_{fc} impacts both the reversible OCV and the activation and ohmic losses. The latter is also heavily influenced by the membrane hydration state, characterized by the water content Λ_{mem} , which is defined as the number of water molecules per sulfonic acid group, the functional part of NAFION[®] responsible for proton conduction.

With the resulting voltage characteristic, the efficiency trend can be directly determined using (2.8), and the electrical power output can be calculated with:

$$P_{fc} = U_{fc} \cdot I_{fc} \quad (2.14)$$

Both are illustrated in Fig. 2.2 (right), which demonstrates that there is a trade-off between achieving maximum power output and maintaining high efficiency in a FC.

The mass flows necessary for the reaction can be determined using Faraday's first law, which states that the electrical current in a cell is proportional to the quantity of electrons exchanged by the reacting substance [26]. Therefore, the amount of hydrogen required for the stoichiometric reaction (2 electrons per molecule) within a PEMFC is:

$$\dot{n}_{H_2,sto} = \frac{I_{fc}}{2 \cdot F} \quad (2.15)$$

For the corresponding reaction to occur on the cathode side with oxygen (4 electrons per molecule), the required reactant amount is:

$$\dot{n}_{O_2,sto} = \frac{I_{fc}}{4 \cdot F} \quad (2.16)$$

To ensure high reaction efficiency, the supply of reactants is typically provided in excess of the stoichiometric amount, meaning a greater mass flow is made available than is necessary for the generation of a given electrical current. This excess is quantified by a stoichiometry coefficient λ for both the anode and cathode sides [49, 66]:

$$\dot{n}_{H_2} = \lambda_{H_2} \cdot \dot{n}_{H_2,sto} \quad (2.17)$$

$$\dot{n}_{O_2} = \lambda_{O_2} \cdot \dot{n}_{O_2,sto} \quad (2.18)$$

2.1.3 Degradation

The individual components of the FC are subject to complex aging mechanisms during operation, leading to both reversible and irreversible performance degradation [62]. Understanding these aging mechanisms and developing mitigation strategies is a broad and complex area of research, especially for automotive applications [17]. Comprehensive overviews of various degradation effects are provided in [58, 93], while potential measures are additionally given in [126]. Although detailed degradation knowledge is not the focus of this work, it is important to provide context to two of the most relevant degradation sources, which are qualitatively considered in this thesis.

One of them is reactant starvation, where an inadequate supply of reactants occurs. Hydrogen starvation causes current reversal, leading to corrosion of the carbon support that carries the catalyst and reduces the available ECSA [93]. A decrease in ECSA increases mass transport losses, as less catalyst surface is available for the FC reaction, limiting the available current. Oxygen starvation, though less severe than hydrogen starvation, can still cause accelerated degradation in the electrode and PEMFC performance when it is exposed to oxygen-starved conditions for a few tens of minutes. This results in similar degradation phenomena in the electrocatalyst through cell reversal, leading to a loss of surface area of the cathode platinum particles. Therefore, avoiding oxygen starvation is also crucial to maintain PEMFC performance. [93, 110] Additionally, the rise in concentration losses reduces cell voltage and therefore the stack's power output, adversely affecting efficiency and achievable maximum power.

Another source of degradation phenomena is poor water management. Dry-wet cycling of the membrane causes shrinking and swelling, which may lead to mechanical fatigue, promote crack formation, and as a result increase unwanted reactant crossover. [87]. Very dry membrane states promote side chain scission in NAFION[®], resulting in a loss of sulfonic acid groups and reduced protonic conductivity. [23] On the other hand, liquid water accumulation can result in reactant starvation, leading to carbon corrosion. This surface oxidation decreases the hydrophobicity of the carbon support, increasing mass transport losses within the CL [129].

Most ageing phenomena occur over slow time scales and do not lead to rapid performance losses. The models used within this work are not equipped to quantify this type of degradation. Therefore, mitigating degradation is only addressed in a scope limited to managing oxygen starvation and the membrane hydration state.

2.1.4 Fuel cell stack

From Fig. 2.2 it is evident that the operating voltage and power output of a single cell are multiple orders of magnitude too low for the automotive application. Consequently, several

hundred cells are connected in series to form a FC stack. Due to the serial connection of n_{fc} individual cells, the resulting stack voltage U_{st} is

$$U_{st} = n \cdot U_{fc}, \quad (2.19)$$

while the current remains the same, i.e. $I_{st} = I_{fc}$. Accordingly, the electrical power in (2.14) and material flows required for cell supply in (2.17)–(2.18) all increase by the factor n .

The BPPs electrically connect the poles of two adjacent cells, creating a series circuit. Within the stack, electrons flow directly from the anode of one cell to the cathode of the neighboring cell. Endplates are attached to both ends of the stack, providing the external poles of the stack unit. These poles enable an external load to be powered by the stack voltage U_{st} . [49, 62] To ensure each cell receives the reactants and coolant required for operation, the gases and the liquid coolant are distributed among the cells using manifolds. This setup connects the cells in parallel within the fluidic domain. [18]

The behavior of a single PEMFC is similar to that of an entire stack composed of multiple cells in series, though variations may exist from cell to cell. This work assumes uniformity among the cells within the FC stack. As a consequence, every cell is considered to have the same temperature, which is also equal to the stack temperature, i.e. $T_{st} = T_{fc}$.

2.2 PEM fuel cell systems

To operate the FC stack effectively, maintaining a consistent supply of external medium flows is crucial. This involves providing the reactants oxygen from ambient air and hydrogen, as well as circulating a cooling medium through the cells. Additionally, regulating operating conditions such as pressure, temperature, and humidity is essential to ensure the system runs robustly, with minimal degradation, and high efficiency. This requires the support of peripheral subsystems, also referred to as BoP components, which consume a portion of the electrical power generated by the stack. The combined system of the FC stack and BoP components is referred to as the FCS. The structure and functionality of an automotive FCS are subject of this section.

Four peripheral subsystems are required to operate the stack, as shown schematically in Fig. 2.3: the air supply subsystem, the hydrogen supply subsystem, the thermal management subsystem and the water management subsystem. Additionally, an energy management subsystem is depicted outside the FCS boundary, which is addressed in the subsequent section.

The **air supply subsystem**, also commonly referred to as the air path, supplies ambient air to the cathode side of the stack and handles its exhaust gas steam. Optimal operation requires precise conditioning of the air's humidity, temperature, and pressure. To provide the oxygen for the reaction, air is drawn in from the environment and compressed using an electrically driven compressor. Supplying more oxygen than consumed by the reaction enhances the polarization

2 Preliminaries on fuel cell technology

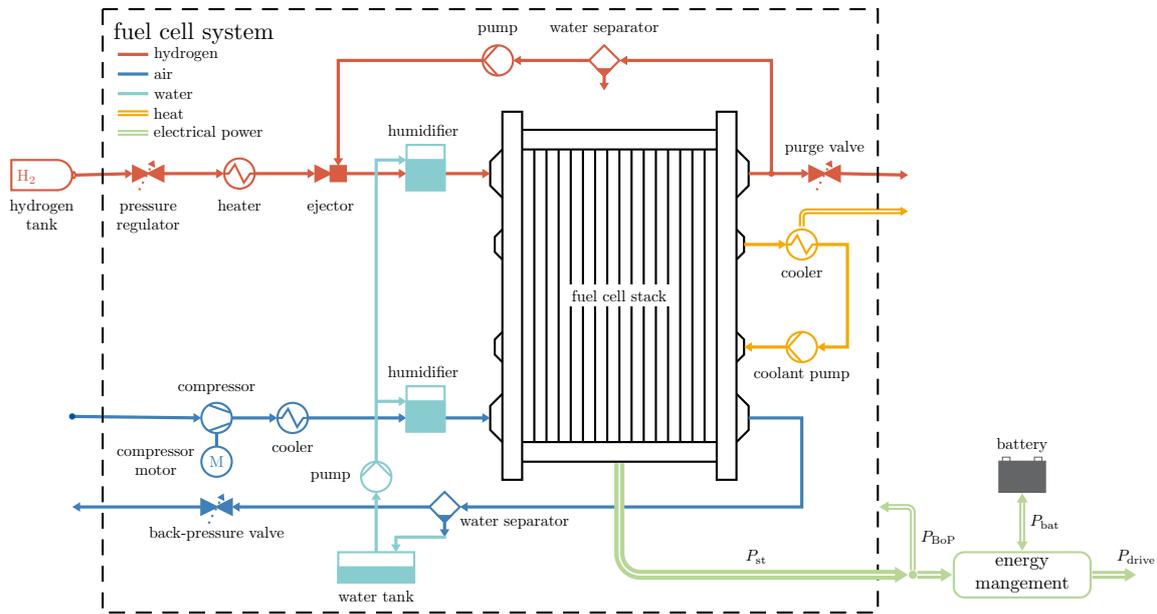


Fig. 2.3: Exemplary illustration of a FCS for automotive applications, adapted from [85].

curve and aids in removing excess water from the cell. After compression, the air is cooled to a stack-compatible temperature level using an intercooler to prevent exceeding the maximum allowable inlet temperature. Excessive temperatures can damage the membrane materials and reduce their proton conductivity. Here, the air is humidified before entering the cathode to prevent the membrane from drying out, utilizing humid exhaust gas from the stack for this purpose. However, commercial systems without external humidification are also available in vehicle applications [56]. A back-pressure valve at the cathode outlet additionally regulates the pressure level in the cathode. Operating pressures of FCSs range from a few millibars to 4 bar above ambient pressure. While increasing air pressure at the cathode improves stack performance, it also increases the power required for the air compressor. With a share of about 80% of the total power of all subsystems, the air compressor significantly impacts the overall efficiency of a FCS [41]. Balancing the power increase of the FC stack and the compressor's power consumption, depending on pressure and oxygen stoichiometry, is crucial. Moreover, in contrast to the hydrogen system, the supply of air to the cathode is directly dependent on the compressor's speed, making the achievable dynamics of a FCS largely determined by the compressor. Additionally, high-pressure FCSs that drive their own compressor exhibit nonminimum phase behavior, where a positive step in requested power causes an inverse response in power output, limiting the closed-loop bandwidth of this loop.

The **hydrogen supply subsystem** provides the stack with the required amount of hydrogen at the appropriate concentration, temperature, and pressure for the electrochemical reaction. Hydrogen is stored in a tank system under pressures up to 700 bar. A pressure regulation system

significantly lowers and regulates the hydrogen pressure to the operating pressure of the anode circuit, typically ranging from 1.5 to 3 bar, ensuring minimal pressure difference between the anode and cathode to avoid damage to the membrane. A heat exchanger, positioned after the pressure regulator, adjusts the gas temperature to the stack temperature. As described in Section 2.1.2, more hydrogen is typically supplied than required for the reaction to ensure high reaction efficiency and prevent hydrogen starvation. To enhance fuel utilization, the anode exhaust gas is usually recirculated. Water vapor diffuses from the cathode to the anode side across the membrane due to a concentration gradient, and the electrolyte is not entirely gas-tight, causing nitrogen to accumulate on the anode side over time. This reduces hydrogen concentration within the loop, inhibiting the reaction and increasing the electrical power required for recirculation. The dilution of hydrogen consequently lowers the FC voltage, therefore its efficiency, as well as FCS efficiency. To address this issue, a purge valve periodically opens to flush out accumulated nitrogen and water vapor from the anode. The so-called purge losses typically account for 1 to 2% of the supplied hydrogen. Furthermore, excessive liquid water in the anode can block the GDLs and flow channels, leading to performance losses. Therefore, a water separator is positioned after the anode outlet to remove excess liquid water. An actively controlled recirculation pump closes the recirculation loop by directing the hydrogen mixture to the fresh hydrogen supply line, where it is mixed before re-entering the stack.

To dissipate the heat produced by the stack and maintain operating temperatures at a safe level, a **thermal management subsystem** is implemented. This subsystem primarily consists of a liquid coolant circuit that connects a radiator with the stack. A pump circulates the coolant through the stack and the coolant circuit. Given the relatively low operating temperatures (around 60°C to 85°C) compared to combustion engines, only a small portion of the heat is released through exhaust gases, with nearly half of the reaction enthalpy having to be absorbed by the coolant. Since the coolant comes into contact with live electrical components, it must be deionized. In many cases, the intercooler is integrated into the stack's thermal circuit to maintain a low temperature difference between the cell walls and the air mass flow entering the cathode. The heat absorbed by the coolant can be utilized energetically, such as for heating the passenger cabin in winter. Any unused heat must be dissipated from the coolant to the environment via a radiator after it has circulated through the stack.

The task of the **water management subsystem** is to maintain proper hydration of the polymer membrane by balancing water production, usage, and removal within the system. This involves components such as humidifiers, which add moisture to the reactant streams, and water separators, which remove excess water from the exhaust streams. As illustrated in Fig. 2.3, we assume external humidification for both anode and cathode inlet streams. The amount of reactant flow and water injected into the anode and cathode flow streams affect the membrane's humidity. Both dry membranes and flooded FCs lead to high polarization losses. As current is drawn from the FC, water molecules are produced in the cathode and dragged from the anode to the cathode by the hydrogen protons through electro-osmotic drag. Increased water concentration in the cathode causes diffusion back to the anode due to the concentration gradient. Various mechanisms can perturb the FC's humidity, such as water generation during load increases,

changes in absolute and partial reactant pressures across the membrane, variations in airflow rate, and fluctuations in stack temperature, which also alter the vapor saturation pressure. These mechanisms indicate strong and nonlinear coupling of the water management subsystem, the reactant flow subsystems, the thermal management subsystem, and the energy management subsystem when it comes to the humidity control task. Maintaining optimal temperature is especially crucial, as it directly impacts the water balance and membrane hydration. Improper humidification control can lead to a 20% to 40% voltage drop [21]. [49, 62, 85]

2.3 Fuel cell vehicle

The powertrain of FC vehicles consists of a hydrogen tank (energy storage), a secondary energy storage device (typically a battery), a FCS (energy converter), multiple voltage converters, an electric motor, a transmission system, and the mechanical drive of the wheels. Thus, FC vehicles are electro-hydrogen hybrids. Their powertrain configurations are typically categorized based on the primary source of driving energy into either a fuel-cell-dominant powertrain or a range-extender powertrain.

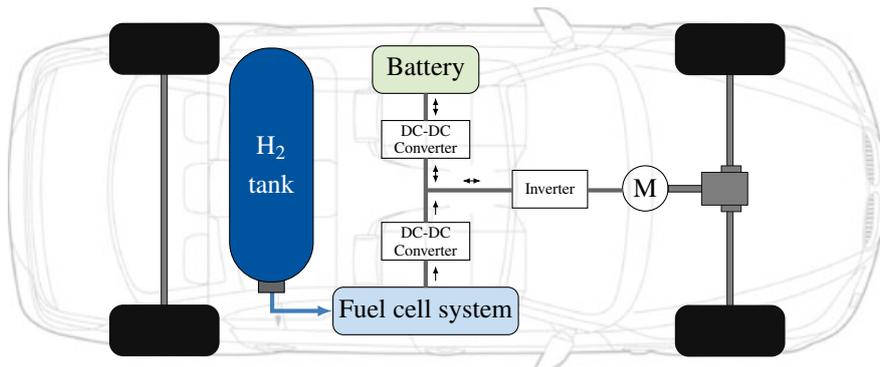


Fig. 2.4: Exemplary illustration of a fuel-cell-dominant power train, adapted from [62].

In a fuel-cell-dominant powertrain (see Fig. 2.4), the driving power demand is met by the FCS, with the battery used mainly for the recuperation of braking energy during deceleration and for power assistance during acceleration. In passenger cars, the FC stack is therefore very powerful (100kW to 150kW), the battery typically has high power density and low capacity (1 kWh to 2kWh), and the hydrogen tank is a high-pressure tank containing several kilograms of H₂ (5kg to 6kg) to achieve ranges of up to 600km. The energy supply is provided by refueling with hydrogen.

In the range-extender powertrain, the driving power demand is met by the battery, while the FCS charges the battery during operation, thereby extending the vehicle's range. Range-extender passenger cars typically have a battery with lower power density and high capacity, a low-power

FC stack (20kW to 30kW), and a small hydrogen pressure tank. Due to the larger battery, range-extender powertrains are often designed as plug-in hybrids, meaning that, in addition to refueling with hydrogen, the vehicle's energy supply is also provided by charging the battery via the power grid.

Voltage converters are necessary in both configurations to connect the different DC voltage levels and to generate AC power for the electric motor. While the FCS is connected to a unidirectional DC-DC converter, the battery is connected to a bidirectional DC-DC converter. [62]

The powertrain considered in this work features a FCS model based on the work by Pukrushpan, Peng, and Stefanopoulou [83]. Their model is derived from the FC stacks used in the FORD[®] P200 FC prototype vehicle [2]. Additionally, we use a battery model based on technical data of the Nickel-metal hydride batteries from the first-generation TOYOTA[®] Mirai, which have a capacity of 1.6kWh. This configuration is considered fuel-cell-dominant, as the batteries are not capable of supplying the full electric power required by the drivetrain.

Apart from supplying and storing energy during acceleration and deceleration, respectively, the batteries also allow for shifts in the load point of the FCS. This means that the amount of power demanded by the vehicle can differ from the power provided by the FCS, enabling it to operate at load points with higher efficiency [64]. With this additional degree of freedom provided by the battery, the operational strategy employed, which coordinates the power distribution between the battery and the FCS, has a significant impact on the overall efficiency and lifespan of the FCS. [49]

Energy management strategies

A hybrid energy supply system, such as in FC vehicles, demands an overarching operational strategy that coordinates the power output of the various energy sources. This control algorithm is commonly referred to as the energy management strategy (EMS). The main aim of the EMS is to minimize the integral fuel consumption over the driving cycles while ensuring that the FC vehicle operates within its operational limits, taking into account various constraints such as power limitations and the battery's state of charge (SoC). Apart from fuel consumption reduction, further objectives might be considered in the energy management problem, such as optimizing for longevity.

An extensive number of scientific publications on EMSs for automotive hybrid systems is available today. In general, EMSs for hybrid electric vehicles (HEVs) can be classified into heuristic and optimization-based methods. These strategies are not limited to FC vehicles but apply to HEVs in general.

Heuristic strategies encompass a group of control algorithms that define a set of rules for controlling power distribution in the hybrid system. These strategies involve rule-based approaches derived from experience, technical understanding of the systems, identified correlations, or

efficiency and cost analyses. Typically, they are implemented using maps, state machines, or fuzzy logic. The primary advantages of heuristic based strategies include their robustness, intuitive implementation, and thus traceability, as well as their low computational effort and low memory usage. However, their effectiveness heavily depends on the formulated rules, which can potentially lead to control behavior that deviates from the optimum to varying degrees. Moreover, these approaches are less attractive from a calibration point of view, as the defined rules are set for a specific vehicle and driving conditions. This imposes a need to reparametrize the strategy for different vehicle configurations, driving cycles, or changes in system behavior.

In optimization-based approaches, power distribution is determined by solving an OCP based on a dynamic model. The goal is to minimize a predefined cost function, which can include factors such as hydrogen consumption or degradation rates of components or systems. These methods are usually more computationally intensive, and depending on the method, may require predictive data about the route. Optimization-based EMSs enable optimal or nearly optimal operation over a driving cycle with respect to the defined cost function. However, the performance of optimization-based controllers heavily depends on the available information about the driving cycle, i.e., on the information related to future torque/power and velocity requests. Generally, the velocity or torque/power request trajectory is an exogenous disturbance input of the HEV control system. Another advantage of these model-based methods is their adaptability and transferability when parameters change, limiting the application effort to reparameterizing the model. Optimization-based strategies can be classified into different categories. Here, we classify them based on their solution method into numerical, analytical, and instantaneous minimization methods.

A well-known optimization-based strategy to the energy management problem of HEVs is dynamic programming (DP). DP is a numerical optimization method that finds a globally optimal solution to OCPs. To achieve this, the entire optimization problem is broken down into subproblems through temporal discretization and solved step by step using Bellman's principle of optimality by minimizing the so-called "cost-to-go". Since DP incorporates the entire optimization horizon, which must be known in advance, it provides the globally optimal solution, apart from interpolation and discretization errors. The result is the same as what would be obtained by trying all possible combinations of control and state variables at each time step and selecting those with the lowest overall value of the objective function. However, the complete information regarding the driving cycle is often not available in real driving conditions. Therefore, deterministic DP and the optimal operating strategies it calculates are typically used offline either as a benchmark for other real-time implementable operating strategies, or for analyzing the optimal operating strategy with the goal of deriving rules for rule-based operating strategies. In this context, causal approaches are also known, which address the prediction of the upcoming driving profile using statistical methods and employ stochastic DP in this regard. Additionally, to represent the optimal solution, the vehicle model used for optimization and the validation model would have to be equivalent, which is an unattainable assumption. Therefore,

the solution obtained by the DP algorithm is rather the lower bound (from the cost perspective) for other real-time control strategies.

Under certain circumstances, the OCP can be solved analytically. The most well-known method for this in the field of operating strategies for hybrid vehicles is Pontryagin's Minimum Principle. Pontryagin's Minimum Principle is a mathematical theorem that formulates the necessary conditions that an optimal solution must satisfy. The basis for formulating these necessary conditions is the so-called Hamiltonian function. For simple OCPs, in the unconstrained case, an analytical solution can be derived from the resulting equations. If the problem under consideration requires the inclusion of additional state constraints, the Hamiltonian function must be extended with a term containing additional optimization variables. This requires a case distinction, making the search for an analytical solution for complex dynamic systems extremely challenging.

In contrast to acausal methods, the following operational strategy approach performs minimization locally, making it entirely causal and suitable for direct implementation in vehicles. The most prominent representative of this category of operational strategies is the equivalent consumption minimization strategy (ECMS). Typically, the absence of information about future driving profiles renders the use of solution methods reliant on such knowledge only marginally practical and feasible. ECMS, however, provides a causal implementation of an EMS. This method determines the power distribution between two energy sources by transforming the global optimization problem into a local optimization problem, dependent solely on currently known parameters. ECMS employs the definition of the Hamiltonian function and its minimization as a necessary condition for an optimal solution. The theoretical approach is equivalent to Pontryagin's Minimum Principle. Unlike DP and other global solution methods, ECMS inevitably leads to suboptimal results but enables practical real-time implementation in the vehicle. The fundamental principle of the ECMS approach is to treat the electric energy storage system merely as an energy buffer, with all energy ultimately originating from the fuel tank. In this context, the energy drawn from the battery during discharge phases is equivalent to a specific amount of chemical energy needed to recharge the battery later. Similarly, charging the battery allows future power withdrawal from the battery. The objective function minimized at every time step is defined as an instantaneous equivalent fuel consumption. It is composed of the actual fuel consumption and the electrical energy usage that is converted into a virtual fuel consumption using an equivalence factor. Although its formulation is straightforward and easy to implement and thus attractive for industry application, ECMS critically relies on the choice of the equivalence factor. Since the perfect tuning of the equivalence factor requires prior knowledge of the driving cycle, adaptive ECMS strategies have been developed to enable the online implementation. These adaptive methods adjust the equivalence factor online to achieve nearly optimal results in real-world applications. There are various adaptation rules used for computing the equivalence factor.

In addition to DP and ECMS, MPC represents an important group of EMSs. MPC can be considered a middle ground between DP and ECMS. Its predictive capability surpasses

instantaneous approaches such as ECMS, which helps to achieve better performance. MPC operates by minimizing a cost function over a finite prediction horizon, accounting for a sequence of future torque/power demand values, thus enabling real-time capability of the control strategy. For MPC control, the future driving information, i.e., the torque/power demand trajectory over a specified time window is required. MPC methods for HEV control differ in the way future driver velocity or torque/power requests are determined. Approaches may range from maintaining the current velocity and corresponding power demand constant over the prediction horizon to assuming precise knowledge of the driving cycle within this horizon. While such MPC approaches cannot be directly implemented in real-world applications due to their idealized assumptions about driving cycles, they are valuable for comparison purposes. Other driving cycle prediction concepts exist. [42, 49, 59]

2.4 Control objectives

The following section aims to answer research question 1 and state the specific performance metrics relevant for this work.

Research question 1:

What are the requirements for control systems in automotive FCSs?

From the elaborations presented in this chapter, we can derive the following threefold top-level goal for controlling automotive FCSs:

1. Dynamic power delivery
2. Avoidance of critical operational limits
3. Efficiency maximization

This can be achieved by actively manipulating the system-level control levers, which entails designing a control strategy that adjusts the FCS inputs to optimize the operating conditions of the stack, such that the threefold goal is satisfied.

With the control levers of the stack and the BoP components introduced in Section 2.2 and the threefold goal, the automotive FCS is a MIMO system. Furthermore, the system is coupled; changing one input of one subsystem affects multiple operating conditions in the system, and one operating condition may be affected by multiple inputs from different subsystems.

For example, water is formed by the FC reaction within the stack and fed into the cells as vapor carried through the inlet gas streams. Thus, both the load and the reactant supply impact the water balance in the FC stack. Another aspect to consider for dynamic power delivery is the nonminimum phase behavior of the compressor-driven FCS. Additionally, the span of different time scales present in the FCS needs to be considered for control design. The time constants

relevant for the control goal range from $\mathcal{O}(10^{-1}\text{s})$ for the manifold dynamics to $\mathcal{O}(10^2\text{s})$ for the temperature dynamics. [18, 48, 85]

In summary, from a control engineering perspective, the automotive FCS is a nonlinear, coupled MIMO system with competing control objectives and constraints, which are necessary for safe transient operation and longevity. Thus, the control algorithm must be able to handle the resulting nonlinear, potentially non-convex, constrained optimization problem, which must be reliably solved within a given sampling time.

Performance metrics

From a practical standpoint, the control concept must compute an efficient power split at vehicle level where efficiency at FCS level translates to an optimal balance between BoP consumption and stack power output, while respecting safety constraints at both system and cell level. This is particularly challenging in dynamic power delivery contexts, as high dynamic operation can push the automotive FCS to its limits.

Power dynamics

The first performance metric quantifies the power dynamics of the closed-loop vehicle response. When a driver wants to accelerate, the system should quickly provide the requested electrical power to the motor. In this work, we evaluate this by the time $t_{90,P_{\text{drive}}}$ it takes the FCS and the battery to deliver 90% of the requested power $P_{\text{drive,ref}}$ [18]:

$$P_{\text{drive}}(t = t_{90,P_{\text{drive}}}) = 0.9 \cdot P_{\text{drive,ref}} \quad (2.20)$$

The corresponding metric for the FCS power is referred to as $t_{90,P_{\text{fcs}}}$. In fuel cell-dominant vehicles with small hybridization batteries, highly dynamic operation of the FCS is required, with power dynamics expected on the scale of one second.

Efficiency

The hydrogen consumption during a load cycle is a key variable at the vehicle level. Reducing consumption increases driving range and lowers operating costs.

The electrical power produced by the stack is given by:

$$P_{\text{st}} = U_{\text{st}} \cdot I_{\text{st}} \quad (2.21)$$

However, driving the compressor can consume up to 20% of the stack power P_{st} [30], reducing the available net power:

$$P_{\text{fcs,net}} = P_{\text{st}} - P_{\text{BoP}} \quad (2.22)$$

where the power consumption of other BoP components is typically an order of magnitude lower than that of the compressor in typical passenger car systems [18].

The total FCS efficiency η_{fcs} for a load cycle lasting from t_0 to t_1 is then defined as [18]:

$$\eta_{\text{fcs}} = \frac{\int_{t_0}^{t_1} P_{\text{fcs,net}} dt}{\Delta \bar{h}_f^0 \cdot \int_{t_0}^{t_1} \dot{n}_{\text{H}_2} dt} \quad (2.23)$$

In this work, we assume the consumed hydrogen molar flow to match the stoichiometric amount needed for the reaction. Thus, \dot{n}_{H_2} can be calculated using (2.15). Consequently, we use a hydrogen excess ratio of $\lambda_{\text{H}_2} = 1$ and neglect minor fuel losses in the system. In addition to η_{fcs} , we quantify the fuel economy of the FC vehicle based on the total equivalent hydrogen consumption during operation $m_{\text{H}_2,\text{eq}}$. This is composed of direct hydrogen consumption by the FCS, $m_{\text{H}_2,\text{fcs}}$, and indirect equivalent hydrogen consumption by the battery, $m_{\text{H}_2,\text{bat}}$ [129]:

$$m_{\text{H}_2,\text{eq}} = m_{\text{H}_2,\text{fcs}} + k \cdot m_{\text{H}_2,\text{bat}} \quad (2.24)$$

where k denotes the correction coefficient and $m_{\text{H}_2,\text{fcs}}$ is calculated from:

$$m_{\text{H}_2,\text{fcs}} = \int_{t_0}^{t_1} \dot{m}_{\text{H}_2} dt = \int_{t_0}^{t_1} M_{\text{H}_2} \cdot \dot{n}_{\text{H}_2,\text{sto}} dt \quad (2.25)$$

with M_{H_2} being the molar mass of hydrogen. Relating the total energy delivered by the vehicle to the total equivalent energy consumed yields an expression for the FCHEV efficiency η_{fchev} :

$$\eta_{\text{fchev}} = \frac{\int_{t_0}^{t_1} P_{\text{drive}} dt}{E_{\text{lhv,eq}}} \quad \text{with} \quad E_{\text{lhv,eq}} = \frac{\Delta \bar{h}_f^0 \cdot m_{\text{H}_2,\text{eq}}}{M_{\text{H}_2}} \quad (2.26)$$

More details on the computation of η_{fchev} is given in A.1.

Safety-critical constraints

Another performance metric is adherence to safety-critical constraints. One key constraint is maintaining a lower bound on the oxygen stoichiometry, defined in (2.18) and also known as oxygen excess ratio (OER). In this work, the minimum OER value of 1.5 must be respected in both steady-state and dynamic operation. Furthermore, we demand that the air-path subsystem operates within the choke and surge boundaries of the compressor. The choke boundary marks the maximum possible mass flow that the compressor can handle, while the surge line represents the minimum operable flow rate for a given pressure ratio and compressor speed.

Additionally, maintaining membrane humidification is critical to avoid drying out. A well-established parameter to quantify membrane hydration is the membrane water content Λ_{mem} ,

which describes the number of water molecules sorbed in the membrane per sulfonic acid group [45, 105]:

$$\Lambda_{\text{mem}} = \frac{N_{\text{H}_2\text{O}}}{N_{\text{SO}_3\text{H}}} \quad (2.27)$$

According to the literature [45, 105], a value of $7 \leq \Lambda_{\text{mem}} \leq 14$ is necessary to prevent membrane drying.

3 Control preliminaries

This chapter covers the essential control concepts required for understanding this thesis. We particularly focus on methodologies for the design and development of model predictive control (MPC) algorithms. First, we formulate the general control problem in continuous time, which typically arises in MPC design, and briefly outline the solution families for such OCPs. This is followed by discretization methods for transcribing the infinite-dimensional OCP into a finite-dimensional nonlinear program (NLP), pertinent to direct optimal control methods. Subsequently, we elaborate on algorithmic techniques for numerically solving the resulting NLPs. Additionally, we address concepts for the practical application of MPC, such as state estimation and zero steady-state offset tracking. Finally, the chapter concludes with a brief introduction to the basic principles of multi-objective Bayesian optimization (MOBO).

3.1 Optimal control

The continuous-time OCP that typically arises in MPC is formulated as follows:

$$\begin{aligned} & \underset{x(\cdot), u(\cdot)}{\text{minimize}} && \int_{t_0}^{t_f} \ell(x(t), u(t)) dt + \ell_f(x(t_f)) \end{aligned} \quad (3.1a)$$

$$\text{subject to} \quad x(t_0) = \hat{x}_0, \quad (3.1b)$$

$$\dot{x}(t) = f(x(t), u(t)), \quad t \in [t_0, t_f], \quad (3.1c)$$

$$h(x(t), u(t)) \leq 0, \quad t \in [t_0, t_f], \quad (3.1d)$$

$$h_f(x(t_f)) \leq 0 \quad (3.1e)$$

Here, $x(t) \in \mathbb{R}^{n_x}$ denotes the differential states, $\dot{x}(t)$ represents the differential state derivatives, and $u(t) \in \mathbb{R}^{n_u}$ denotes the control inputs at time t . The objective in (3.1a) consists of the stage cost function $\ell(\cdot) : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ and the terminal cost function $\ell_f(\cdot) : \mathbb{R}^{n_x} \rightarrow \mathbb{R}$. These functions are typically chosen by the control designer to achieve the desired control performance. The optimization problem depends on the actual or estimated current state of the controlled system, \hat{x}_0 , through the initial value condition in (3.1b). The equality constraints in (3.1c) arise from the system dynamics model $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, described here by an explicit ordinary differential equation (ODE) system. Additionally, the inequality constraints $h : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_{\text{ineq}}}$ (3.1d) and $h_f : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_{\text{ineq},f}}$ (3.1e) denote the path constraints and

terminal constraints, respectively. The decision variables are the state and control trajectories $x(\cdot)$ and $u(\cdot)$, which are functions of continuous time. Our primary focus is on the solution $u^*(t, \hat{x}_0) \forall t \in [t_0, t_f]$, representing the locally optimal control trajectory to be applied as a function of the current system state \hat{x}_0 . The continuous-time OCP is an infinite-dimensional optimization problem. It involves infinite-dimensional decision variables and an infinite number of constraints, as the time index t runs through infinitely many values $t \in [t_0, t_f]$, where t_f is the prediction time length. [88, 118]

3.1.1 Solution methods for optimal control

There are three basic approaches to numerically solve continuous-time OCPs:

- DP is based on Bellman’s principle of optimality, which suggests constructing the optimal policy by recursively minimizing the cost-to-go function backwards in time. This process leads to the Hamilton-Jacobi-Bellman equation, a partial differential equation in state space. This equation is the continuous-time analog of the DP algorithm in discrete time. The main advantage of DP is that it yields a globally optimal solution, if one can be found. However, despite the existence of numerical methods for approximating solutions, this approach severely suffers from Bellman’s curse of dimensionality and is thus restricted to small state dimensions. [10]
- Indirect methods for optimal control derive the necessary conditions for optimality for an infinite-dimensional solution to the OCP (3.1), often based on the Pontryagin maximum principle, which results in a boundary value problem (BVP) that must be solved numerically. Once the functions $x(\cdot)$ and $u(\cdot)$ are found, they need to be sampled at the control system’s rate. These methods are often described as ‘first optimize, then discretize’. The two major drawbacks of indirect methods are the complexity of solving a BVP, which can be challenging in general cases, and the difficulty of handling arbitrary path constraints. [15]
- Direct methods for optimal control, also known as ‘first discretize, then optimize’ approaches, solve a finite-dimensional NLP problem that corresponds to a discrete approximation of the original infinite-dimensional OCP. This involves introducing a grid on the time interval $[t_0, t_f]$ divided into N intervals, within which the continuous-time dynamics are simulated. This discretization transforms the problem into a finite-dimensional one with decision variables $w \in \mathbb{R}^{n_w}$ instead of function space, solvable by numerical optimization methods. It can be shown that as the discretization becomes finer, this direct approach converges to the original continuous-time solution.

One of the most significant advantages of direct methods over indirect ones is their ability to easily treat all types of constraints, such as the inequality path constraints in the formulation above. For solving constrained OCPs in real-world applications, direct methods are currently

the most widely successfully used techniques. Consequently, the focus for the rest of this thesis will be on direct methods.

All direct methods involve a finite-dimensional parameterization of the control trajectory, but they differ significantly in how they handle the state trajectory. Three commonly used direct methods are single shooting, multiple shooting, and direct collocation. [12, 20, 33] The following sections elaborate on the first two methods.

3.1.2 Single shooting

In this context, the term ‘shooting’ refers to the simulation of a dynamical system. The idea in single shooting is to discretize the continuous control input $u(\cdot)$ on a grid of N intervals, obtaining the finite-dimensional control trajectory $u_{0:N-1}^\top = [u_0^\top, u_1^\top, \dots, u_{N-1}^\top]$, where $u_k \in \mathbb{R}^{n_u}$. Given this control trajectory, we simulate the state trajectories on all grid intervals at once. Thus, the decision variables in single shooting are solely the control inputs u_0, \dots, u_{N-1} , which are used to recursively eliminate the states x_1, \dots, x_N , as follows:

$$x_0(\hat{x}_0, u_{0:N-1}) = \hat{x}_0 \quad (3.2)$$

$$x_{k+1}(\hat{x}_0, u_{0:N-1}) = f_d(x_k(\hat{x}_0, u_{0:N-1}), u_k), \quad k = 0, \dots, N-1, \quad (3.3)$$

The control inputs are typically assumed to be piecewise constant, which aligns with practical implementations where zero-order hold signals are used for analog-to-digital conversion. NLPs obtained from single shooting read as:

$$\begin{aligned} & \underset{u_{0:N-1}}{\text{minimize}} && \sum_{k=0}^{N-1} \ell(x_k(\hat{x}_0, u_{0:N-1}), u_k) + \ell_f(x_N(\hat{x}_0, u_{0:N-1})) \\ & \text{subject to} && h(x_k(\hat{x}_0, u_{0:N-1}), u_k) \leq 0, \quad k = 0, \dots, N-1, \\ & && h_f(x_N(\hat{x}_0, u_{0:N-1})) \leq 0 \end{aligned}$$

Since nearly all states are eliminated by forward simulation, the variable space of the NLP is reduced from $(N+1) \cdot n_x + N \cdot n_u$ to $N \cdot n_u$. Note that the system dynamics in (3.3) are implicitly satisfied by definition and are no longer constraints in the optimization problem. This reduced problem can now be addressed using Newton-type methods. Single shooting is also referred to as the sequential approach because the simulation problem and optimization problem are solved sequentially, one after the other. [33, 118]

3.1.3 Multiple shooting

In contrast to single shooting, multiple shooting does not only introduce a finite-dimensional discretization for the control input $u(\cdot)$, but also for the states $x(\cdot)$, typically on the same grid.

The NLP arising from a discretization of an OCP based on multiple shooting typically reads as:

$$\underset{x_{0:N}, u_{0:N-1}}{\text{minimize}} \quad \sum_{k=0}^{N-1} \ell(x_k, u_k) + \ell_f(x_N) \quad (3.4a)$$

$$\text{subject to} \quad x_0 = \hat{x}_0, \quad (3.4b)$$

$$x_{k+1} = f_d(x_k, u_k), \quad k = 0, \dots, N-1, \quad (3.4c)$$

$$h(x_k, u_k) \leq 0, \quad k = 0, \dots, N-1, \quad (3.4d)$$

$$h_f(x_N) \leq 0, \quad (3.4e)$$

with state trajectory $x_{0:N}^\top = [x_0^\top, \dots, x_N^\top]$ where $x_k \in \mathbb{R}^{n_x}$. Thus, the state trajectories are additional optimization variables, and ensuring their continuity, as enforced by (3.4c), is left to the NLP solver. When addressing this problem directly in a nonlinear optimization framework, the variables in $x_{0:N}$ and $u_{0:N-1}$ generally represent a feasible state and control trajectory only upon convergence. In multiple shooting, simulation and optimization are interwoven, so this approach is often referred to as simultaneous approach. [33, 118]

While solving the reduced NLP from single shooting might seem advantageous, the sparsity structure from multiple shooting discretization can be exploited efficiently by dedicated solver techniques. In fact, it has been shown that the cost per Newton iteration can be made equal for both approaches due to the inherent sparsity. Additionally, the main advantages of multiple shooting over single shooting are greater flexibility in problem initialization and improved convergence properties, especially for unstable systems. [3]

3.1.4 Numerical simulation

Direct optimal control methods rely on an accurate discretization of the original continuous-time OCP, resulting in a finite-dimensional NLP. In the context of NMPC, solving these NLPs under strict timing constraints can be challenging [33]. One potential bottleneck is the numerical simulation of the nonlinear dynamics. [88] In this thesis, we focus on deterministic nonlinear dynamical systems from a state-space perspective. To simulate these systems on a digital computer, we need to evaluate the evolution of the state at discrete points in time $t_k = k \cdot T_s$, for $k = 0, 1, \dots$, where T_s denotes the constant step length between t_k and t_{k+1} . Thus, a discrete-time formulation of the continuous-time dynamics is required.

Linear continuous-time dynamics in state-space form are represented by the following equations:

$$\begin{aligned} \dot{x}(t) &= A \cdot x(t) + B \cdot u(t), \quad \forall t \in [0, \infty) \\ y(t) &= C \cdot x(t) + D \cdot u(t), \quad \forall t \in [0, \infty) \end{aligned}$$

where A is the system matrix, B the input matrix, C the output matrix, and D the feedthrough matrix. Assuming the input is piecewise constant, the matrix exponential allows us to transform from continuous-time to discrete-time state-space model as follows:

$$\begin{aligned}x_{k+1} &= A_d \cdot x_k + B_d \cdot u_k, \quad k = 0, 1, \dots, \\y_k &= C_d \cdot x_k + D_d \cdot u_k, \quad k = 0, 1, \dots,\end{aligned}$$

with

$$\begin{aligned}A_d &= e^{A \cdot T_s}, \\B_d &= \int_0^{T_s} e^{A \cdot (T_s - \tau)} \cdot B \, d\tau = (A_d - I) \cdot A_d^{-1} \cdot B, \\C_d &= C, \\D_d &= D\end{aligned} \tag{3.5}$$

However, only a few systems can be realistically modeled by linear equations. Often, linear dynamical systems arise as a linearization of some nonlinear system.

A basic nonlinear dynamic system is characterized by an initial value problem (IVP) of an ODE as follows:

$$\dot{x}(t) = f(x(t), u(t)), \quad \forall t \in [0, T], \quad x(0) = \hat{x}_0,$$

with $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ and initial value $\hat{x}_0 \in \mathbb{R}^{n_x}$. Here, $T > 0$ denotes a finite integration interval over which the state evolution is considered. For the existence and uniqueness of a solution, we assume that f is continuously differentiable and T sufficiently small. The latter assumption is satisfied by the sampling time T_s , which is the time interval of interest in the discretization context at hand. Unlike linear systems, there is no general closed-form expression to obtain discrete-time dynamics from continuous-time dynamics for nonlinear systems. Instead, we rely on numerical integration routines to compute the simulation of differential equations, providing a numerical approximation of this solution at discrete time points. [118] To this end, the sampling interval $[0, T_s]$ is subdivided into $N_{\text{int}} \in \mathbb{N}$ subintervals of equal length $T_{\text{int}} = T_s/N_{\text{int}}$. The resulting step size T_{int} defines the integration grid used for the numerical propagation of the continuous-time dynamics.

Runge-Kutta (RK) methods are the most important class of single-step integration schemes, commonly employed in MPC. They form a family of integrators of different orders and come in both explicit and implicit forms. The order of an integrator indicates, for example, that if an integrator has order p and the step size T_{int} is reduced by the factor of 2, the integration error decreases by 2^p . RK methods can be classified by the number of stages s . Starting from the sampled state x_k , the numerical integration over one sampling interval is initialized as $\tilde{x}_0 = x_k$. A general s -stage RK method is then applied to propagate the state within the sampling interval according to

$$\tilde{x}_{j+1} = \tilde{x}_j + T_{\text{int}} \cdot \sum_{i=1}^s b_i \cdot \ell_i, \quad j = 0, \dots, N_{\text{int}} - 1,$$

where the intermediate stage evaluations ℓ_i are defined by

$$\ell_i = f \left(t_k + j \cdot T_{\text{int}} + c_i \cdot T_{\text{int}}, \tilde{x}_j + T_{\text{int}} \cdot \sum_{r=1}^s a_{ir} \cdot \ell_r, u_k \right), \quad i = 1, \dots, s.$$

After N_{int} integration steps, the state at the next sampling instant is obtained as $x_{k+1} := \tilde{x}_{N_{\text{int}}}$. The coefficients a_{ir} , b_i , and c_i can be organized in a compact form known as the Butcher tableau. Within the class of RK methods, explicit Runge-Kutta (ERK) schemes evaluate the intermediate stage values sequentially, resulting in a strictly lower triangular Butcher tableau. In contrast, implicit Runge-Kutta (IRK) schemes are characterized by a dense Butcher tableau, which requires the solution of a system of equations for the intermediate stage values. The advantage of IRK methods lies in their superior stability properties, particularly for stiff dynamical systems, albeit at the cost of increased computational effort [92].

3.2 Numerical optimization

Numerical optimization methods are fundamental to every MPC implementation, and the algorithmic choices significantly impact the reliability and performance of the resulting MPC controller [92]. Thus, this section aims to introduce several key concepts from the field of numerical optimization relevant to solving the optimization problems discussed in this thesis.

Once the finite-dimensional optimization problem is formulated, it needs to be solved. Here, we consider a compact NLP formulation that encapsulates the NLPs in (3.4), with all optimization variables collected into the variable $w \in \mathbb{R}^{n_w}$ and parameter $\hat{x}_0 \in \mathbb{R}^{n_x}$. This standard NLP form reads:

$$\underset{w}{\text{minimize}} \quad J(w) \tag{3.6a}$$

$$\text{subject to} \quad c(\hat{x}_0, w) = 0, \tag{3.6b}$$

$$g(w) \leq 0, \tag{3.6c}$$

where $J: \mathbb{R}^{n_w} \rightarrow \mathbb{R}$, $c: \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_{\text{eq}}}$, and $g: \mathbb{R}^{n_w} \rightarrow \mathbb{R}^{n_{\text{ineq}}}$, which we assume to be at least twice continuously differentiable in all arguments. [88]

The goal of the optimization procedure is to reliably and efficiently find an approximation of the solution $w^*(\hat{x}_0)$ for a given value of \hat{x}_0 . In this thesis, we solve optimization problems such as the NLP in (3.6) using Newton-type methods. A Newton-type algorithm proceeds to find solutions to the set of equations given by the necessary first-order conditions for optimality. The first order necessary conditions for optimality of the above optimization problem are known as the Karush-Kuhn-Tucker (KKT) conditions.

Before stating these, we need to make some definitions. At a feasible point w , an inequality constraint with index $i \in \{1, \dots, n_{\text{ineq}}\}$ is considered active if and only if $g_i(w) = 0$. Further,

the linear independence constraint qualification (LICQ) is satisfied if and only if the gradients of all active inequalities, $\nabla_w g_i(w) \in \mathbb{R}^{n_w}$, and the gradients of the equality constraints, $\nabla_w c_j(w) \in \mathbb{R}^{n_w}$ for $j \in \{1, \dots, n_{\text{eq}}\}$, form a linearly independent set of vectors. Next, we define the Lagrangian function corresponding to the NLP (3.6). To this end, we introduce Lagrange multipliers $\nu \in \mathbb{R}^{n_{\text{eq}}}$ and $\mu \in \mathbb{R}^{n_{\text{ineq}}}$ associated with the equality and inequality constraints, respectively. The Lagrangian is then defined as:

$$\mathcal{L}(w, \nu, \mu) := J(w) + \nu^\top \cdot c(\hat{x}_0, w) + \mu^\top \cdot g(w)$$

The Lagrangian can be interpreted as an extension to the cost function where we linearly penalize constraint deviations.

With that, the KKT conditions are defined as follows: If w^* is a local minimizer of the optimization problem defined in (3.6) and if LICQ holds at w^* , then there exist multiplier vectors ν^* and μ^* such that

$$\nabla_w \mathcal{L}(w^*, \nu^*, \mu^*) = 0 \quad (3.7)$$

$$c(\hat{x}_0, w^*) = 0 \quad (3.8)$$

$$0 \geq g(w^*) \perp \mu^* \geq 0 \quad (3.9)$$

where the last condition, known as the complementarity condition, states not only that all components of $g(w^*)$ are non-positive and all components of μ^* are non-negative, but also that the products $\mu_i^* \cdot g_i(w^*)$ are zero for each $i \in \{1, \dots, n_{\text{ineq}}\}$. The latter implies that the two vectors are orthogonal. Thus, each pair $(g_i(w^*), \mu_i^*) \in \mathbb{R}^2$ must lie on a non-smooth, L-shaped subset of \mathbb{R}^2 that comprises only the negative x-axis, the positive y-axis, and the origin. Any triple (w^*, ν^*, μ^*) satisfying the KKT conditions (3.7)–(3.9) and the LICQ is called a KKT point, regardless of local optimality. In general, the existence of multipliers that satisfy the KKT conditions is just a necessary condition for local optimality of a point w^* at which LICQ holds. Only in the special case of a convex optimization problem do the KKT conditions serve as both a necessary and sufficient condition for global optimality. For the general case, additional conditions on the second-order derivatives of the problem functions are needed to formulate sufficient conditions for local optimality. [92]

The Newton-type algorithm proceeds to find solutions to the KKT complementarity system (3.7–3.9) by iterating on the triple

$$\xi^k := (w^k, \nu^k, \mu^k)$$

at iteration k until a solution is reached, up to a specified accuracy. Two major approaches exist: sequential quadratic programming (SQP) and interior point methods, which differ in their treatment of inequalities. [118]

3.2.1 Interior point methods

Interior point or barrier methods are among the most powerful algorithms for nonlinear optimization, particularly suited for large-scale programming problems [75]. Nonlinear interior point methods address the non-smoothness of the KKT conditions (3.7)–(3.9) by formulating an approximate, smooth root-finding problem. This smooth problem represents the necessary optimality conditions of an equality-constrained optimization problem that approximates the original problem as follows:

$$\nabla_w \mathcal{L}(w, v, \mu) = 0 \quad (3.10)$$

$$c(\hat{x}_0, w) = 0 \quad (3.11)$$

$$g_i(w) \cdot \mu_i - \tau = 0 \quad i = 1, \dots, n_{\text{ineq}} \quad (3.12)$$

The key idea is to approximate the non-smooth complementarity condition in (3.9) by the smooth version $\mu_i \cdot g_i(w) = \tau$ for $i = 1, \dots, n_{\text{ineq}}$, where $\tau > 0$ is referred to as the barrier parameter. The resulting perturbed KKT system is then solved for a sequence of positive barrier parameter values for which $\tau \rightarrow 0$. Solving this smooth set of nonlinear equations with a Newton-type method results in structured linear systems that can be efficiently solved using sparsity-exploiting linear algebra solvers. By adjusting the barrier parameter in each iteration such that

$$\lim_{j \rightarrow \infty} \tau_j = 0$$

the method ultimately converges (under certain conditions) to (w^*, v^*, μ^*) . [88, 92, 118] The set of solutions $(\bar{x}(\tau), \bar{v}(\tau), \bar{\mu}(\tau))$ for $\tau \in (0, \infty)$ is called the central path.

The interior point KKT conditions (3.10)–(3.12) are equivalent to the first-order necessary conditions of the so-called barrier problem:

$$\underset{w}{\text{minimize}} \quad J(w) + \sum_{i=1}^{n_{\text{eq}}} v_i \cdot c_i(\hat{x}_0, w) - \tau \cdot \sum_{i=1}^{n_{\text{ineq}}} \log(-g_i(w)) \quad (3.13a)$$

$$\text{subject to} \quad c_i(\hat{x}_0, w) = 0 \quad i = 1, \dots, n_{\text{eq}} \quad (3.13b)$$

The first-order necessary conditions to (3.13) are

$$\nabla J(w) + \sum_{i=1}^{n_{\text{eq}}} \nabla c_i(\hat{x}_0, w) \cdot \lambda_i + \tau \cdot \sum_{i=1}^{n_{\text{ineq}}} \frac{1}{-g_i(w)} \cdot \nabla g_i(w) = 0$$

$$c_i(\hat{x}_0, w) = 0 \quad i = 1, \dots, n_{\text{eq}}$$

and with $\mu_i = \tau / -g_i(w)$, they are equivalent to the interior point KKT conditions. [89]

For further details on this family of optimization algorithms, we refer the reader to [11, 36]. A widely used general-purpose interior point method for sparse NLPs is implemented in the open-source code IPOPT [123].

3.2.2 Sequential quadratic programming

Instead of the smoothing technique used in interior point methods, SQP relies on sequentially approximating the NLP by a quadratic subproblem. Unlike interior point techniques, which do not easily benefit from a good initial guess [101], the natural warm-starting capabilities of SQP methods make them particularly popular for embedded optimization [33]. Therefore, this thesis focuses on the family of SQP methods for nonlinear programming.

An SQP method applied to the general NLP in (3.6) successively linearizes the KKT conditions around the current iterate ξ^k and solves the following quadratic programming (QP) subproblem in each iteration k :

$$\underset{\Delta w}{\text{minimize}} \quad \frac{1}{2} \cdot \Delta w^\top \cdot H \cdot \Delta w + \Delta_w J(w^k)^\top \cdot \Delta w \quad (3.14a)$$

$$\text{subject to} \quad c(w^k, \hat{x}_0) + c_w(w^k, \hat{x}_0) \cdot \Delta w = 0 \quad | \nu_{\text{QP}}, \quad (3.14b)$$

$$g(w^k) + g_w(w^k) \cdot \Delta w \leq 0 \quad | \mu_{\text{QP}}, \quad (3.14c)$$

where $c_w := \frac{\partial c}{\partial w}(w^k)$ and $g_w := \frac{\partial g}{\partial w}(w^k)$ denote the constraint Jacobian matrices. In the case of an exact SQP method, the matrix $H = \nabla_w^2 \mathcal{L}(w, \nu, \mu)$ represents the Hessian of the Lagrangian.

The current primal and dual variables are then updated based on the solution of the QP (3.14) as follows:

$$\begin{aligned} w^{k+1} &= w^k + \alpha \cdot \Delta w \\ \nu^{k+1} &= \nu^k + \alpha \cdot (\nu_{\text{QP}} - \nu^k) \\ \mu^{k+1} &= \mu^k + \alpha \cdot (\mu_{\text{QP}} - \mu^k), \end{aligned}$$

where the step size α can be determined using a globalization strategy [24, 75]. Due to the real-time constraints in embedded optimization, we omit globalization techniques in this thesis and instead utilize constant step sizes.

Hessian approximation techniques

Exact Hessian-based SQP methods rely on computing the second order derivatives, $H = \nabla_w^2 \mathcal{L}(w, \nu, \mu)$, which can be computationally expensive. As a result, many SQP method variants use Hessian approximation techniques. Although these methods typically converge more slowly, they offer considerably reduced computational cost per iteration.

One prominent class of such SQP algorithms employs exact constraint Jacobians but approximates the Hessian matrix H using updated formulas based on first-order derivative information [82]. These are known as Quasi-Newton methods, with the Broyden-Fletcher-Goldfarb-Shanno (BFGS) update being the most widely used [75]. However, the potential for a large number of Hessian updates per iteration makes these methods less suitable for embedded optimization [118].

A popular Hessian approximation technique for embedded optimization is the Generalized Gauss-Newton (GGN) method [13]. This method is applicable to optimization problems with a (nonlinear) least squares type objective $J(w) = \frac{1}{2} \cdot \|r(w)\|_2^2$ which is common in OCPs. The function $r(\cdot)$ is often referred to as the residual function, since it typically represents an error or misfit. [88] The GGN method linearizes the objective ‘inside the norm’, resulting in the following objective for the QP subproblem (3.14):

$$\begin{aligned} J_{\text{GN}}(w^k) &= \frac{1}{2} \cdot \|r(w^k) + \nabla_w r(w^k) \cdot \Delta w\|_2^2 \\ &= \frac{1}{2} \cdot \Delta w^\top \cdot \nabla_w r(w^k)^\top \cdot \nabla_w r(w^k) \cdot \Delta w + r(w^k)^\top \cdot \nabla_w r(w^k) \cdot \Delta w + \text{constant} \\ &= \frac{1}{2} \cdot \Delta w^\top \cdot H_{\text{GN}}(w^k) \cdot \Delta w + \nabla_w J(w^k)^\top \cdot \Delta w + \text{constant} \end{aligned}$$

where $r_w(w^k) := \frac{\partial r}{\partial w}(w^k)$ denotes the Jacobian of the residual at the current iterate w^k . Thus, the GGN method approximates the Hessian of the Lagrangian as:

$$\nabla_w^2 \mathcal{L}(w^k, v^k, \mu^k) \approx H_{\text{GN}}(w^k) = \nabla_w r(w^k)^\top \cdot \nabla_w r(w^k) \quad (3.15)$$

Note that the Gauss-Newton (GN) Hessian approximation $H_{\text{GN}}(w^k)$ does not depend on the dual variables, hence it is also called a multiplier-free approach. This is a significant advantage in practice, as it means that its performance does not depend on a good initial guess for the Lagrange multipliers. The GGN method provides a good Hessian approximation as long as the residual evaluations $r(\cdot)$ remain small or the problem functions have small second order derivatives. This is typically the case in parameter estimation problems or tracking formulations of optimal control, two important application examples. [88] The GGN method usually results in a (fast) linear convergence rate. Besides being a multiplier-free approach and cheap to compute compared to an exact Hessian, another advantage is that the GN Hessian is positive definite by definition, if $\frac{\partial r}{\partial w}(w^k)$ has full rank, and therefore does not require convexification techniques for regularization.

3.2.3 Quadratic programming

If in the general NLP formulation (3.6) the constraints c, g are affine, and the objective J is a linear-quadratic function, the resulting problem simplifies to a QP of the following form:

$$\underset{w}{\text{minimize}} \quad \frac{1}{2} \cdot w^\top \cdot H \cdot w + w^\top \cdot q \quad (3.16a)$$

$$\text{subject to} \quad A_{\text{eq}} \cdot w + b_{\text{eq}} = 0, \quad (3.16b)$$

$$A_{\text{ineq}} \cdot w + b_{\text{ineq}} \leq 0. \quad (3.16c)$$

If the Hessian matrix H is positive semi-definite, i.e. if $\forall w \in \mathbb{R}^{n_w} : w^\top \cdot H \cdot w \geq 0$, the QP (3.16) is a convex QP. Convex QPs possess unique minimizers and can be solved to global optimality in polynomial time. Conversely, non-convex QPs may exhibit multiple local minima.

Numerical algorithms for solving convex QPs, which also need to be solved as subproblems in the SQP approach for solving NLPs, can be categorized into three primary classes: active-set methods, interior-point methods, and first-order methods.

Active-set methods

Assuming a convex QP, the following KKT conditions are necessary and sufficient for global optimality:

$$\begin{aligned} H \cdot w^* + q + A_{\text{eq}}^\top \cdot v^* + A_{\text{ineq}}^\top \cdot \mu^* &= 0 \\ A_{\text{eq}} \cdot w^* + b_{\text{eq}} &= 0 \\ A_{\text{ineq}} \cdot w^* + b_{\text{ineq}} &\leq 0 \\ \mu^* &\geq 0 \\ \mu_i^* \cdot (A_{\text{ineq}} \cdot w^* + b_{\text{ineq}})_i &= 0 \end{aligned}$$

for $i = 1, \dots, n_{\text{ineq}}$. Active-set methods operate on the principle that if the set of active constraints at the solution is known, the optimization problem can be reduced to solving an equality-constrained optimization problem. This simplification makes the problem significantly easier to solve compared to handling a QP with inequality constraints. With

$$\begin{aligned} \mathbb{A} &\in \{1, \dots, n_{\text{ineq}}\} \quad \text{"Active"} \\ \mathbb{I} &= \{1, \dots, n_{\text{ineq}}\} \setminus \mathbb{A} \quad \text{"Inactive"} \end{aligned}$$

w^* is a global minimizer of the QP if and only if there exist an index set \mathbb{A} and \mathbb{I} and a vector $\mu_{\mathbb{A}}^*$ so that:

$$H \cdot w^* + q + A_{\text{eq}}^\top \cdot v^* + A_{\text{ineq}, \mathbb{A}}^\top \cdot \mu_{\mathbb{A}}^* = 0 \quad (3.17)$$

$$A_{\text{eq}} \cdot w^* + b_{\text{eq}} = 0 \quad (3.18)$$

$$A_{\text{ineq},\mathbb{A}} \cdot w^* + b_{\text{ineq},\mathbb{A}} = 0 \quad (3.19)$$

$$A_{\text{ineq},\mathbb{I}} \cdot w^* + b_{\text{ineq},\mathbb{I}} \leq 0 \quad (3.20)$$

$$\mu_{\mathbb{A}}^* \geq 0 \quad (3.21)$$

and

$$\mu^* = \begin{bmatrix} \mu_{\mathbb{A}}^* \\ \mu_{\mathbb{I}}^* \end{bmatrix} \text{ with } \mu_{\mathbb{I}}^* = 0$$

Thus, an active-set method iteratively refines a guess for the active set, often referred to as the working set, and solves the equality constrained problem (3.17)–(3.19) in each iteration until (3.20) and (3.21) are also satisfied. The arising equality-constrained problems are particularly easy to solve when dealing with affine inequality constraints, as is the case in QPs. A major advantage of active-set strategies is that they can very efficiently be warm-started under circumstances where a series of related problems have to be solved, e.g. in the context of MPC or within an SQP method in the context of NMPC.

An active-set QP solver that has shown to be practical in the context of (N)MPC is qpOASES [35]. Its active set updates are based on the fact that subsequent problems in real-time MPC lie close to each other. qpOASES is best suited for small and dense problems. [118]

Interior point methods

The basic idea behind interior point methods was briefly introduced in Section 3.2.1. For NLPs, solving the KKT system involves dealing with nonlinear equations, which typically require sophisticated numerical techniques such as Newton’s method for nonlinear systems. In contrast, the KKT system of a QP is a linear system and can often be solved using more straightforward linear algebra techniques.

As discussed, the presence of inequality constraints complicates the optimization task due to the introduction of complementary slackness conditions, which are inherently non-smooth. The fundamental concept behind interior point methods is to reformulate the optimization problem to circumvent the inequality constraints.

Interior point QP solvers are often based on direct sparse linear algebra solvers. Given a special sparsity structure of the QP, which arises from applying multiple shooting to the OCP, even more efficient solutions can be achieved. An example of such a structure-exploiting interior point solver based on a Riccati recursion for linear system solutions is HPMP [39], and its successor HPIPM [37]. These solvers employ a strategy similar to the structure-exploiting strategy for linear MPC problems presented in [90], but additionally offer linear algebra kernels for small-to-midsize matrices, implemented in BLASFEO [38], optimized for various computer architectures. This enabled an additional speedup compared to existing methods. [118]

3.3 Model predictive control

In the preceding sections, we introduced optimization methods for solving OCPs in order to find an optimal state and control trajectory. To carry out a certain control task, one could simply apply the computed optimal control trajectory to the system in an open-loop manner. In an ideal scenario, this implementation would yield optimal results as defined by the OCP formulation. For any real-world application, however, one has to deal with numerous uncertainties, such as modeling errors, unforeseen external disturbances, and imperfect information about the current state of the system. Thus, we aim to observe the system and adjust our control solutions online by implementing feedback control, thereby closing the loop as the system under control progresses in a dynamically changing environment. MPC is an advanced dynamic optimization-based strategy for feedback control that can be used to control a large class of systems. The scheme is based on the solution of an OCP at each sampling instant, for which a standard formulation is provided by (3.1). This optimization problem depends on the current state of the system \hat{x}_0 , which can be either measured directly or estimated. The model in (3.1c) should accurately represent the system dynamics of interest. Depending on whether linear or nonlinear system dynamics are employed, MPC gives rise to problems that can be categorized into two broad categories: linear MPC (LMPC) and NMPC.

For the objective in (3.1a), the classic formulation for reference tracking is given as follows:

$$\begin{aligned}\ell(x(t), u(t)) &= \|x(t)\|_Q^2 + \|u(t)\|_R^2 \\ \ell_f(x(t_f)) &= \|x(t_f)\|_{Q_f}^2\end{aligned}\tag{3.22}$$

This corresponds to regulating the dynamic system to the origin. In Section 3.3.4, we will discuss other reference tracking formulations which are more practical for real-world tasks. Reference tracking formulations contrast with more general economic MPC formulations, where the stage cost may represent minimization of risk or directly maximizing profit. Another type of MPC formulations are time-optimal control formulations.

To implement the MPC scheme, it is necessary to decide on a sampling frequency with which the OCPs are solved. The sampling time T_s should be short enough to capture the fast dynamics of the process being controlled, while also being long enough to ensure that the required computations can be performed within the available time. The resulting receding horizon control strategy can be summarized as follows:

1. Measure or estimate the system state \hat{x}_0 at the current time.
2. Calculate an (approximately) optimal control trajectory by using a prediction of the future behavior of the system.
3. Apply the first optimal control input u_0^* to the system.
4. Repeat step 1-3 after the fixed sampling time T_s .

In step 2, we typically solve an OCP of type (3.1) using any of the direct methods discussed in Section 3.1. This concept is further illustrated in Fig. 3.1, which depicts the output and control trajectories at two consecutive sampling time points. The figure demonstrates how the use of the current state estimate introduces feedback into the process, thereby closing the loop. As shown in Fig. 3.1, the planned control trajectory and the corresponding predicted outputs change from one time instant to the next due to deviations of the state estimate from the model-based prediction. [88, 118]

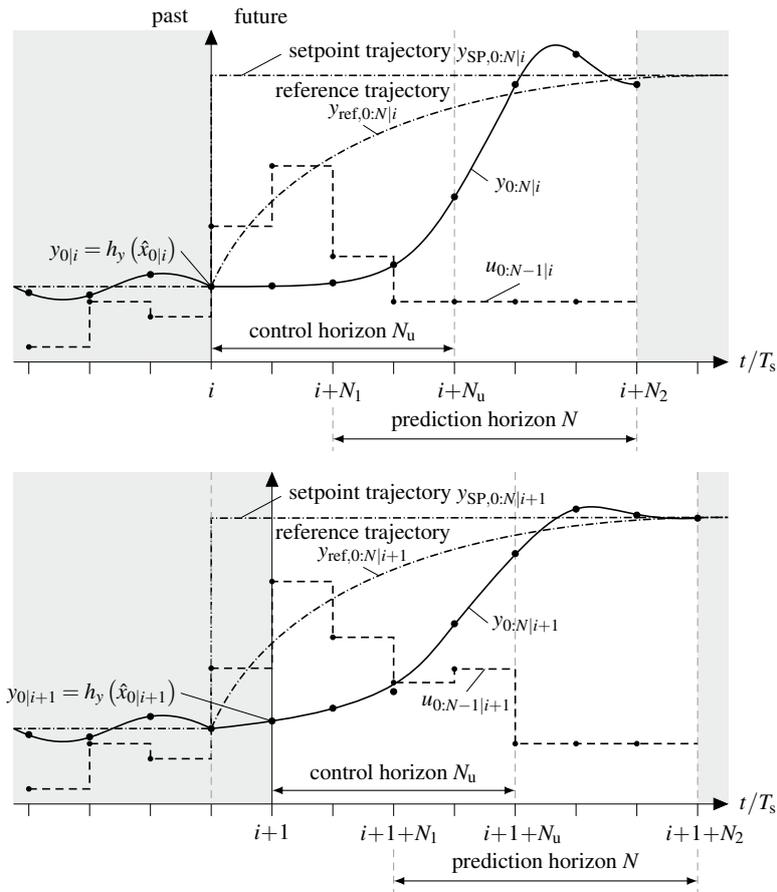


Fig. 3.1: The principle of model predictive control. In the general case, the control horizon N_u can differ from the prediction horizon N . In this work, we use identical control and prediction horizons with the lower limit $N_1 = 0$.

3.3.1 Linear model predictive control

The most basic and common LMPC formulation is characterized by a deterministic linear prediction model, affine constraints, and a quadratic cost function. The linear prediction

model can be derived by choosing a linear system identification approach or linearizing a nonlinear model at a specific operating point. Similarly, the constraints can be either affine or linearized.

LMPC gives rise to QPs. By using positive definite weighting matrices Q_k and R_k , the convexity of the QP is ensured. Convex QPs can be solved very efficiently, and if a solution is found, it is guaranteed to be globally optimal.

There are two different approaches to formulate the optimization problem: the sparse formulation and the dense formulation. The sparse formulation results from the application of multiple shooting (Section 3.1.3), while the dense formulation from the application of single shooting (Section 3.1.2).

Sparse formulation

As explained in Section 3.1.3, this approach incorporates the state equations of the linear model directly as equality constraints in the optimization problem. Consequently, the optimization variables include both the system states and the actuated values. The optimization problem can be formulated as follows:

$$\underset{x_{0:N}, u_{0:N-1}}{\text{minimize}} \quad \frac{1}{2} \cdot \sum_{k=0}^{N-1} \begin{bmatrix} x_k \\ u_k \end{bmatrix}^\top \cdot H_k \cdot \begin{bmatrix} x_k \\ u_k \end{bmatrix} + 2 \cdot \zeta_k^\top \cdot \begin{bmatrix} x_k \\ u_k \end{bmatrix} + \frac{1}{2} \cdot x_N^\top \cdot Q_N \cdot x_N + \zeta_{x,N}^\top \cdot x_N \quad (3.23a)$$

$$\text{subject to} \quad x_0 = \hat{x}_0, \quad (3.23b)$$

$$x_{k+1} = A_k \cdot x_k + B_k \cdot u_k, \quad k = 0, \dots, N-1, \quad (3.23c)$$

$$C_{x,k} \cdot x_k + C_{u,k} \cdot u_k \leq b_k, \quad k = 0, \dots, N-1, \quad (3.23d)$$

$$C_{x,N} \cdot x_N \leq b_N, \quad (3.23e)$$

with

$$H_k = \begin{bmatrix} Q_k & S_k^\top \\ S_k & R_k \end{bmatrix}, \quad \zeta_k = \begin{bmatrix} \zeta_{x,k} \\ \zeta_{u,k} \end{bmatrix}$$

Here, $x_k \in \mathbb{R}^{n_x}$ represents the state vectors, $u_k \in \mathbb{R}^{n_u}$ denotes the controls, and the cost matrices and vectors are $Q_k \in \mathbb{R}^{n_x \times n_x}$, $S_k \in \mathbb{R}^{n_u \times n_x}$, $R_k \in \mathbb{R}^{n_u \times n_u}$, $\zeta_{x,k} \in \mathbb{R}^{n_x}$, $\zeta_{u,k} \in \mathbb{R}^{n_u}$. The cost matrices Q_k and R_k are to be chosen by the control designer, often as positive-definite diagonal matrices to ensure that QP (3.23) is convex. The constraints include dynamic constraints with matrices $A_k \in \mathbb{R}^{n_x \times n_x}$ and $B_k \in \mathbb{R}^{n_x \times n_u}$, inequality constraints with $C_{x,k} \in \mathbb{R}^{n_{\text{ineq}} \times n_x}$, $C_{u,k} \in \mathbb{R}^{n_{\text{ineq}} \times n_u}$ and $C_{x,N} \in \mathbb{R}^{n_{\text{ineq},f} \times n_x}$, and an initial constraint with $\hat{x}_0 \in \mathbb{R}^{n_x}$.

When merging all optimization variables in the vector w

$$w = [x_0^\top \quad u_0^\top \quad \dots \quad x_{N-1}^\top \quad u_{N-1}^\top \quad x_N^\top]^\top$$

and formulating the cost function and constraints in (3.23) in terms of w using the matrices

$$\begin{aligned}
 H_S &= \begin{bmatrix} H_0 & & & 0 \\ & H_1 & & \\ & & \ddots & \\ & & & H_{N-1} \\ 0 & & & & Q_N \end{bmatrix} & q_S &= \begin{bmatrix} \zeta_0 \\ \zeta_1 \\ \vdots \\ \zeta_{N-1} \\ \zeta_{x,N} \end{bmatrix} \\
 A_{\text{eq},S} &= \begin{bmatrix} I & & & & & & 0 \\ A_0 & B_0 & -I & & & & \\ & & A_1 & B_1 & -I & & \\ & & & \ddots & \ddots & \ddots & \\ 0 & & & & A_{N-1} & B_{N-1} & -I \end{bmatrix} & b_{\text{eq},S} &= \begin{bmatrix} \hat{x}_0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \\
 A_{\text{ineq},S} &= \begin{bmatrix} C_{x,0} & C_{u,0} & & & & & 0 \\ & & C_{x,1} & C_{u,1} & & & \\ & & & \ddots & \ddots & & \\ & & & & C_{x,N-1} & C_{u,N-1} & \\ 0 & & & & & & C_{x,N} \end{bmatrix} & b_{\text{ineq},S} &= \begin{bmatrix} b_0 \\ b_1 \\ \vdots \\ b_{N-1} \\ b_N \end{bmatrix}
 \end{aligned}$$

the stage-wise formulation of the LMPC in (3.23) is transformed into the standard QP formulation (3.16):

$$\underset{w}{\text{minimize}} \quad \frac{1}{2} \cdot w^\top \cdot H_S \cdot w + w^\top \cdot q_S + \text{const.} \quad (3.24a)$$

$$\text{subject to} \quad A_{\text{eq},S} \cdot w = b_{\text{eq},S}, \quad (3.24b)$$

$$A_{\text{ineq},S} \cdot w \leq b_{\text{ineq},S} \quad (3.24c)$$

In this formulation, all matrices exhibit a banded sparse structure. The dimensions of the matrices and vectors are as follows [89]:

$$\begin{aligned}
 w &\in \mathbb{R}^{n_x \cdot (N+1) + n_u \cdot N} \\
 H_S &\in \mathbb{R}^{n_x \cdot (N+1) + n_u \cdot N \times n_x \cdot (N+1) + n_u \cdot N} \\
 q_S &\in \mathbb{R}^{n_x \cdot (N+1) + n_u \cdot N} \\
 A_{\text{eq},S} &\in \mathbb{R}^{n_x \cdot (N+1) \times n_x \cdot (N+1) + n_u \cdot N} \\
 b_{\text{eq},S} &\in \mathbb{R}^{n_x \cdot (N+1)} \\
 A_{\text{ineq},S} &\in \mathbb{R}^{n_{\text{ineq}} \cdot N + n_{\text{ineq},f} \times n_x \cdot (N+1) + n_u \cdot N} \\
 b_{\text{ineq},S} &\in \mathbb{R}^{n_{\text{ineq}} \cdot N + n_{\text{ineq},f}}
 \end{aligned}$$

Dense formulation

In the dense formulation, equivalent to the single shooting approach (Section 3.1.2), the model equations for the system dynamics in the QP (3.23) are utilized to eliminate the state variables x_k for $k = 1, \dots, N$ as follows:

$$\begin{aligned}
x_1 &= A \cdot x_0 + B \cdot u_0 \\
x_2 &= A \cdot x_1 + B \cdot u_1 = A \cdot (A \cdot x_0 + B \cdot u_0) + B \cdot u_1 \\
&= A^2 \cdot x_0 + A \cdot B \cdot u_0 + B \cdot u_1 \\
x_3 &= A \cdot x_2 + B \cdot u_2 = A \cdot (A^2 \cdot x_0 + A \cdot B \cdot u_0 + B \cdot u_1) + B \cdot u_2 \\
&= A^3 \cdot x_0 + A^2 \cdot B \cdot u_0 + A \cdot B \cdot u_1 + B \cdot u_2 \\
&\vdots \\
\Rightarrow x_k &= A^k \cdot x_0 + A^{k-1} \cdot B \cdot u_0 + A^{k-2} \cdot B \cdot u_1 + \dots + A \cdot B \cdot u_{k-2} + B \cdot u_{k-1} \\
&= A^k \cdot x_0 + \sum_{i=0}^{k-1} A^i \cdot B \cdot u_{k-i-1}
\end{aligned} \tag{3.25}$$

The resulting QP has $n_u \cdot N$ optimization variables, collected in $u_{0:N-1}$, and is formulated as:

$$\underset{u_{0:N-1}}{\text{minimize}} \quad \frac{1}{2} \cdot u_{0:N-1}^\top \cdot H_D \cdot u_{0:N-1} + q_D^\top \cdot u_{0:N-1} \tag{3.26a}$$

$$\text{subject to} \quad C_D \cdot u_{0:N-1} + c_D \leq 0 \tag{3.26b}$$

Here, the Hessian matrix and gradient vector are defined as $H_D \in \mathbb{R}^{n_u \cdot N \times n_u \cdot N}$ and $q_D \in \mathbb{R}^{n_u \cdot N}$, respectively. The constraints are given by $C_D \in \mathbb{R}^{n_{\text{ineq}} \cdot N + n_{\text{ineq},f} \times n_u \cdot N}$ and $c_D \in \mathbb{R}^{n_{\text{ineq}} \cdot N + n_{\text{ineq},f}}$. Due to the dependency of the state variables on not only the current system input but all past system inputs, as shown in (3.25), the matrices in the optimization problem (3.26) are dense, meaning that most elements of the matrices are non-zero.

The dense formulation involves fewer optimization variables and constraints, resulting in much smaller matrices. However, it can still lead to higher computational complexity and memory requirements because dense linear algebra techniques need to be employed. Decomposing the dense matrices is numerically expensive, potentially making the sparse formulation faster to solve. The sparse banded structure of the Hessian, the equality and inequality constraint matrices result in an overall sparse structure of the resulting QP. This can be exploited by tailored numerical methods in QP solver, such as efficient lower-upper decomposition. Particularly in large-scale problems, the sparsity structure can lead to faster convergence times and reduced memory usage. Additionally, the sparse formulation provides better conditioning for unstable systems by avoiding matrix potentiation. For LTV MPC, the constraints and system matrices change at each time step. With a sparse structure, less time is needed to construct the necessary matrices.

In general, the dense formulation is advantageous for systems with fewer control inputs, shorter prediction horizons, or simpler dynamics. The sparse formulation is better suited for large-scale problems with long prediction horizons and complex dynamics. [89]

Linear time-varying model predictive control

The preceding elaborations assume a linear prediction model. However, real-world systems often exhibit nonlinear dynamics. Depending on the system and use case, a linear representation may suffice for the control purposes of interest. The linear system representation can be either obtained from a simplified modeling approach, linear system identification methods, or linearizing the nonlinear dynamics, including the constraints if necessary. Once a linear representation is obtained, it can be used to formulate the QP by computing the required dense or sparse matrices. When the linear prediction model is derived from linearization, the linearization step occurs once for the so-called LTI case and the resulting model is kept constant for the whole control task. As a result, the QP matrices need to be computed only once. Another approach is LTV MPC, where the system matrices and thus the QP matrices are only valid for the current sampling instant. In this case, the system dynamics are either described by an LTV model, or, as considered in this work, LTV MPC is used to approximately take into account nonlinear system behavior. This involves linearizing a nonlinear system model $\dot{x} = f(x, u)$ through Taylor series expansion around each current operating point (OP) given by x_{OP}, u_{OP} . The linearized continuous-time state-space model at the current OP is:

$$\dot{x} = f(x, u) \approx f(x_{OP}, u_{OP}) + \underbrace{\frac{\partial f}{\partial x} \Big|_{x_{OP}, u_{OP}}}_{=:A} \cdot (x - x_{OP}) + \underbrace{\frac{\partial f}{\partial u} \Big|_{x_{OP}, u_{OP}}}_{=:B} \cdot (u - u_{OP}) \quad (3.27)$$

For a stationary OP, $f(x_{OP}, u_{OP}) = 0$. Otherwise, the non-stationary term must be considered in the state prediction. In the LTV case, a new linearized LTI state-space model results from linearization in every time step, which is used for prediction within MPC. Consequently, the QP matrices have to be recalculated at each time step. The approach involves the following steps:

1. Obtain a new measurement or estimation for x_k .
2. Build one linear model around the measured or estimated value x_k .
3. Use this model for an LTI prediction over the entire prediction horizon.
4. Repeat this procedure in every sampling step.

Since the prediction remains linear in the optimization variables, the LTV MPC results in one QP solved at each time step, same as for the LTI case. Compared to LTI MPC, LTV MPC increases computational load due to the additional linearization step at each time step. [89]

The linearization can occur before or after discretization. In the first-linearize-then-discretize approach, the continuous-time nonlinear state-space is first linearized and then discretized using the matrix exponential, as given in (3.5). In the first-discretize-then-linearize approach, the continuous-time nonlinear state-space is discretized using a numerical integration scheme before linearizing the resulting discrete-time nonlinear dynamics to obtain the sensitivities. Given that the numerical integration provides an accurate representation of the nonlinear dynamics for the sampling time of interest, the second approach may offer higher accuracy, albeit with increased computational complexity. [47]

Another aspect to consider is that the linearization in (3.27) can be expressed in absolute or relative variables:

$$\dot{x} \approx A \cdot x + B \cdot u + f(x_{\text{OP}}, u_{\text{OP}}) - A \cdot x_{\text{OP}} - B \cdot u_{\text{OP}} \quad (3.28)$$

$$\dot{x}_{\text{rel}} \approx A \cdot x_{\text{rel}} + B \cdot u_{\text{rel}} \quad (3.29)$$

where $x_{\text{rel}} = x - x_{\text{OP}}$ and $u_{\text{rel}} = u - u_{\text{OP}}$. Depending on the choice, the resulting QP matrices differ. With relative variables (3.29), the QP solution has to be adjusted accordingly before being applied to the controlled system. The standard formulations given in Section 3.3.1 result from using relative variables.

In this work, we linearize after discretization and compute the prediction matrices for absolute variables (3.28).

3.3.2 Nonlinear model predictive control

Many physical systems cannot be accurately represented by linear or linearized dynamics. This motivated a generalization to allow for nonlinear dynamic equations. NMPC has the ability to directly handle nonlinear dynamics, constraints, and non-convex objectives, allowing design requirements to be naturally translated into mathematical statements. This results in an optimization problem that is an NLP instead of a QP.

In this thesis, we consider NMPC problems in a multiple shooting context, giving rise to problems like the NLP in (3.4). In Section 3.2, we provided an overview of major Newton-type optimization algorithms for solving general NLPs. The algorithm relevant for this thesis is the SQP method.

When applying the SQP method to the multiple shooting type NLP in (3.4), each iteration solves the following structured QP subproblem:

$$\begin{aligned} \text{minimize} \quad & \sum_{k=0}^{N-1} \frac{1}{2} \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix}^\top \cdot H_k \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} + q_k^\top \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} + \frac{1}{2} \cdot \Delta x_N^\top \cdot H_N \cdot \Delta x_N + q_N^\top \cdot \Delta x_N \\ \Delta x_{0:N}, \Delta u_{0:N-1} \end{aligned} \quad (3.30a)$$

$$\text{subject to} \quad \Delta x_0 - \Delta \hat{x}_0 = 0, \quad (3.30b)$$

$$f_{d,k} + \frac{\partial f_d(\bar{x}_k, \bar{u}_k)}{\partial (\bar{x}_k, \bar{u}_k)} \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} - \Delta x_{k+1} = 0, \quad k = 0, \dots, N-1, \quad (3.30c)$$

$$h_k + \frac{\partial h(\bar{x}_k, \bar{u}_k)}{\partial (\bar{x}_k, \bar{u}_k)} \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} \leq 0, \quad k = 0, \dots, N-1, \quad (3.30d)$$

$$h_f + \frac{\partial h_f(\bar{x}_N)}{\partial \bar{x}_N} \cdot \Delta x_N \leq 0, \quad (3.30e)$$

where $\bar{x}_{0:N} = [\bar{x}_0^\top, \dots, \bar{x}_N^\top]^\top$ and $\bar{u}_{0:N-1} = [\bar{u}_0^\top, \dots, \bar{u}_{N-1}^\top]^\top$ denote the current state and control trajectory which form the linearization point for the problem functions. The vectors are defined as $q_k := \nabla \ell(\bar{x}_k, \bar{u}_k)$ and $q_N := \nabla \ell_f(\bar{x}_N)$, $\Delta \hat{x}_0 = \hat{x}_0 - \bar{x}_0$, $f_{d,k} := f_d(\bar{x}_k, \bar{u}_k) - \bar{x}_{k+1}$, $h_k := h(\bar{x}_k, \bar{u}_k)$ and $h_f := h_f(\bar{x}_N)$. The terminal cost matrix $H_N := \nabla_x^2 \ell_f(\bar{x}_N)$ is defined, while the stage Hessian blocks H_k for $k = 0, \dots, N-1$ depend on the approximation technique used.

The Lagrangian of the NLP in (3.4) is given by:

$$\begin{aligned} \mathcal{L}(x_{0:N}, u_{0:N-1}, v_{-1:N-1}, \mu_{0:N}) &= \sum_{k=0}^{N-1} \ell(x_k, u_k) + \ell_f(x_N) + v_{-1}^\top \cdot (x_0 - \hat{x}_0) + \mu_N^\top \cdot h_f(x_N) \\ &\quad + \sum_{k=0}^{N-1} v_k^\top \cdot (f_d(x_k, u_k) - x_{k+1}) + \sum_{k=0}^{N-1} \mu_k^\top \cdot h(x_k, u_k), \end{aligned}$$

where v_k for $k = 0, \dots, N-1$ are the multipliers corresponding to the continuity constraints in (3.4c), v_{-1} denotes the multiplier of the initial value condition in (3.4b), and μ_k for $k = 0, \dots, N$ are the multipliers corresponding to the inequality constraints in (3.4d) and (3.4e).

For an exact Hessian SQP method, the block matrices correspond to the following second order derivatives $H_k := \nabla_{(x_k, u_k)}^2 \mathcal{L}(\bar{x}_{0:N}, \bar{u}_{0:N-1}, \bar{v}_{-1:N-1}, \bar{\mu}_{0:N})$. In case of the common (non-linear) least-squares type objective $\ell(x_k, u_k) = \frac{1}{2} \cdot \|r(x_k, u_k)\|_2^2$, the Gauss-Newton Hessian approximation is given by $H_{GN,k} := \nabla r(\bar{x}_k, \bar{u}_k) \cdot \nabla r(\bar{x}_k, \bar{u}_k)^\top$.

The QP subproblem in (3.30) has a particular sparsity structure, see Section 3.3.1. There are two common approaches for solving the band-structured QP subproblems (3.30) arising in optimal control. The first approach exploits the sparsity structure by employing tailored direct linear algebra routines within the convex QP solver. The second approach, known as condensing, eliminates all state deviations via the continuity constraints in (3.30c) and solves instead a dense QP of significantly smaller dimension. Condensing is not applied within this thesis. For further details, please refer to the following references: [6, 14].

During the runtime of the process, at each time step of the control algorithm, a new instance of the NLP has to be solved. From one time to the next, the structure stays the same, but certain parameters, such as the initial state, change and the reference might change. Especially with

the SQP method, the solution from the previous time step can be leveraged to provide a good initial guess. Specifically, the previous solution is reused, shifted by one time step to account for the progression in time. Both the preceding solution for the optimization variables and for the Lagrange multipliers can be used for initialization. Typically, the last state and actuated value in the prediction horizon are set as $x_N := x_{N-1}$ and $u_{N-1} := u_{N-2}$, respectively. Similarly, the last Lagrange multipliers can be initialized.

While local minimizers can be determined using the SQP method, they do not necessarily correspond to global minimizers. In the case of non-convex optimization in the context of an online algorithm for NMPC, local minimizers are generally acceptable. Ideally, a local minimizer is obtained within each sampling step. However, with very short sampling times, a fully converged solution may not be obtainable. Instead, the so-called real-time iteration (RTI) scheme [32] can be employed. In this approach, a single SQP step is performed at each time step, and the resulting suboptimal solution is applied to the plant. Thus, convergence to the optimal solution is not realized in every time step, but over multiple time steps during runtime. [88, 118]

3.3.3 State estimation

If the current state \hat{x}_0 for the initial state constraint in the MPC optimization problem is not known, we first need to *estimate* it based on the available measurement information. Kalman filters (KFs) are often used to estimate the internal states of a system in the presence of uncertain and indirect measurements. Named after Rudolf E. Kálmán [60], the KF is an optimal estimation algorithm, that addresses the general problem of trying to estimate the state x of a discrete-time controlled process that is governed by a linear stochastic system. A derivation of the discrete-time KF is provided in [103]. Here, we discuss a nonlinear extension of the KF, the EKF, a widely used nonlinear state estimation technique originally proposed by Stanley Schmidt [9]. The idea is to first linearize the nonlinear system, and then applying the linear KF equations to the linearized system.

Extended Kalman Filter

Suppose the process with state vector $x \in \mathbb{R}^{n_x}$ is governed by a nonlinear stochastic difference equation given by:

$$x_k = f_d(x_{k-1}, u_{k-1}, w_{p,k-1}) \quad (3.31)$$

with a measurement $y_m \in \mathbb{R}^{n_{y_m}}$ described by:

$$y_{m,k} = \phi_m(x_k, v_{m,k}) \quad (3.32)$$

3 Control preliminaries

The random variables w_p and v_m represent the process and measurement noise, respectively. They are assumed to be white, uncorrelated and normally distributed:

$$\begin{aligned} p(w_p) &\sim N(0, Q_{e,k}) \\ p(v_m) &\sim N(0, R_{e,k}) \end{aligned}$$

where $Q_{e,k}$ and $R_{e,k}$ are the known covariance matrices for the process and measurement noise, respectively. Our goal is to estimate the state x_k based on our knowledge of the system dynamics and the available noisy measurements $y_{m,k}$. We denote the a priori estimate as \hat{x}_k^- and the a posteriori estimate as \hat{x}_k . Both are estimates of x_k , but \hat{x}_k^- is the estimate of x_k before incorporating the measurement $y_{m,k}$, while \hat{x}_k is the estimate of x_k after incorporating $y_{m,k}$. Similarly, P_k^- denotes the covariance of the estimation error of \hat{x}_k^- , and P_k denotes the covariance of the estimation error of \hat{x}_k .

When performing a Taylor series expansion of the state equation (3.31) around $x_{k-1} = \hat{x}_{k-1}$ and $w_{p,k-1} = 0$, we obtain:

$$\begin{aligned} x_k &= f_d(\hat{x}_{k-1}, u_{k-1}, 0) + \left. \frac{\partial f_d}{\partial x} \right|_{(\hat{x}_{k-1}, u_{k-1}, 0)} \cdot (x_{k-1} - \hat{x}_{k-1}) + \left. \frac{\partial f_d}{\partial w_p} \right|_{(\hat{x}_{k-1}, u_{k-1}, 0)} \cdot w_{p,k-1} \\ &= f_d(\hat{x}_{k-1}, u_{k-1}, 0) + A_{k-1} \cdot (x_{k-1} - \hat{x}_{k-1}) + L_{k-1} \cdot w_{p,k-1} \end{aligned} \quad (3.33)$$

Similarly, when we linearize the measurement equation (3.32) around $x_k = \hat{x}_k^-$ and $v_{m,k} = 0$, we obtain:

$$\begin{aligned} y_{m,k} &= \phi_m(\hat{x}_k^-, 0) + \left. \frac{\partial \phi_m}{\partial x} \right|_{(\hat{x}_k^-, 0)} \cdot (x_k - \hat{x}_k^-) + \left. \frac{\partial \phi_m}{\partial v_m} \right|_{(\hat{x}_k^-, 0)} \cdot v_{m,k} \\ &= \phi_m(\hat{x}_k^-, 0) + C_{m,k} \cdot (x_k - \hat{x}_k^-) + V_k \cdot v_{m,k} \end{aligned} \quad (3.34)$$

This results in a linear state-space system in (3.33) that relates the state at the previous time step $k-1$ to the state at the current time step k , and a linear measurement in (3.34) that relates the state x_k to the measurement $y_{m,k}$. This allows us to use the standard KF equations to estimate the state, resulting in the following algorithm for the discrete-time EKF:

1. Initialize the filter with \hat{x}_0 and P_0 .
2. For $k = 1, 2, \dots$, perform the following steps:
 - (a) Compute the partial derivative matrices:

$$\begin{aligned} A_{k-1} &= \left. \frac{\partial f_d}{\partial x} \right|_{(\hat{x}_{k-1}, u_{k-1}, 0)} \\ L_{k-1} &= \left. \frac{\partial f_d}{\partial w_p} \right|_{(\hat{x}_{k-1}, u_{k-1}, 0)} \end{aligned}$$

- (b) Perform the time update of the state estimate and estimation-error covariance (prediction step):

$$P_k^- = A_{k-1} \cdot P_{k-1} \cdot A_{k-1}^\top + L_{k-1} \cdot Q_{e,k-1} \cdot L_{k-1}^\top$$

$$\hat{x}_k^- = f_d(\hat{x}_{k-1}, u_{k-1}, 0)$$

- (c) Compute the partial derivative matrices:

$$C_{m,k} = \left. \frac{\partial \phi_{m,k}}{\partial x} \right|_{(\hat{x}_k^-, 0)}$$

$$V_k = \left. \frac{\partial \phi_{m,k}}{\partial v_m} \right|_{(\hat{x}_k^-, 0)}$$

- (d) Perform the measurement update of the state estimate and estimation-error covariance (update step):

$$K_k = P_k^- \cdot C_{m,k}^\top \cdot \left(C_{m,k} \cdot P_k^- \cdot C_{m,k}^\top + V_k \cdot R_{e,k} \cdot V_k^\top \right)^{-1}$$

$$\hat{x}_k = \hat{x}_k^- + K_k \cdot [y_{m,k} - \phi_{m,k}(\hat{x}_k^-, 0)]$$

$$P_k = (I - K_k \cdot C_{m,k}) \cdot P_k^- \cdot (I - K_k \cdot C_{m,k})^\top + K_k \cdot R_{e,k} \cdot K_k^\top$$

$$= (I - K_k \cdot C_{m,k}) \cdot P_k^-$$

The first expression for P_k above is known as the Joseph stabilized version of the covariance measurement update equation. Formulated by Peter Joseph in the 1960s, it is shown to be more stable and robust than the simpler expression for P_k . The Joseph form guarantees that P_k remains symmetric positive definite, provided that P_k^- is symmetric positive definite. In contrast, while the third expression for P_k is computationally simpler, it does not ensure symmetry or positive definiteness. [103]

A key limitation of the EKF is that the probability distributions of random variables cease to be normal after nonlinear transformations. Consequently, the EKF is a pragmatic estimator that approximates the optimal solution of Bayes' rule by linearizing the nonlinear dynamics and measurement equations. This approximation may not always be optimal but provides a practical method for state estimation in nonlinear systems. [125]

3.3.4 Offset-free reference tracking

This section delves into the topic of reference tracking. The concepts are presented for linear systems. However, they can be extended to nonlinear systems by using the same optimization problems subject to a nonlinear instead of a linear model, and incorporating a nonlinear observer, such as an EKF. Throughout the section, it is assumed that the reference values are piecewise constant over time.

MPC problem formulation

In Section 3.3, we introduced the (continuous-time) cost function in (3.22). The corresponding standard MPC problem, which aims to regulate the system state to the origin while penalizing the input, is formulated as follows:

$$\begin{aligned}
 & \underset{x_{0:N-1}, u_{0:N-1}}{\text{minimize}} && \frac{1}{2} \cdot \sum_{k=0}^{N-1} \|x_k\|_Q^2 + \|u_k\|_R^2 \\
 & \text{subject to} && x_0 = \hat{x}_0, \\
 & && x_{k+1} = A \cdot x_k + B \cdot u_k, \quad k = 0, \dots, N-1, \\
 & && C_x \cdot x_k + C_u \cdot u_k \leq b, \quad k = 0, \dots, N-1
 \end{aligned} \tag{3.35}$$

Note that we excluded the terminal cost and terminal constraint. While this omission can increase the set of feasible solutions, it may compromise system stability. Therefore, the prediction horizon N must be carefully chosen to ensure stable feedback control.

Instead of regulating to the origin, a common practical task is to track a given non-zero output reference y_{ref} . An intuitive adaption of the cost function to $\sum_{k=0}^{N-1} \|y_k - y_{\text{ref}}\|_Q^2 + \|u_k\|_R^2$ will not achieve the desired outcome. This is because a trade-off arises between tracking the reference values and penalizing the actuated values, except in the special case where the reference can be achieved with zero steady-state input, i.e. when $y_\infty = y_{\text{ref}}$ for $u_\infty = 0$, a condition that is generally not satisfied. For reference values that require $u_\infty \neq 0$, a steady-state offset will exist, even without considering constraints. The two cost terms cannot simultaneously be minimized below a certain threshold, making the steady-state solution dependent on the weighting matrices Q and R .

One approach to achieve appropriate reference tracking is to adopt the so-called delta formulation. Here, the optimization problem takes into account the change of the actuated values Δu . For instance, the following optimization problem can be used:

$$\begin{aligned}
 & \underset{\substack{x_{0:N-1}, u_{0:N-1}, \\ \Delta u_{0:N-1}}}{\text{minimize}} && \frac{1}{2} \cdot \sum_{k=0}^{N-1} \|y_k - y_{\text{ref}}\|_Q^2 + \|\Delta u_k\|_R^2 \\
 & \text{subject to} && x_0 = \hat{x}_0, \\
 & && u_0 = u_{-1} + \Delta u_0, \\
 & && x_{k+1} = A \cdot x_k + B \cdot u_k, \quad k = 0, \dots, N-1, \\
 & && u_k = u_{k-1} + \Delta u_k, \quad k = 0, \dots, N-1, \\
 & && y_k = C \cdot x_k, \quad k = 0, \dots, N-1, \\
 & && C_x \cdot x_k + C_u \cdot u_k \leq b, \quad k = 0, \dots, N-1,
 \end{aligned} \tag{3.36}$$

where u_{-1} denotes the actuated value from the previous time step. For a stable closed-loop control system, the system can converge to a stationary point $(u, x, y) \rightarrow (u_\infty, x_\infty, y_\infty)$ with $\Delta u_\infty =$

0. The stationary point is characterized as the minimizer of the optimization problem, ensuring $y_\infty = y_{\text{ref}}$ if the reference is reachable. Consequently, the delta formulation can naturally achieve offset-free reference tracking. The steady state (u_∞, x_∞) of the system must satisfy the following condition to ensure this behavior:

$$\begin{aligned} x_\infty &= A \cdot x_\infty + B \cdot u_\infty \\ C \cdot x_\infty &= y_{\text{ref}} \end{aligned} \quad (3.37)$$

This can be written as follows:

$$\begin{bmatrix} I - A & -B \\ C & 0 \end{bmatrix} \cdot \begin{bmatrix} x_\infty \\ u_\infty \end{bmatrix} = \begin{bmatrix} 0 \\ y_{\text{ref}} \end{bmatrix} \quad (3.38)$$

For square systems, which have as many outputs as actuated values, with a full-row rank matrix, there exists one unique solution. In this case, the delta formulation will regulate the system to this specific point, provided the closed-loop system is stable and the point is reachable. However, for non-square systems, two different initial conditions but same reference values might yield different solutions for x_∞ and u_∞ . Additionally, if the optimization problem is formulated as given in (3.36), where both u_k and Δu_k are treated as optimization variables in addition to x_k , the resulting problem is larger and may increase the computational cost, depending on the chosen solver and the exploitation of sparsity. In this case, the number of optimization variables increases from $N \cdot (n_x + n_u)$ to $N \cdot (n_x + 2n_u)$.

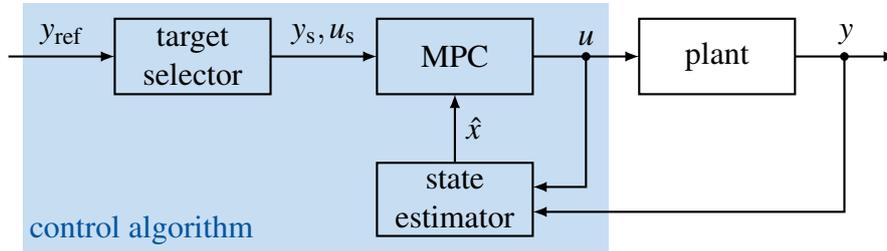


Fig. 3.2: Two-layered control structure: MPC regulator, state estimator, and target selector.

An alternative approach is a two-layered control structure as depicted in Fig. 3.2, where the desired steady-state value for the system states and actuated values are computed explicitly by solving an optimization problem using a target selector:

$$\underset{x_s, u_s}{\text{minimize}} \quad \frac{1}{2} \cdot (\|C \cdot x_s - y_{\text{ref}}\|_{Q_s}^2 + \|u_s\|_{R_s}^2) \quad (3.39a)$$

$$\text{subject to} \quad \begin{bmatrix} I - A & -B \\ C & 0 \end{bmatrix} \cdot \begin{bmatrix} x_s \\ u_s \end{bmatrix} = \begin{bmatrix} 0 \\ y_{\text{ref}} \end{bmatrix}, \quad (3.39b)$$

$$C_x \cdot x_s + C_u \cdot u_s \leq b \quad (3.39c)$$

The target selector considers only the steady-state dynamics. The objective is to find the optimal steady-state values for the actuated values and system states. The definition of the optimal target

values depends on the specific task. The cost function depicted in (3.39a) minimizes the steady-state tracking error and control effort. The definition of the cost function enables calculation of a unique steady state, even for an overactuated system. The steady-state values should lead to the reference values, satisfying (3.39b), while considering system constraints (3.39c). Feasibility issues may arise if the reference y_{ref} cannot be achieved in steady state due to the imposed constraints, resulting in an infeasible optimization problem. This can generally be managed by relaxing the constraints through the addition of slack variables.

The solution of the target selector serves as input for the dynamic regulator, where the entire system dynamics, including the transient behavior, are considered. The dynamic regulator calculates the actuated values applied to the plant. Compared to the delta formulation, this approach requires solving two optimization problems, with the steady-state target values u_s and x_s changing in dependence on the reference values.

Tracking the desired output references in the dynamic regulator is achieved by replacing the states and actions in the original problem (3.35) with the deviations from these setpoints: $\Delta x_k = x_k - x_s$ and $\Delta u_k = u_k - u_s$. This coordinate transformation allows for reusing the standard MPC formulation. A problem with the standard formulation combined with a target selector arises for nonminimum phase systems. Once the system, expressed in deviations around the desired setpoints, reaches the origin, the stationary solution corresponds to the desired reference. However, the optimal transient behavior of the states may result in undesirable transient behavior for the outputs in nonminimum phase systems. Additionally, for systems with direct feedthrough, there is no cost formulation in terms of states and actions that is equivalent to directly penalizing the objective $\|y - y_{\text{ref}}\|_Q^2 = \|C \cdot x + D \cdot u - y_{\text{ref}}\|_Q^2$ due to cross-coupling terms between states and actions. Thus, we introduce an alternative formulation for the regulation problem:

$$\begin{aligned}
 & \underset{x_{0:N-1}, u_{0:N-1}}{\text{minimize}} && \frac{1}{2} \cdot \sum_{k=0}^{N-1} \|y_k - y_{\text{ref}}\|_Q^2 + \|u_k - u_s\|_R^2 \\
 & \text{subject to} && x_0 = \hat{x}_0, \\
 & && x_{k+1} = A \cdot x_k + B \cdot u_k, \quad k = 0, \dots, N-1, \\
 & && y_k = C \cdot x_k + D \cdot u_k, \quad k = 0, \dots, N-1, \\
 & && C_x \cdot x_k + C_u \cdot u_k \leq b, \quad k = 0, \dots, N-1
 \end{aligned} \tag{3.40}$$

All in all, this procedure allows for offset-free reference tracking if the reference is reachable. Additionally, even for an overactuated system, the system is always steered to the same steady-state values, independent of its initial condition. [76, 89]

Disturbance rejection

Given that the reference point is reachable despite the existence of constraints, another common objective in closed-loop control is to design a feedback controller that compensates for

asymptotically constant, nonzero disturbances due to plant-model mismatch, ensuring the tracking error vanishes in steady state. In classical unconstrained control, an integrating mode is introduced to eliminate the resulting tracking offset. In both constrained and unconstrained systems, we aim to prevent the system from settling at an unconstrained steady state that exhibits an offset in the controlled variables.

We distinguish between the number n_{y_m} of measured outputs, the number n_y of outputs we desire to track, and the number n_d of disturbances. The MPC algorithm only achieves offset-free control when the tracked outputs match their predictions from the controller's internal model. In practice, various unmodeled and unmeasured disturbances affect the system inputs, states, and measured outputs, causing deviations between the prediction and the actual tracked outputs. To compensate for disturbances, one approach is to (i) model the disturbance, (ii) use output measurements and the model to estimate the state and the disturbance, and (iii) determine control inputs that use the disturbance estimate to eliminate the tracking offset.

Consider the discrete-time linear time-invariant system:

$$\begin{aligned} x_{k+1} &= A \cdot x_k + B \cdot u_k \\ y_{m,k} &= C_m \cdot x_k \\ y_k &= T_y \cdot y_{m,k} \end{aligned} \quad (3.41)$$

where $x_k \in \mathbb{R}^{n_x}$, $u_k \in \mathbb{R}^{n_u}$ and $y_{m,k} \in \mathbb{R}^{n_{y_m}}$ are the state, input, measured output vector, respectively. The controlled variables $y \in \mathbb{R}^{n_y}$ are a linear combination of the measured variables for which offset-free behavior is sought. Without loss of generality, we assume T_y has full row rank. To achieve offset-free performance we augment the system state with an integrating disturbance d to capture the plant-model mismatch in steady state:

$$d_{k+1} = d_k \quad (3.42)$$

The original system (3.41) is augmented with the constant, nonzero disturbance model (3.42). The augmented system model used by the state estimator to estimate both states and disturbances is given by:

$$\begin{aligned} \begin{bmatrix} x_{k+1} \\ d_{k+1} \end{bmatrix} &= \begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} \cdot \begin{bmatrix} x_k \\ d_k \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \cdot u_k \\ y_{m,k} &= \begin{bmatrix} C_m & C_d \end{bmatrix} \cdot \begin{bmatrix} x_k \\ d_k \end{bmatrix} \end{aligned} \quad (3.43)$$

We choose B_d and C_d to determine how the integrating disturbance affects the states and measured outputs, with the only restriction being that the augmented system must be observable. From the Hautus observability condition, the augmented system (3.43) is observable if and only if the matrix in (3.44) has full column rank for all eigenvalues:

$$\text{rank} \begin{bmatrix} A - \lambda \cdot I & B_d \\ 0 & I - \lambda \cdot I \\ C_m & C_d \end{bmatrix} = n_x + n_d \quad \forall \lambda \in \mathbb{C} \quad (3.44)$$

The first set of columns is linearly independent with rank n_x if and only if the unaugmented system (A, C_m) is observable. The second set of columns is linearly independent from the first except possibly for $\lambda = 1$. Checking the Hautus condition for $\lambda = 1$ yields the following condition:

$$\text{rank} \begin{bmatrix} A - I & B_d \\ C_m & C_d \end{bmatrix} = n_x + n_d \quad (3.45)$$

For condition (3.45) to be satisfied, the number of disturbances in d needs to be smaller than or equal to the number of available measurements in y_m , $n_d \leq n_{y_m}$. A pair of matrices (B_d, C_d) such that (3.45) is satisfied always exists. Intuitively, this means that at steady state and given $y_{m,s}$, d_s must be uniquely defined:

$$\begin{bmatrix} A - I & B_d \\ C_m & C_d \end{bmatrix} \cdot \begin{bmatrix} x_s \\ d_s \end{bmatrix} = \begin{bmatrix} 0 \\ y_{m,s} \end{bmatrix}$$

The state and disturbance estimator is designed based on the augmented model as follows:

$$\begin{bmatrix} \hat{x}_{k+1} \\ \hat{d}_{k+1} \end{bmatrix} = \begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} \cdot \begin{bmatrix} \hat{x}_k \\ \hat{d}_k \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \cdot u_k + \begin{bmatrix} L_x \\ L_d \end{bmatrix} \cdot (-y_{m,k} + C_m \cdot \hat{x}_k + C_d \cdot \hat{d}_k)$$

where \hat{x} and \hat{d} are estimates of the state and disturbance. The error dynamics are:

$$\begin{aligned} \begin{bmatrix} x_{k+1} - \hat{x}_{k+1} \\ d_{k+1} - \hat{d}_{k+1} \end{bmatrix} &= \begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} \cdot \begin{bmatrix} x_k - \hat{x}_k \\ d_k - \hat{d}_k \end{bmatrix} - \begin{bmatrix} L_x \\ L_d \end{bmatrix} \cdot (-C_m \cdot x_k - C_d \cdot d_k + C_m \cdot \hat{x}_k + C_d \cdot \hat{d}_k) \\ &= \left(\begin{bmatrix} A & B_d \\ 0 & I \end{bmatrix} + \begin{bmatrix} L_x \\ L_d \end{bmatrix} \cdot [C_m \quad C_d] \right) \cdot \begin{bmatrix} x_k - \hat{x}_k \\ d_k - \hat{d}_k \end{bmatrix} \end{aligned}$$

where $L = \begin{bmatrix} L_x \\ L_d \end{bmatrix}$ is chosen to ensure that the error dynamics are stable and converge to zero, i.e. the estimator is stable. Assuming the observer is stable, we have:

$$\det \left(\begin{bmatrix} A + L_x \cdot C_m - \lambda \cdot I & B_d + L_x \cdot C_d \\ L_d \cdot C_m & I + L_d \cdot C_d - \lambda \cdot I \end{bmatrix} \right) \neq 0$$

This must also hold for $\lambda = 1$, i.e. a pole at $(1, 0)$:

$$\det \left(\begin{bmatrix} A + L_x \cdot C_m - I & B_d + L_x \cdot C_d \\ L_d \cdot C_m & L_d \cdot C_d \end{bmatrix} \right) \neq 0 \quad (3.46)$$

For (3.46) to hold, the last n_d rows of the matrix must be full row rank. Thus, a necessary condition is that $L_d \in \mathbb{R}^{n_d \times n_{y_m}}$ has full row rank, requiring $n_d \leq n_{y_m}$. This also follows from the fact that the disturbance dynamics need to converge to zero:

$$L_d \cdot \left(y_{m,k} - [C_m \quad C_d] \cdot \begin{bmatrix} \hat{x} \\ \hat{d} \end{bmatrix} \right) = 0$$

For L_d to have only the zero vector in its nullspace, $n_d \geq n_{y_m}$ is required. Since $n_d \leq n_{y_m}$ from (3.45), we conclude $n_d = n_{y_m}$. The steady state of the observer then satisfies:

$$\begin{bmatrix} A-I & B \\ C_m & 0 \end{bmatrix} \cdot \begin{bmatrix} \hat{x}_\infty \\ u_\infty \end{bmatrix} = \begin{bmatrix} -B_d \cdot \hat{d}_\infty \\ y_{m,\infty} - C_d \cdot \hat{d}_\infty \end{bmatrix}$$

Consequently, the observer output $C_m \cdot \hat{x}_\infty + C_d \cdot \hat{d}_\infty$ tracks the measurement $y_{m,\infty}$ without offset. Note that with $n_d = n_{y_m}$, observability of the augmented system in (3.41) is given for the simple choice $C_d = I$ if and only if (A, C_m) is observable and $\det(A - I - B_d \cdot C_m) \neq 0$. If the system has no integrators, then $B_d = 0$ can be chosen as long as $\det(A - I) \neq 0$.

The estimator provides \hat{x}_k and \hat{d}_k at each time k . To find suitable control inputs, we can modify the steady-state target problem to account for the nonzero disturbance, setting $\hat{d}_s = \hat{d}_k$:

$$\underset{x_s, u_s}{\text{minimize}} \quad \frac{1}{2} \cdot (\|C \cdot x_s - (y_{\text{ref}} - T_y \cdot C_d \cdot \hat{d}_s)\|_{Q_s}^2 + \|u_s\|_{R_s}^2) \quad (3.47a)$$

$$\text{subject to} \quad \begin{bmatrix} I-A & -B \\ T_y \cdot C_m & 0 \end{bmatrix} \cdot \begin{bmatrix} x_s \\ u_s \end{bmatrix} = \begin{bmatrix} B_d \cdot \hat{d}_s \\ y_{\text{ref}} - T_y \cdot C_d \cdot \hat{d}_s \end{bmatrix}, \quad (3.47b)$$

$$C_x \cdot x_s + C_u \cdot u_s \leq b \quad (3.47c)$$

If (3.47c) considers constraints on the measured outputs y_m or tracked outputs y , the disturbance \hat{d}_s needs to be considered in the formulation of b . All the concepts discussed in this section can be extended to systems with direct feedthrough. Given the steady-state target, the same dynamic regulation problem as presented in (3.40) is used for the regulator. This approach ensures the references and constraints are met statically, but not dynamically. The approach is illustrated in Fig. 3.3.

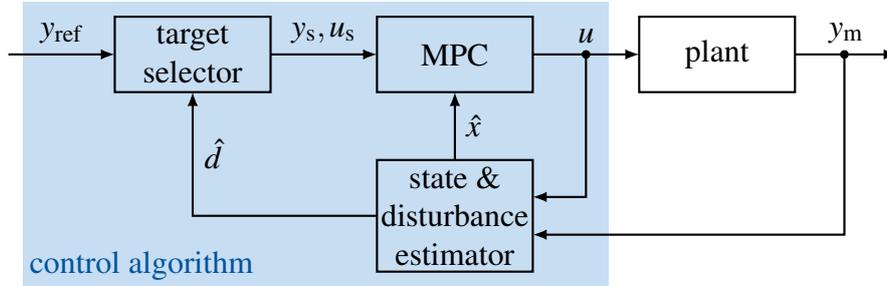


Fig. 3.3: Two-layered control structure: MPC regulator, state and disturbance estimator, and target selector.

Another approach to achieve offset-free reference tracking is to use the MPC problem formulation in (3.36) and modify the reference as follows:

$$\tilde{y}_{\text{ref}} = y_{\text{ref}} - T_y \cdot C_d \cdot \hat{d}_k \quad (3.48)$$

However, if constraints on the measured or tracked outputs are relevant, this approach does not consider the disturbances in those constraints. [16, 76, 89, 92]

3.4 Bayesian optimization

In this section, we introduce Bayesian optimization (BO), an effective method for finding the extrema of objective functions that are expensive to evaluate. Mathematically, BO addresses the problem of identifying the global maximizer (or minimizer) of an unknown objective function Φ , formulated as follows:

$$\theta^* = \arg \max_{\theta \in \Theta} \Phi(\theta) \quad (3.49)$$

where Θ represents the search space of interest. In global optimization, Θ is often a compact subset of \mathbb{R}^{n_θ} , but the BO framework can also be applied to non-continuous search spaces. Typically, black box optimization requires that all dimensions have bounds on the search space. Here, we assume these bounds are axis-aligned, so the search space forms a hyperrectangle of dimension n_θ . We further assume the black-box function Φ has no closed-form expression but can be evaluated at any point θ in the domain. This evaluation produces noise-corrupted (stochastic) outputs $z \in \mathbb{R}$ such that $\mathbb{E}[z|\Phi(\theta)] = \Phi(\theta)$. Therefore, we can only observe Φ through unbiased noisy point-wise observations z . [19, 100]

In this context, we employ a sequential search algorithm. At each iteration n , the algorithm selects a location θ_{n+1} to query Φ and observe z_{n+1} . After N queries, it provides a final recommendation $\bar{\theta}_N$, which is its best estimate of the optimizer. The BO framework is very efficient in terms of the number of function evaluations required, making it particularly useful when evaluations of Φ are costly, derivatives with respect to θ are unavailable, and Φ is potentially nonconvex and multimodal. In these scenarios, BO effectively leverages the complete optimization history to enhance the search efficiency. [100]

Fundamentally, BO is a sequential, model-based approach to solving the optimization problem (3.49). The data efficiency of BO arises from its ability to incorporate prior belief about the problem, guiding the sampling process and balancing exploration and exploitation of the search space. Based on Bayes' theorem, this prior model is sequentially refined through Bayesian posterior updating as new data points are observed:

$$P(\Phi(\theta) | \mathcal{D}_{n+1}) \propto P(z_{n+1} | \theta_{n+1}, \Phi(\theta)) \cdot P(\Phi(\theta) | \mathcal{D}_n)$$

Here, $\mathcal{D}_n = \{(\theta_i, z_i)\}_{i=1}^n$ represents the set of observed data points at iteration n , $P(\Phi(\theta) | \mathcal{D}_n)$ is the prior distribution, i.e. our prior belief about the space of possible objective functions, before observing (θ_{n+1}, z_{n+1}) , and $P(z_{n+1} | \theta_{n+1}, \Phi(\theta))$ is the likelihood of observing the new data point z_{n+1} given the new input θ_{n+1} and the current model of the objective function Φ . The posterior distribution $P(\Phi(\theta) | \mathcal{D}_{n+1})$ represents our updated beliefs about the objective function based on the new data. Using this probabilistic model, BO constructs acquisition functions $a_n : \Theta \mapsto \mathbb{R}$ to guide the search for the optimum. The algorithm selects the next query point θ_{n+1} by maximizing a_n , with index n reflecting the implicit dependence on the current data. The maxima are located where the uncertainty in the posterior, i.e. its variance,

is high (exploration) and/or where the model prediction, i.e. the posterior mean, is optimal (exploitation). This approach effectively balances exploration and exploitation, and aims to minimize the number of objective function evaluations. [19, 100]

Fig. 3.4 shows a typical run of BO on a 1D problem. The optimization process begins with two noisy evaluations of the cost function Φ . Based on these observations, a model, commonly referred to as the surrogate, is built to predict Φ at new points. This surrogate model is typically implemented using Gaussian process regression (GPR). While other surrogate models exist, they are not discussed here. The probabilistic predictions from the Gaussian process (GP) help construct the acquisition function, which determines where to sample the objective function next. The acquisition function considers the mean and variance of the predictions to model the utility of sampling. At each iteration, the acquisition function is maximized, and the objective is sampled at the maximum point. The GP is then updated with this new data, and the process is repeated until the desired number of iterations is reached or a stopping criterion is met. [19] These steps are illustrated in Fig. 3.5 and a pseudocode of the BO framework is given in Algorithm 1.

Algorithm 1 Bayesian optimization [100]

- 1: **for** $n = 1, 2, \dots$, **do**
 - 2: select new θ_{n+1} by optimizing acquisition function a

$$\theta_{n+1} = \arg \max_{\theta \in \Theta} a(\theta; \mathcal{D}_n)$$
 - 3: query objective function to obtain z_{n+1}
 - 4: augment data $\mathcal{D}_{n+1} = \{\mathcal{D}_n, (\theta_{n+1}, z_{n+1})\}$
 - 5: update statistical model
 - 6: **end for**
-

A GP is a distribution over functions, defined by its mean function and covariance function (also known as the kernel). Unlike a typical function that returns a scalar $\Phi(\theta)$ for an input θ , a GP provides the mean and variance of a normal distribution over the possible values of Φ at θ . In BO, defining the prior and particularly the covariance function is crucial. Important kernels include the squared exponential kernel [91] and the Matérn kernel [71, 107]. The choice of kernel hyperparameters significantly affects the performance. The selection of acquisition functions is also critical in BO. Common acquisition functions include the upper confidence bound [106], expected improvement [72], probability of improvement [65], and Thompson sampling [111]. [19, 100] A detailed discussion of kernels and acquisition functions is beyond the scope of this thesis.

3 Control preliminaries

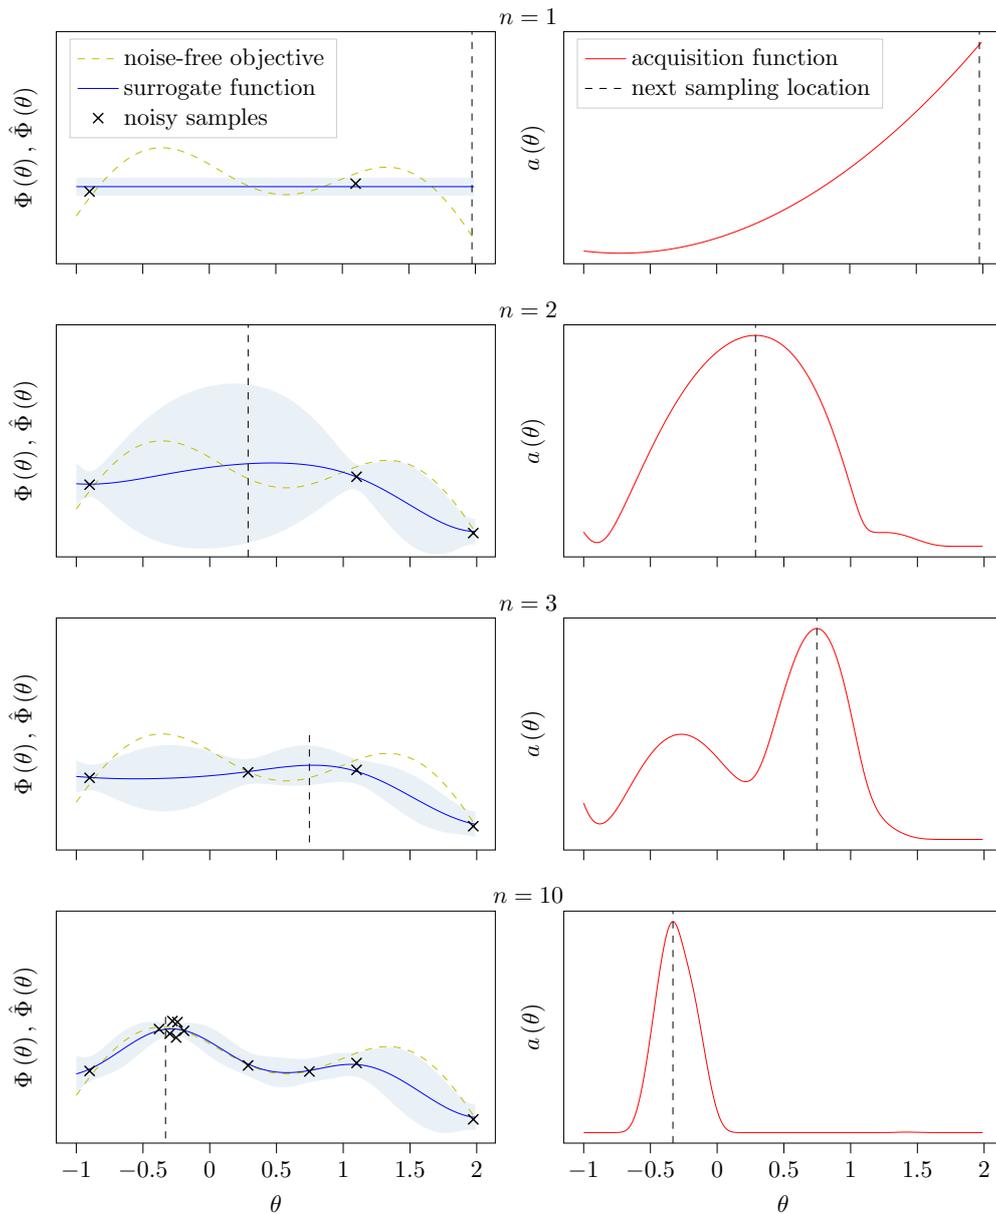


Fig. 3.4: Illustration of the BO procedure over selected iterations. The left column plots show the estimated mean and confidence intervals of the objective function using a probabilistic model. While the actual objective function is displayed here for illustration, it remains unknown in practice. The right column plots depict the acquisition functions. The acquisition function peaks in regions where the model predicts a high objective value (exploitation) and where prediction uncertainty is high (exploration). [100]

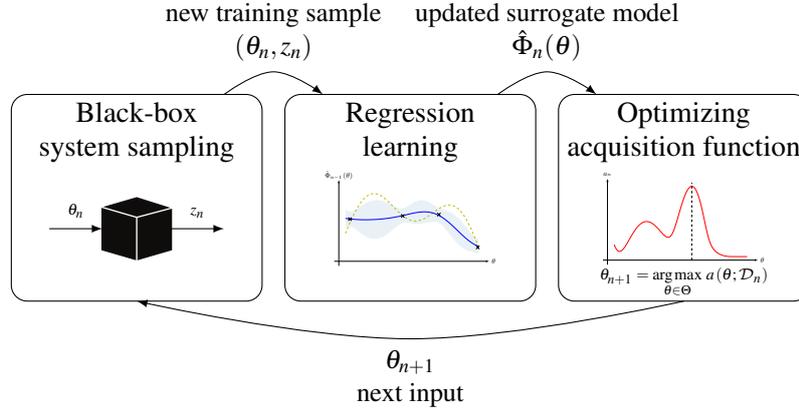


Fig. 3.5: Illustration of the steps in BO, adapted from [29].

Multi-objective optimization

In multi-objective optimization (MOO), the aim is to find a compromise between multiple objectives rather than a single optimal value. Thus, the MOO problem seeks to maximize (or minimize) a vector-valued objective function $\Phi(\cdot)(\theta) : \Theta \subset \mathbb{R}^{n_\theta} \rightarrow \mathcal{Y} \subset \mathbb{R}^{m_{\text{obj}}}$:

$$\Theta^* = \arg \max_{\theta \in \Theta} \Phi(\theta) = \arg \max_{\theta \in \Theta} \left[\Phi_1(\theta), \dots, \Phi_{m_{\text{obj}}}(\theta) \right]^\top$$

Here, $m_{\text{obj}} \geq 2$ represents the number of conflicting objectives, with each objective $\Phi_i : \Theta \rightarrow \mathbb{R}$. Θ is the design space, and $\mathcal{Y} = \Phi(\Theta)$ is the performance space. [1] Usually, no single solution maximizes all m_{obj} objectives simultaneously. Instead, we seek a set of optimal solutions called the Pareto set $\mathcal{P}_\Theta \subseteq \Theta$ in the design space. The corresponding set of optimal objective values in the performance space is known as the Pareto front $\mathcal{P}_\mathcal{Y} = \Phi(\mathcal{P}_\Theta) \subset \mathbb{R}^{m_{\text{obj}}}$. [63]

To understand this better, we need to clarify the concept of dominance. Given two points $\theta, \theta' \in \Theta$ with associated objective vectors $\Phi(\theta)$ and $\Phi(\theta')$, $\Phi(\theta)$ is said to dominate, or Pareto-dominate, $\Phi(\theta')$, denoted as $\Phi(\theta) \succeq \Phi(\theta')$, if and only if $\Phi_i(\theta) \geq \Phi_i(\theta')$ for all $i = 1, \dots, m_{\text{obj}}$. The dominance is strict, denoted as $\Phi(\theta) \succ \Phi(\theta')$, if $\Phi(\theta) \succeq \Phi(\theta')$ and there exists at least one $i \in \{1, \dots, m_{\text{obj}}\}$ such that $\Phi_i(\theta) > \Phi_i(\theta')$. In terms of sets, a set $\mathcal{A} \subset \mathbb{R}^{m_{\text{obj}}}$ is said to Pareto-dominate a set $\mathcal{B} \subset \mathbb{R}^{m_{\text{obj}}}$, written $\mathcal{A} \succeq \mathcal{B}$, if and only if $\forall \mathbf{b} \in \mathcal{B}, \exists \mathbf{a} \in \mathcal{A}$ such that $\mathbf{a} \succeq \mathbf{b}$. [40, 124]

Finding all the non-dominated points leads to the Pareto frontier. A Pareto-optimal solution is a point $\theta^* \in \Theta$ for which it is impossible to find another solution $\theta \in \Theta$ such that $\Phi(\theta) \succeq \Phi(\theta^*)$. It represents an optimal compromise in the sense that it is not possible to find a competitor being better or equal to in all objectives simultaneously. Thus, within the set of Pareto-optimal solutions, improving one objective leads to worsening another. The Pareto set \mathcal{P}_Θ is defined as [40]:

$$\mathcal{P}_\Theta = \{ \theta \in \Theta : \nexists \theta' \in \Theta, \Phi(\theta') \succeq \Phi(\theta) \}$$

The Pareto front $\mathcal{P}_{\mathcal{Y}}$ is the image of the Pareto set and contains only non-dominated solutions [40]:

$$\mathcal{P}_{\mathcal{Y}} = \Phi(\mathcal{P}_{\Theta}) = \{\Phi(\theta) \in \mathcal{Y} : \nexists \Phi(\theta') \in \mathcal{Y}, \Phi(\theta') \preceq \Phi(\theta)\}$$

Given a Pareto frontier, the context of the optimization will help decide which solution to pick. The goal of an MOO algorithm is therefore to identify an approximate Pareto frontier of the true Pareto frontier within a specified budget of function evaluations. [27]

To measure the quality of an approximated Pareto front, the hypervolume indicator is commonly used in MOO. The hypervolume indicator $I_H(\hat{\mathcal{P}}_{\mathcal{Y}}; \mathbf{r})$ of a non-dominated set $\hat{\mathcal{P}}_{\mathcal{Y}} \subset \mathbb{R}^{m_{\text{obj}}}$ is the m_{obj} -dimensional volume lower-bounded by a reference point $\mathbf{r} \in \mathbb{R}^{m_{\text{obj}}}$, which is dominated by at least one $\Phi \in \hat{\mathcal{P}}_{\mathcal{Y}}$ [40]:

$$I_H(\hat{\mathcal{P}}_{\mathcal{Y}}; \mathbf{r}) = \text{Vol}\left(\bigcup_{\Phi \in \hat{\mathcal{P}}_{\mathcal{Y}}} \{\mathbf{z} : \mathbf{r} \preceq \mathbf{z} \preceq \Phi\}\right)$$

The reference point is typically provided by the practitioner based on domain knowledge. [128] A larger hypervolume indicates a better approximation of the true Pareto front by $\hat{\mathcal{P}}_{\mathcal{Y}}$. Thus, our aim is to find the set $\hat{\mathcal{P}}_{\mathcal{Y}}$ that maximises the hypervolume. See Fig. 3.6 for an illustration of some of the concepts introduced in this section.

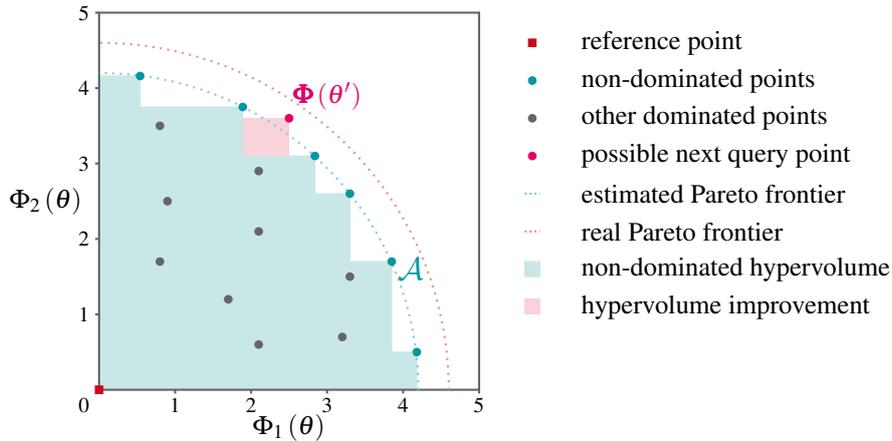


Fig. 3.6: Illustration of MOO concepts with $m_{\text{obj}} = 2$ objectives, adapted from [40]: The non-dominated objective vectors are depicted as blue circles. The hypervolume indicator of this non-dominated set \mathcal{A} with respect to the red reference point is the surface of the blue shaded region. The pink rectangle is the hypervolume improvement brought by $\Phi(\theta')$.

MOO problems are often addressed using evolutionary algorithms such as NSGA-II. [28] However, evolutionary algorithms generally suffer from high sample-complexity, making them unsuitable for scenarios with limited evaluation budgets. [27]

Multi-objective bayesian optimization

To avoid the slow convergence of evolutionary algorithms, Bayesian methods have been extended to perform efficient global optimization in a multi-objective setting. Thus, MOBO is designed to solve multi-objective problems that are expensive to evaluate. [63]

In general, m_{obj} GPs $\hat{\Phi}_i(\theta)$ are fitted to each objective $\Phi_i(\theta)$ independently. The algorithm conforms to the outline of Fig. 3.5 and Algorithm 1, except that m_{obj} surrogates $\hat{\Phi}_1(\theta), \dots, \hat{\Phi}_m(\theta)$ and m_{obj} objective functions are now considered, and that an empirical Pareto set $\hat{\mathcal{P}}_{\Theta}$ and Pareto front $\hat{\mathcal{P}}_{\mathcal{Y}}$ are returned. As in the single objective case, the problem is cast into the sequential optimization of an acquisition function used for determining $\theta_{n+1} \in \Theta$, the most promising next iterate to be evaluated. A popular acquisition function for MOBO is the expected hypervolume improvement (EHVI). In EHVI, the multi-objective improvement metric is the growth of the hypervolume indicator. Given the posterior GPs $\hat{\Phi}(\theta) := [\hat{\Phi}_1(\theta), \dots, \hat{\Phi}_m(\theta)]^\top$ based on the current observations $\mathcal{D}_n = \{(\theta_i, \mathbf{z}_i)\}_{i=1}^n$ with each observation $(\theta_i, \mathbf{z}_i) \in \mathbb{R}^{n_\theta} \times \mathbb{R}^{m_{\text{obj}}}$, the EHVI acquisition function is defined by:

$$a_n(\theta) = \mathbb{E}_{\hat{\Phi}(\theta)|\mathcal{D}_n} \left[I_H \left(\hat{\mathcal{P}}_{\mathcal{Y}} \cup \{ \hat{\Phi}(\theta) \}; \mathbf{r} \right) - I_H \left(\hat{\mathcal{P}}_{\mathcal{Y}}; \mathbf{r} \right) \right]$$

It represents the expected change in the hypervolume of the dominated region in the performance space when a new candidate θ is evaluated. Thus, the formula calculates the expectation of the improvement in hypervolume between the current Pareto front and the Pareto front updated with the new candidate point. Acquisition functions like EHVI aim at providing new non-dominated points while balancing exploitation and exploration, and eventually have the goal of approximating the Pareto frontier entirely. [27] This is achieved by the iterative process illustrated in Fig. 3.7.

While the algorithm is carried out offline and does not require real-time execution, this is acceptable in the present context since each closed-loop evaluation is computationally expensive and potentially safety-critical. Consequently data-efficient global optimization is particularly valuable. One limitation of BO algorithms is that they are typically most effective for problems with a limited number of optimization variables, generally ranging from 3 to 10 variables. In this work, we apply MOBO for control parameter tuning. Thus, selected outputs from closed-loop control evaluations serve as the black-box functions we aim to optimize, while the control design parameters of our MPC framework are the BO optimization variables. For implementation, we utilize the Ax platform [78] for its out-of-the-box application of BO algorithms.

In the context of MPC, the optimization problem is usually mathematically well-defined, and suboptimal solutions are accepted if computational time is limited to ensure real-time feasibility of the algorithm. In contrast, BO is a global optimization technique that does not require a known mathematical description of the optimization problem. It is particularly useful for hyperparameter tuning in machine learning, experimental design, and other areas

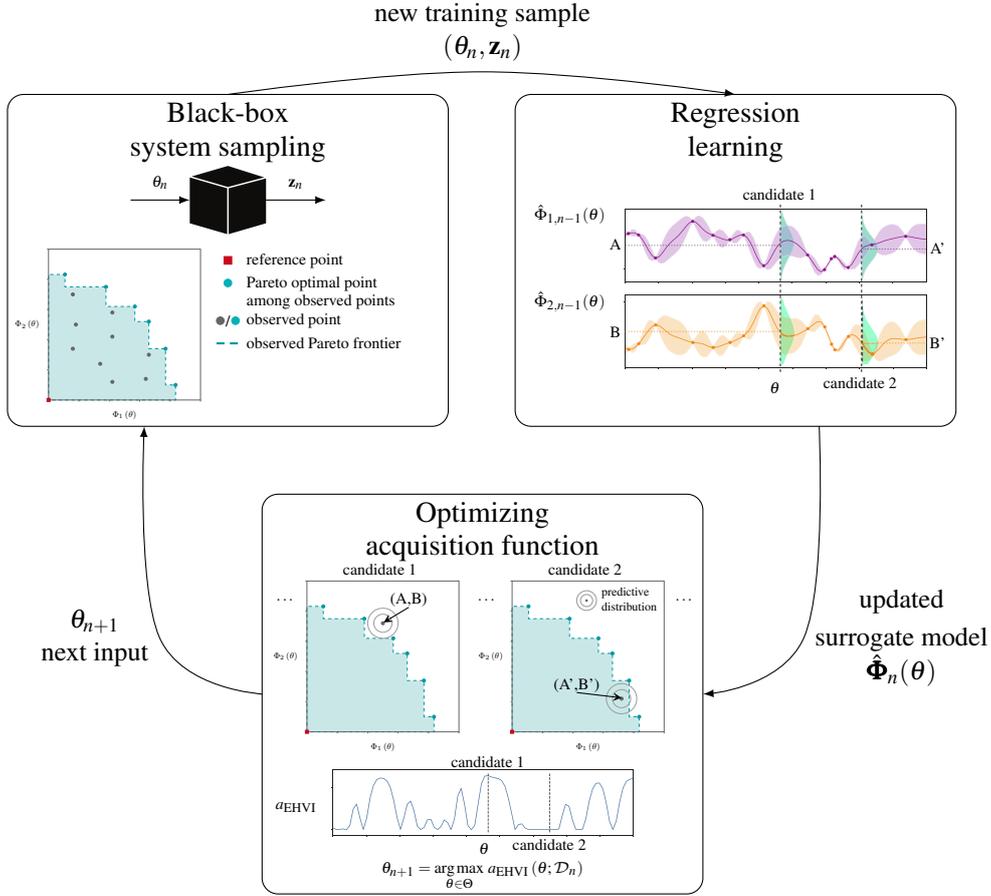


Fig. 3.7: Illustration of the MOBO framework with $m_{\text{obj}} = 2$ objectives, adapted from [61]. At iteration n , the black-box system is evaluated at θ_n , yielding a noisy observation \mathbf{z}_n of the objective vector. Based on the accumulated data, surrogate models are learned for each objective and an acquisition function, here based on the expected hypervolume improvement, is optimized to select the next query point θ_{n+1} .

where the function to be optimized is unknown and the function evaluations are expensive or time-consuming. These applications typically do not require real-time optimization.

These aspects also apply to the use case considered in this thesis. The control performance of the proposed MPC framework in closed-loop with a comprehensive plant model is a black-box function of the control design parameters, that we want to optimize globally. In our case, we aim to optimize multiple competing objectives, hence the application of MOBO. While the algorithm is applied offline and does not require real-time optimization, we still benefit from the efficient sampling characteristics of BO algorithms to reduce closed-loop control evaluations.

4 Design of model predictive control framework

This chapter details the design of each component of the MPC framework. It begins with an overview of the hierarchical control structure in relation to the control task. Next, the model equations used for control-oriented modeling in the prediction models are presented. This is followed by the formulation of the OCP and its problem transcription for each controller. Special focus is given to the NMPC design of the low-level tracking controller, where real-time capability is a primary goal. Finally, the design of the state and disturbance estimator is addressed.

4.1 Control task and control structure

The control task considered in this work for automotive FCSs in a FCHEV includes:

- i. Optimizing the power distribution between the FCS and the battery given a power demand $P_{\text{drive,ref}}$
- ii. Manipulating the system-level control levers of the FCS and the battery load current to meet their respective optimized power references, $P_{\text{fcs,bus,ref}}$ and $P_{\text{bat,bus,ref}}$

To address this task, we adopt the control structure depicted in Fig. 4.1, comprising a high-level control layer and a low-level control layer.

The high-level control layer includes a power-split NMPC and a target selector. The NMPC manages the power distribution between the FCS and the battery system, executing on a 1 s sampling time. The underlying economic cost function maximizes FCS efficiency while maintaining a stable battery SoC. An equality constraint enforces adherence to the total power demand $P_{\text{drive,ref}}$, if feasible. This optimization problem is based on a dynamic single-state battery model and a static map that relates the FCS power to its efficiency. More details on the design of the power-split NMPC are provided in Section 4.3.2.

The target selector, based on a static FCS model, employs a static optimization to derive a closed-form solution for the maximum admissible FCS efficiency, given the required FCS power output $P_{\text{fcs,bus,ref}}$. From this solution, we also obtain the compressor's optimal operating

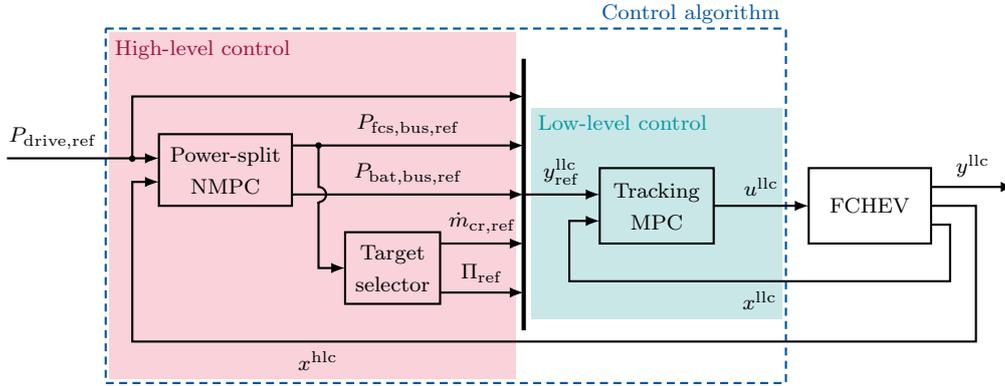


Fig. 4.1: Simplified schematic of the control structure.

point, defined by the corrected compressor mass flow rate $\dot{m}_{cr,ref}$ and pressure ratio Π_{ref} . The design of the target selector is explained in Section 4.3.1.

The low-level control comprises a tracking MPC that directly manipulates the FCHEV. This controller operates with a faster sampling rate, integrating a dynamic model of the FCS (see Section 4.2.1 to Section 4.2.4) and a static model of the battery (see Section 4.2.5). The references for the low-level MPC, y_{ref}^{llc} , include the total power demand, the optimal power distribution from the power-split NMPC, and the compressor's optimal operating point. Its control objectives are:

- i. Precise tracking of the power demand
- ii. Maintaining adequate OER
- iii. Avoiding compressor surge and choke
- iv. Following the optimized compressor setpoints
- v. Ensuring real-time feasibility

The design of the tracking MPC as an NMPC is discussed in Section 4.4.

4.2 Control-oriented modeling

The control-oriented model of the FCHEV consists of a FCS and a battery, as depicted in Fig. 4.2. While the prediction model of the high-level control considers the FCS statically and the battery system dynamically, the prediction model utilized in the low-level control utilizes the dynamic model of the FCS and a static model of the battery system. This section provides the relevant physical relations for both control layers.

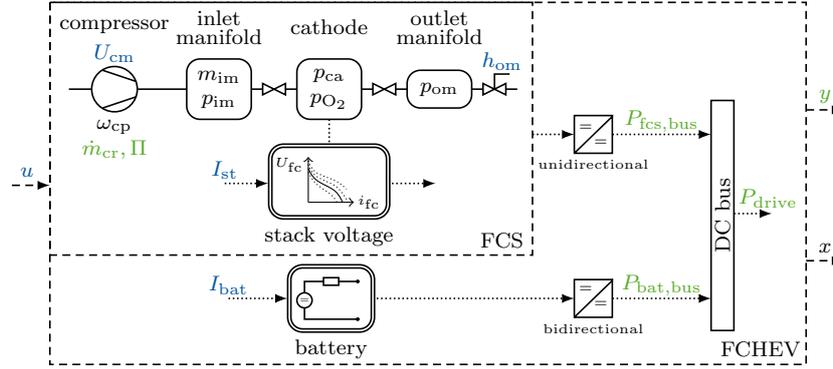


Fig. 4.2: Schematic overview of the model structure used for control-oriented modeling the FCHEV, focusing on air-path dynamics without humidity considerations.

Generally, a FCS consists of the FC stack itself, as well as its peripheral subsystems, such as the oxygen supply subsystem, hydrogen supply subsystem, water management subsystem, and thermal management subsystem. In this work, we focus on the control levers associated with the oxygen supply subsystem which will be referred to as the air-supply or air-path system, as oxygen is supplied by pressurizing the ambient air through an electric compressor. The air-path components, including the compressor, inlet manifold, cathode, and outlet manifold, as well as the stack voltage, are modeled based on [83]. The zero-dimensional model treats each air-path component as a single control volume, akin to a continuously-stirred tank reactor, without any spatial variations. The dynamic states of the model stem from this lumped-parameter approach to the air-path system, while the stack voltage as well as the battery are described by static algebraic expressions. For completeness, we provide most of the model equations in the subsequent sections.

4.2.1 Compressor model

A lumped rotational inertia J_{cp} is used to represent the dynamic behavior of the compressor with the compressor rotational speed ω_{cp} :

$$J_{cp} \cdot \frac{d\omega_{cp}}{dt} = \tau_{cm} - \tau_{cp} \quad (4.1)$$

Here, τ_{cm} is the compressor motor torque and τ_{cp} is the load torque required to drive the compressor. The compressor motor torque is calculated using a simplified static DC motor model:

$$\tau_{cm} = \frac{\eta_{cm} \cdot k_t}{R_{cm}} \cdot (U_{cm} - k_v \cdot \omega_{cp}), \quad (4.2)$$

where k_t , R_{cm} and k_v are motor constants, η_{cm} is the motor mechanical efficiency and the voltage U_{cm} is the motor control command. The torque required to drive the compressor is

calculated using the thermodynamic equation

$$\tau_{cp} = \frac{c_p}{\omega_{cp}} \cdot \frac{T_{atm}}{\eta_{cp}} \cdot \left(\Pi^{\frac{\gamma-1}{\gamma}} - 1 \right) \cdot \dot{m}_{cp}^{out}, \quad (4.3)$$

where c_p is the constant-pressure specific heat capacity and γ is the ratio of the specific heat capacities of air, η_{cp} is the compressor efficiency, $\Pi = p_{im}/p_{atm}$ is the ratio of the pressure inside the inlet manifold (Section 4.2.2) and the atmospheric pressure, and T_{atm} is the atmospheric temperature. The compressor flow is

$$\dot{m}_{cp}^{out} = \dot{m}_{cr} \cdot \frac{\delta}{\theta},$$

where $\theta = T_{atm}/288\text{K}$ and $\delta = p_{atm}/1\text{atm}$. The corrected compressor flow $\dot{m}_{cr} = f(\Pi, \omega_{cp})$ and the efficiency $\eta_{cp} = f(\dot{m}_{cr}, \Pi)$ are modeled using static maps. Thermodynamic equations are used to calculate the exit air temperature

$$T_{cp}^{out} = T_{atm} + \frac{T_{atm}}{\eta_{cp}} \cdot \left[\left(\frac{p_{im}}{p_{atm}} \right)^{\frac{\gamma-1}{\gamma}} - 1 \right],$$

and the power consumed by the compressor motor is

$$P_{cm} = \frac{1}{R_{cm}} \cdot (U_{cm}^2 - k_v \cdot \omega_{cp} \cdot U_{cm}). \quad (4.4)$$

As stated earlier, the corrected compressor flow \dot{m}_{cr} is modeled via a static map. More specifically, we used a polynomial fit to approximate the model in [85]. The reference model as well as the resulting fit are given in Fig. 4.3. We applied the same approach for the compressor efficiency η_{cp} . The efficiency data used in [85] is depicted in Fig. 4.4 and the corresponding polynomial fit is plotted in Fig. 4.5.

4.2.2 Inlet manifold model

The inlet manifold includes pipe and stack manifold volumes between the compressor and the FCs. Its dynamics are governed by mass continuity and energy conservation laws, which yield differential equations for the inlet manifold mass m_{im} and pressure p_{im} :

$$\frac{dm_{im}}{dt} = \dot{m}_{cp}^{out} - \dot{m}_{im}^{out} \quad (4.5)$$

$$\frac{dp_{im}}{dt} = \frac{\gamma \cdot R_a}{V_{im}} \cdot (\dot{m}_{cp}^{out} \cdot T_{cp}^{out} - \dot{m}_{im}^{out} \cdot T_{im}) \quad (4.6)$$

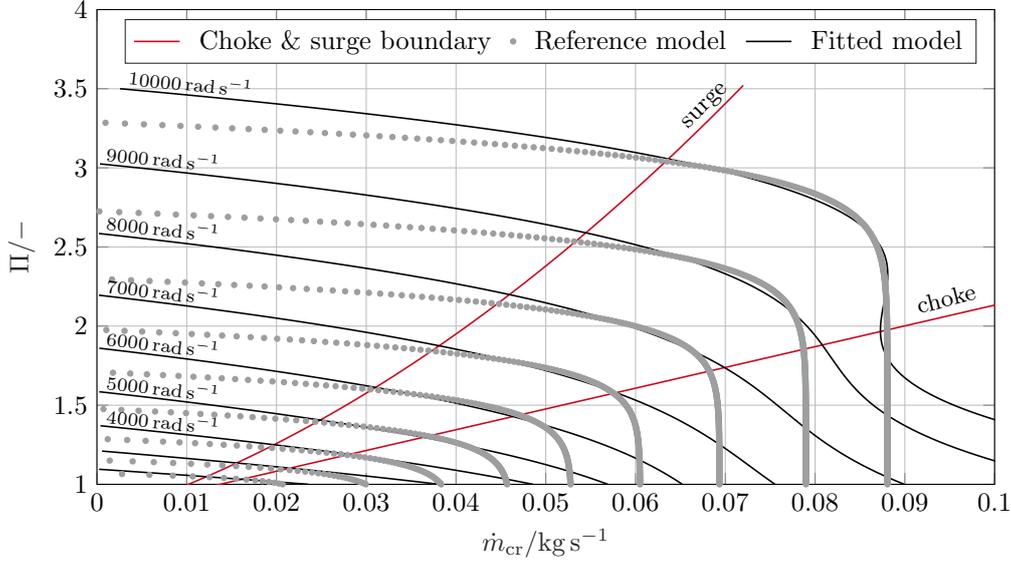


Fig. 4.3: Compressor map.

Here, R_a is the gas constant for ambient dry air, V_{im} is the inlet manifold volume, and T_{im} is the temperature of the flow inside the manifold, computed from the ideal gas law.

$$T_{im} = \frac{p_{im} \cdot V_{im}}{m_{im} \cdot R_a}$$

The exit flow, \dot{m}_{im}^{out} , is calculated as a function of p_{im} and the cathode pressure p_{ca} (Section 4.2.3.2) using a linearized nozzle flow equation with a nozzle constant $k_{ca,in}$:

$$\dot{m}_{im}^{out} = k_{ca,in} \cdot (p_{im} - p_{ca}) \quad (4.7)$$

The oxygen mass fraction in the inlet manifold's dry air, $y_{O_2}^{im}$, is derived from the oxygen mole fraction in the incoming ambient dry air, $x_{O_2}^{atm}$, using the molar masses of oxygen M_{O_2} and nitrogen M_{N_2} :

$$y_{O_2}^{im} = \frac{x_{O_2}^{atm} \cdot M_{O_2}}{x_{O_2}^{atm} \cdot M_{O_2} + (1 - x_{O_2}^{atm}) \cdot M_{N_2}} \quad (4.8)$$

Utilizing $y_{O_2}^{im}$, we can then determine the outgoing mass flow rate of oxygen, $\dot{m}_{O_2,im}^{out}$:

$$\dot{m}_{O_2,im}^{out} = y_{O_2}^{im} \cdot \dot{m}_{im}^{out}$$

4.2.3 Fuel cell stack model

In this study, the modeling of the FC stack is focused on two components: the stack voltage model and the cathode flow model. The original model [83] encompasses a broader range of

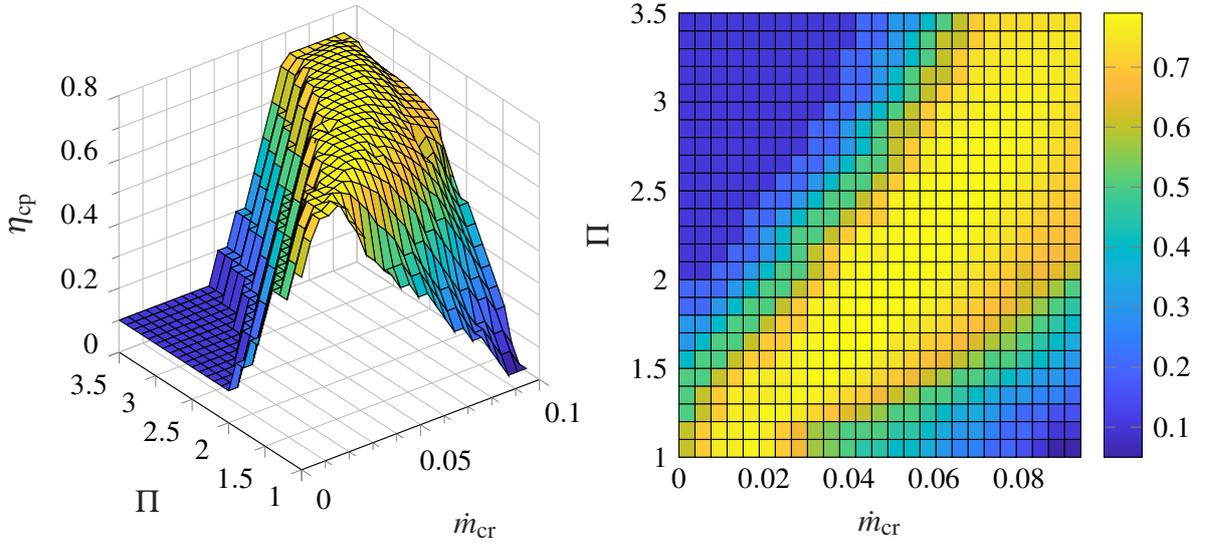


Fig. 4.4: Compressor efficiency data published in [85].

submodels, including an anode flow model and a membrane hydration model. In this adapted version, these are omitted to concentrate on the air-path dynamics.

4.2.3.1 Stack voltage model

Given that the FC stack comprises multiple FCs connected in series, and assuming that all cells are identical, the stack voltage U_{st} is calculated as the product of the number of cells and the cell voltage

$$U_{st} = n_{fc} \cdot U_{fc}, \quad (4.9)$$

and the stack current, I_{st} , is equal to the cell current. The combined effect of thermodynamics, kinetics, and ohmic resistance determines the cell's output voltage

$$U_{fc} = E_{ocv} - U_{act} - U_{ohm} - U_{conc}, \quad (4.10)$$

where E_{ocv} is the OCV, and U_{act} , U_{ohm} and U_{conc} are activation, ohmic and concentration overpotentials. These are calculated using a combination of physical and empirical relationships given in [83].

As we assume instantaneous electrochemical reaction, the FC voltage is a static function of the current density i_{fc} and the dynamically varying stack variables of oxygen p_{O_2} and hydrogen partial pressures p_{H_2} , cathode pressure p_{ca} , stack temperature T_{st} , and membrane humidity Λ_{mem} . Here, our focus is on the dynamics of the air-path variables, specifically p_{O_2} and p_{ca} . Consequently, we assume perfect control of the anode pressure, i.e. $p_{an} = p_{ca}$. We further simplify the model by assuming lack of vapor and, thus, posit that the hydrogen partial

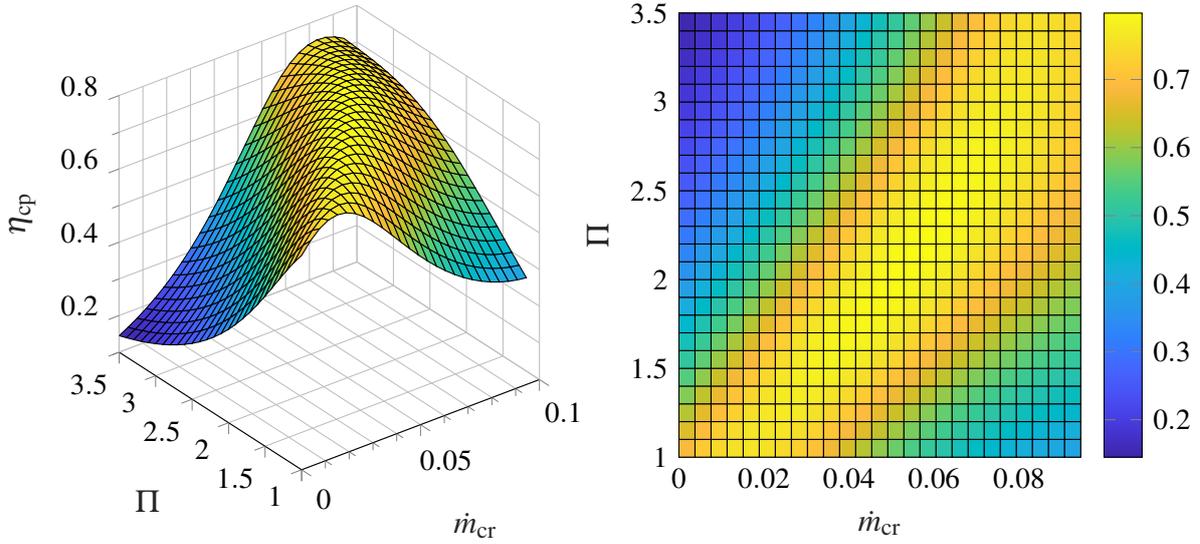


Fig. 4.5: Fit of compressor efficiency η_{cp} .

pressure mirrors the anode pressure, i.e. $p_{H_2} = p_{an}$. Additionally, we assume that T_{st} and Λ_{mem} are accurately controlled with negligible delay to constant values of 80°C and 12.5, respectively.

Similar to the reversible OCV given in (2.4), the OCV in (4.10) is calculated as follows:

$$E_{ocv} = 1.229 - 8.5 \cdot 10^{-4} \cdot (T_{fc} - 298.15) + 4.3085 \cdot 10^{-5} \cdot T_{fc} \cdot \left[\ln(p_{H_2}) + \frac{1}{2} \cdot \ln(p_{O_2}) \right]$$

where the FC temperature T_{fc} , assumed equal to the stack temperature T_{st} , is expressed in Kelvin, and reactant partial pressures p_{H_2} and p_{O_2} are expressed in atm. The relationship between the activation overpotential, U_{act} , and the current density is described by the Tafel equation, which is approximated by:

$$U_{act} = E_0 + E_a \cdot (1 - e^{-c_1 \cdot i_{fc}}) \quad (4.11)$$

The activation overvoltage depends on temperature and oxygen partial pressure. The values for E_0 , E_a and c_1 and their dependency on oxygen partial pressure and temperature can be determined from nonlinear regression of experimental data. The voltage drop caused by the ohmic overpotential, U_{ohm} , is proportional to the stack current:

$$U_{ohm} = i_{fc} \cdot R_{ohm}$$

The resistance, R_{ohm} , depends strongly on membrane humidity and cell temperature. It is proportional to membrane thickness t_{mem} and inversely proportional to the membrane conductivity, σ_{mem} , i.e.:

$$R_{ohm} = \frac{t_{mem}}{\sigma_{mem}}$$

The membrane conductivity is a function of membrane humidity and temperature in the form:

$$\sigma_{\text{mem}} = (b_{11} \cdot \Lambda_{\text{mem}} - b_{12}) \cdot \exp\left(b_2 \cdot \left(\frac{1}{303} - \frac{1}{T_{\text{fc}}}\right)\right) \quad (4.12)$$

where the value of t_{mem} , b_{11} and b_{12} for the NAFION[®] 117 membrane are used, and b_2 is adjusted to fit FC data in. An equation to approximate the concentration overpotential, U_{conc} , is given by:

$$U_{\text{conc}} = i_{\text{fc}} \cdot \left(c_2 \frac{i_{\text{fc}}}{i_{\text{fc,max}}}\right)^{c_3} \quad (4.13)$$

where c_2 , c_3 and $i_{\text{fc,max}}$ are constants that depend on temperature and reactant partial pressure and can be determined empirically. The coefficients in equations (4.11), (4.12) and (4.13) are given in A.2.

4.2.3.2 Cathode flow model

This model describes the cathode airflow behavior, developed using the principle of mass conservation along with the thermodynamic properties of air. Mass continuity is applied to balance the mass of dry air and the mass of oxygen within the cathode volume. Neglecting the presence of water vapor, the mass of dry air is equal to the total mass in the cathode. Balancing the total mass, m_{ca} , and the mass of oxygen, m_{O_2} , and utilizing the ideal gas law, we obtain the differential equations for the total pressure and the partial pressure of oxygen in the cathode, p_{ca} and p_{O_2} , respectively.

$$\frac{dp_{\text{ca}}}{dt} = \frac{R_a \cdot T_{\text{ca}}}{V_{\text{ca}}} \cdot (\dot{m}_{\text{ca}}^{\text{in}} - \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{ca}}^{\text{out}}) \quad (4.14)$$

$$\frac{dp_{\text{O}_2}}{dt} = \frac{R \cdot T_{\text{ca}}}{M_{\text{O}_2} \cdot V_{\text{ca}}} \cdot (\dot{m}_{\text{O}_2}^{\text{in}} - \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2}^{\text{out}}) \quad (4.15)$$

Here, R is the universal gas constant, V_{ca} is the cathode volume, and the cathode air temperature T_{ca} is assumed equal to the stack temperature, T_{st} . The specific gas constant of dry air in (4.14), R_a , is calculated as follows:

$$R_a = \frac{R}{M_a} = \frac{R}{x_{\text{O}_2}^{\text{atm}} \cdot M_{\text{O}_2} + (1 - x_{\text{O}_2}^{\text{atm}}) \cdot M_{\text{N}_2}}, \quad (4.16)$$

where the usage of the atmospheric molar fraction of oxygen $x_{\text{O}_2}^{\text{atm}}$ instead of the molar fraction of oxygen in the cathode $x_{\text{O}_2}^{\text{ca}}$ introduces some minor inaccuracy.

The inlet and outlet mass flow rates are determined from the inlet and outlet flow conditions and using thermodynamic properties. The cathode inlet flow rate and condition are established in Section 4.2.2. The calculations for the cathode outlet flow rate and condition are

$$\dot{m}_{\text{ca}}^{\text{out}} = k_{\text{ca,out}} \cdot (p_{\text{ca}} - p_{\text{om}}),$$

$$\dot{m}_{\text{O}_2, \text{ca}}^{\text{out}} = y_{\text{O}_2}^{\text{ca}} \cdot \dot{m}_{\text{ca}}^{\text{out}}.$$

Here, $\dot{m}_{\text{ca}}^{\text{out}}$ is obtained analogously to (4.7), with nozzle constant $k_{\text{ca}, \text{out}}$ and outlet manifold pressure p_{om} (Section 4.2.4). Similarly, $y_{\text{O}_2}^{\text{ca}}$ is determined as in (4.8), with

$$x_{\text{O}_2}^{\text{ca}} = \frac{p_{\text{O}_2}}{p_{\text{ca}}}$$

Electrochemical principles are used to calculate the rate of oxygen consumption, $\dot{m}_{\text{O}_2}^{\text{react}}$, from the stack current I_{st} :

$$\dot{m}_{\text{O}_2}^{\text{react}} = \frac{M_{\text{O}_2} \cdot n_{\text{fc}}}{4 \cdot F} \cdot I_{\text{st}},$$

where F is the Faraday constant. The OER λ_{O_2} corresponds to the ratio between the oxygen supplied and the oxygen reacted:

$$\lambda_{\text{O}_2} = \frac{\dot{m}_{\text{O}_2, \text{im}}^{\text{out}}}{\dot{m}_{\text{O}_2}^{\text{react}}} \quad (4.17)$$

4.2.4 Outlet manifold model

Contrary to the inlet manifold, where temperature variations are significant, the outlet manifold temperature T_{om} is assumed constant and equal to the temperature of the flow leaving the cathode, which is the stack temperature T_{st} . The outlet manifold pressure, p_{om} , is then derived using mass conservation and the ideal gas law, assuming isothermal conditions:

$$\frac{dp_{\text{om}}}{dt} = \frac{R_a \cdot T_{\text{om}}}{V_{\text{om}}} \cdot (\dot{m}_{\text{ca}}^{\text{out}} - \dot{m}_{\text{om}}^{\text{out}}), \quad (4.18)$$

where V_{om} is the outlet manifold volume. The simplification made in (4.16) is applied here, too.

The air flow rate exiting the outlet manifold, $\dot{m}_{\text{om}}^{\text{out}}$, is calculated via the nonlinear nozzle equation for subcritical flow. It is dependent on p_{om} and the back-pressure valve opening area, $h_{\text{om}} \cdot A_{\text{T}}$, with C_{D} being the discharge coefficient of the nozzle:

$$\dot{m}_{\text{om}}^{\text{out}} = \frac{C_{\text{D}} \cdot h_{\text{om}} \cdot A_{\text{T}} \cdot p_{\text{om}}}{\sqrt{R \cdot T_{\text{om}}}} \cdot \left(\frac{p_{\text{atm}}}{p_{\text{om}}} \right)^{\frac{1}{\gamma}} \cdot \Psi \quad (4.19)$$

$$\Psi = \sqrt{\frac{2 \cdot \gamma}{\gamma - 1} \cdot \left[1 - \left(\frac{p_{\text{atm}}}{p_{\text{om}}} \right)^{\frac{\gamma - 1}{\gamma}} \right]} \quad (4.20)$$

4.2.5 Battery model

For the control purposes within this thesis, the battery is modeled using a static equivalent circuit model, which excludes RC branches and thus neglects voltage dynamics. The intentional simplification involves representing the battery with an OCV source, U_{ocv} , connected in series with an internal ohmic resistance, R_i , as illustrated in Fig. 4.6.

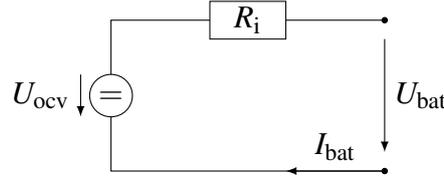


Fig. 4.6: Static equivalent circuit model.

The terminal voltage of the battery, denoted by U_{bat} , is the OCV minus the voltage drop due to the internal resistance:

$$U_{bat} = U_{ocv} - R_i \cdot I_{bat}$$

The OCV is modeled as the sum of the reference OCV potential E_0 , minus the polarization voltage E_{pol} , plus an exponential voltage term E_{exp} :

$$U_{ocv} = E_0 - E_{pol} + E_{exp} \quad (4.21)$$

where the constant voltage, E_0 , represents the OCV at full capacity. Both E_{pol} and E_{exp} are static functions of the battery SoC x_{SoC} :

$$E_{pol} = \frac{K}{x_{SoC}},$$

$$E_{exp} = A \cdot \exp(-B \cdot Q \cdot (1 - x_{SoC})), \quad (4.22)$$

where K is the polarization resistance coefficient, Q is the battery capacity, A is the exponential zone amplitude, and B is the inverse of the exponential zone time constant. [112]

The SoC of the battery is the amount of electrical charge stored in the battery relative to its capacity. Here, we calculate the SoC as follows:

$$x_{SoC}(t) = x_{SoC}(t_0) - \frac{1}{Q} \cdot \int_{t_0}^t I_{bat}(t) dt$$

where the current sign convention is set such that $I_{bat} > 0$ denotes discharging and $I_{bat} < 0$ charging. The model assumes the same characteristics for the charge and discharge cycles and

the charging/discharging current, I_{bat} , can be obtained from the net power output of the battery system, $P_{\text{bat,net}}$, defined by:

$$P_{\text{bat,net}} = U_{\text{bat}} \cdot I_{\text{bat}} \cdot n_{\text{packs}} \cdot n_{\text{bat,cells}} = (U_{\text{ocv}} \cdot I_{\text{bat}} - R_i \cdot I_{\text{bat}}^2) \cdot n_{\text{packs}} \cdot n_{\text{bat,cells}} \quad (4.23)$$

The elementary battery cells are connected in series, with n_{packs} packs, each consisting of $n_{\text{bat,cells}}$ cells. By solving the quadratic expression in (4.23) for the battery current we obtain:

$$I_{\text{bat}} = \frac{1}{2 \cdot R_i} \cdot \left(U_{\text{ocv}} - \sqrt{U_{\text{ocv}}^2 - \frac{4 \cdot P_{\text{bat,net}} \cdot R_i}{n_{\text{packs}} \cdot n_{\text{bat,cells}}}} \right)$$

The SoC dynamics can then be expressed as [59]:

$$\dot{x}_{\text{SoC}} = -\frac{I_{\text{bat}}}{Q} = -\frac{1}{2 \cdot Q \cdot R_i} \cdot \left(U_{\text{ocv}} - \sqrt{U_{\text{ocv}}^2 - \frac{4 \cdot P_{\text{bat,net}} \cdot R_i}{n_{\text{packs}} \cdot n_{\text{bat,cells}}}} \right) \quad (4.24)$$

4.2.6 Power equations

The power generated by the FC stack, P_{st} , is the product of the stack voltage and the stack current:

$$P_{\text{st}} = U_{\text{st}} \cdot I_{\text{st}}$$

To calculate the net power output of the FCS, $P_{\text{fcs,net}}$, we deduct the power used by the compressor motor, P_{cm} , and a fixed auxiliary power, P_{aux} , where P_{aux} accounts for the consumption of other actuators in the FCS:

$$P_{\text{fcs,net}} = P_{\text{st}} - P_{\text{cm}} - P_{\text{aux}}$$

The power losses that occur in the DC-DC converter are approximated by a second-order polynomial as follows:

$$P_{\text{fcs,bus}} = P_{\text{fcs,max}} \cdot \left(\sum_{i=1}^2 a_i \cdot \frac{P_{\text{fcs,net}}^i}{P_{\text{fcs,max}}^i} \right), \quad (4.25)$$

where $P_{\text{fcs,bus}}$ is the FCS power available at the output of the DC-DC converter, $P_{\text{fcs,max}}$ is the expected maximum FCS power, and a_i are polynomial coefficients obtained from DC-DC converter efficiency data.

For the battery system, the net power, $P_{\text{bat,net}}$, is the product of the battery voltage and current, factoring in the quantity of battery packs and cells, see (4.23). Contrary to the FCS, losses within the battery are already included in the battery voltage by considering the internal ohmic resistance:

$$P_{\text{bat,net}} = P_{\text{bat}} - P_{\text{bat,loss}} = (U_{\text{ocv}} \cdot I_{\text{bat}} - R_i \cdot I_{\text{bat}}^2) \cdot n_{\text{packs}} \cdot n_{\text{bat,cells}}$$

Analogous to the FCS, the power losses during DC-DC conversion are considered for the battery system. Given the battery's ability to manage bidirectional power flow, the DC-DC converter is designed to accommodate this. The polynomial is therefore utilized for both directions, which introduces a non-smooth transition at zero power flow. To manage this transition, the two polynomials are smoothly connected using a sigmoid function:

$$P_{\text{bat,bus}} = P_{\text{bat,max}} \cdot f_{\text{sigmoid}} \left(\sum_{i=1}^2 a_i \cdot \frac{P_{\text{bat,net}}^i}{P_{\text{bat,max}}} \right), \quad (4.26)$$

with $P_{\text{bat,bus}}$ being the battery system power available at the output of the DC-DC converter, and $P_{\text{bat,max}}$ is the expected maximum battery system power.

Finally, (4.27) simply sums the powers of the FCS and the battery system at the DC bus to yield the total delivered power P_{drive} .

$$P_{\text{drive}} = P_{\text{fcs,bus}} + P_{\text{bat,bus}} \quad (4.27)$$

4.3 High-level control

The high-level control consists of a power-split NMPC and a target selector. The power-split NMPC determines the optimal power distribution between the FCS and the battery system, while the target selector determines the optimal states and inputs of the FCS given the FCS power optimized by the power-split NMPC. Together, these two components determine the references for the low-level control which comprises a tracking MPC. We will start by introducing the target selector as it provides a closed-form solution for the FCS efficiency utilized in the power-split NMPC, introduced shortly after.

4.3.1 Target selector

The target selector calculates the optimal operating point, i.e., the optimal steady-state states x_s and inputs u_s , for the FCS with respect to an economic cost function, while ensuring compliance with a desired FCS power reference $P_{\text{fcs,bus,ref}}$ and system constraints.

The relevant system equations are the ODEs describing the FCS given in (4.14), (4.6), (4.18), (4.15), (4.5), and (4.1). For clarity and readability, the equations are listed below with their left-hand sides set to zero.

$$0 = \frac{R_a \cdot T_{\text{st}}}{V_{\text{ca}}} \cdot (\dot{m}_{\text{ca}}^{\text{in}} - \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{ca}}^{\text{out}}) \quad (4.28)$$

$$0 = \frac{\gamma \cdot R_a}{V_{\text{im}}} \cdot (\dot{m}_{\text{cp}}^{\text{out}} \cdot T_{\text{cp}}^{\text{out}} - \dot{m}_{\text{im}}^{\text{out}} \cdot T_{\text{im}}) \quad (4.29)$$

$$0 = \frac{R_a \cdot T_{om}}{V_{om}} \cdot (\dot{m}_{ca}^{out} - \dot{m}_{om}^{out}) \quad (4.30)$$

$$0 = \frac{R \cdot T_{st}}{M_{O_2} \cdot V_{ca}} \cdot (\dot{m}_{O_2}^{in} - \dot{m}_{O_2}^{react} - \dot{m}_{O_2}^{out}) \quad (4.31)$$

$$0 = \dot{m}_{cp}^{out} - \dot{m}_{im}^{out} \quad (4.32)$$

$$0 = \frac{1}{J_{cp}} \cdot (\tau_{cm} - \tau_{cp}) \quad (4.33)$$

For brevity, we will refer to the steady-state FCS dynamics as

$$0 = f(x_s, u_s, p) \quad (4.34)$$

with the following states and inputs:

$$x = [p_{ca} \ p_{im} \ p_{om} \ p_{O_2} \ m_{im} \ \omega_{cp}]^\top \in \mathbb{R}^6 \quad (4.35)$$

$$u = [I_{st} \ U_{cm} \ h_{om}]^\top \in \mathbb{R}^3 \quad (4.36)$$

The relevant parameter for the FCS model is the stack temperature T_{st} . The optimization problem solved within the target selector aims to minimize the lower heating value power P_{lhv} :

$$J_{economic} = P_{lhv} = \frac{\Delta \bar{h}_f^0 \cdot \dot{m}_{react}}{M_{H_2}} = \frac{\Delta \bar{h}_f^0 \cdot n_{fc}}{2 \cdot F} \cdot I_{st}$$

Given a desired FCS power $P_{fcs,bus}$, minimizing P_{lhv} is equivalent to maximizing the FCS efficiency:

$$\eta_{fcs} = \frac{P_{fcs,bus}}{P_{lhv}}$$

This is done considering the following inequality constraints $h(x_s, u_s, p, \varepsilon_s)$. Firstly, the steady-state dynamics in (4.34) are implemented as soft constraints due to potential infeasibility issues:

$$-\varepsilon_{s,sys} \leq f(x_s, u_s, T_{st}) \leq \varepsilon_{s,sys}$$

Another equality constraint, implemented with slack variables, ensures that the FCS output power $P_{fcs,bus}$ matches the FCS power optimized by the power-split NMPC, referred to as $P_{fcs,bus,ref}$:

$$-\varepsilon_{s,power} \leq P_{fcs,bus} - P_{fcs,bus,ref} \leq \varepsilon_{s,power}$$

Furthermore, the steady-state target problem includes the safety-critical constraints of the air path, which comprise the surge and choke compressor boundaries as well as a lower bound on the OER:

$$h_{choke} = 13.125 \cdot \dot{m}_{cr} + 0.82 - \Pi \leq p_{bo,choke} \quad (4.37)$$

$$h_{surge} = \Pi - 278.6852 \cdot \dot{m}_{cr}^2 - 17.8817 \cdot \dot{m}_{cr} - 0.7888 \leq p_{bo,surge} \quad (4.38)$$

$$h_{\text{oer}} = \lambda_{\text{O}_2, \text{min}} \cdot \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2, \text{im}}^{\text{out}} \leq -\dot{m}_{\text{O}_2}^{\text{react}} \cdot p_{\text{bo, oer}} \quad (4.39)$$

We implement these constraints with the option to tighten them using dedicated back-off parameters $p_{\text{bo,choke}}$, $p_{\text{bo,surge}}$, and $p_{\text{bo,oer}}$. When the back-off parameters are set to zero, we refer to the target selector as the nominal target selector. The steady-state target problem also includes box constraints on the states and inputs:

$$x_{\text{min}} \leq x_s \leq x_{\text{max}} \quad (4.40)$$

$$u_{\text{min}} \leq u_s \leq u_{\text{max}} \quad (4.41)$$

Lastly, all slack variables are demanded to be greater or equal to zero:

$$\varepsilon_s \geq 0$$

They are penalized with a quadratic and linear penalty norm in the objective function to allow for exact penalties. In summary, the steady-state target problem is formulated as follows:

$$\begin{aligned} & \underset{x_s, u_s, \varepsilon_s}{\text{minimize}} && P_{\text{lhv}} + \|\varepsilon_s\|_{S_q} + |\varepsilon_s|_{S_1} \\ & \text{subject to} && h(x_s, u_s, p, \varepsilon_s) \leq 0. \end{aligned} \quad (4.42)$$

The steady-state target problem is solved at the same frequency as the sampling time of the power-split NMPC, i.e. every 1 s. This frequency represents the maximum possible rate of change for the power reference $P_{\text{fcs, bus, ref}}$. Considering the entire optimization problem, the parameters p comprise:

$$p = [T_{\text{st}} \quad p_{\text{bo}}^{\top}]^{\top} = [T_{\text{st}} \quad p_{\text{bo,choke}} \quad p_{\text{bo,surge}} \quad p_{\text{bo,oer}}]^{\top} \in \mathbb{R}^4$$

With 6 states and 3 inputs, the optimization problem has 9 degrees of freedom. Given 7 (relaxed) equality constraints from the steady-state system dynamics and the required FCS power delivery, 2 degrees of freedom remain. We choose these remaining degrees of freedom to be the operating point in the compressor map, specifically the corrected compressor mass flow (\dot{m}_{cr}) and the compressor pressure ratio (Π). Therefore, the solution of the steady-state target problem is used to calculate the optimal compressor operating point.

Given a desired FCS reference and a specific compressor operating point, the steady-state FCS model yields a unique solution. Consequently, the main purpose of the target selector is to determine an efficient operating point in the compressor map that meets the given power reference. This approach helps to address the issue of overactuation of the FCS in the low-level control.

The input parameters to the steady-state target problem in (4.42) are the parameters p and the FCS power reference $P_{\text{fcs, bus, ref}}$. The solutions to the optimization problem can be generated offline for a grid of these input parameters and saved in a map. We do this for the nominal target selector, without the back-off parameters, for the dimension of the map to remain

manageable. In this case, the solutions depend only on the stack temperature T_{st} and the FCS power reference $P_{fcs,bus,ref}$. These precomputed solutions allow us to obtain a closed-form expression for the FCS efficiency, which depends solely on the two mentioned variables. Selected FCS efficiency curves are plotted in Fig. 4.7. These are generated by solving the steady-state target problem on a grid with FCS output power ranging from 0kW to 100kW in increments of 2kW, and stack temperature ranging from 2°C to 100°C in increments of 5°C. Note that the efficiency curves stop before the maximum FCS power output specified for the grid (100kW), and the maximum FCS power output they reach depends on the stack temperature. Since we observed that the resulting efficiency curves reach a minimum before 100kW and then increase again, we truncate the efficiency curves at this minimum point, assuming that the results beyond the minimum point are not reasonable. This results in a different maximum admissible FCS power output for every stack temperature.

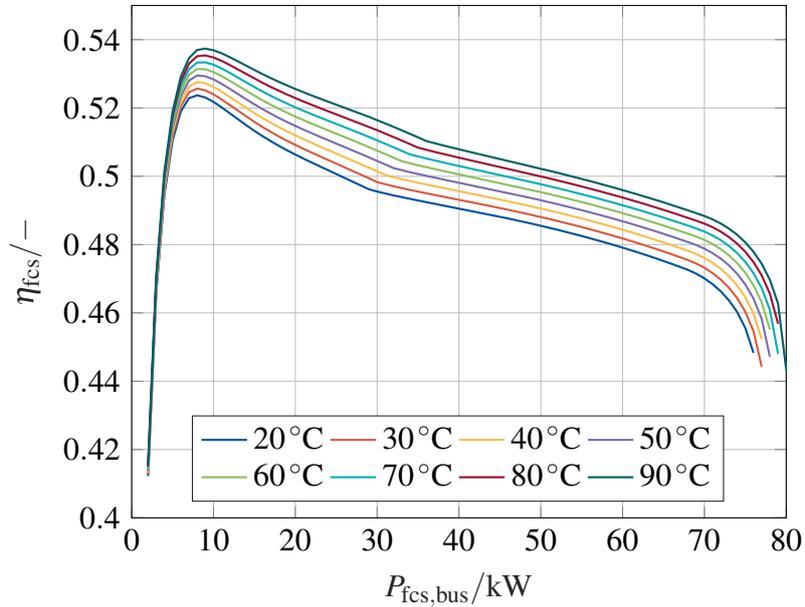


Fig. 4.7: Resulting FCS efficiency curves η_{fcs} from the nominal steady-state target problem, plotted over corresponding FCS output powers $P_{fcs,bus}$ for different stack temperatures T_{st} .

The closed-form solution for the FCS efficiency as a function of FCS output power and stack temperature is particularly useful for formulating the optimization problem solved within the power-split NMPC, as explained in the upcoming section.

4.3.2 Power-split NMPC

Prediction model formulation

The controller-internal model of the power-split NMPC mainly consists of the battery model described in Section 4.2.5. The relevant equations are (4.21)–(4.22) and (4.24), which are summarized below for better clarity and readability:

$$U_{\text{ocv}} = E_0 - \frac{K}{x_{\text{SoC}}} + A \cdot \exp(-B \cdot Q \cdot (1 - x_{\text{SoC}}))$$

$$\dot{x}_{\text{SoC}} = -\frac{1}{2 \cdot Q \cdot R_i} \cdot \left(U_{\text{ocv}} - \sqrt{U_{\text{ocv}}^2 - \frac{4 \cdot P_{\text{bat,net}} \cdot R_i}{n_{\text{packs}} \cdot n_{\text{bat,cells}}}} \right) \quad (4.43)$$

To avoid complex values, the square root function in (4.43) is adapted as explained in A.3. The main purpose of this model is to provide an estimate of the change in battery SoC as a function of the battery output power, including the power losses due to power electronics. These losses are considered through the following relation between battery net power output, $P_{\text{bat,net}}$, and power output at the DC bus, $P_{\text{bat,bus}}$:

$$P_{\text{bat,net}} = \frac{P_{\text{bat,bus}}}{\eta_{\text{DCDC}}},$$

where η_{DCDC} can be derived from (4.26). Altogether, this provides us with a dynamic relationship between the battery output power we can harness, $P_{\text{bat,bus}}$, and the battery SoC, x_{SoC} . The second part of the prediction model used in the power-split NMPC considers the FCS statically via the efficiency map described in Section 4.3.1:

$$\eta_{\text{fcs}} = \frac{P_{\text{fcs,bus}}}{P_{\text{lhv}}} = \phi_{\text{map}}(P_{\text{fcs,bus}}, T_{\text{st}}),$$

which provides a static relationship between the usable output power of the FCS and its efficiency. Thus, the only system state of the high-level prediction model is the battery SoC and the inputs to the model are the FCS power, $P_{\text{fcs,bus}}$, and the battery system power, $P_{\text{bat,bus}}$. The stack temperature, T_{st} , is treated as a known run-time parameter and assumed constant over the prediction horizon. In short, the prediction model can be expressed as follows:

$$\begin{aligned} \dot{x} &= f(x, u) \\ y &= \phi(u, p) \end{aligned} \quad (4.44)$$

with

$$x = x_{\text{SoC}} \in \mathbb{R} \quad (4.45)$$

$$u = [P_{\text{fcs,bus}} \quad P_{\text{bat,bus}}]^\top \in \mathbb{R}^2 \quad (4.46)$$

$$y = \eta_{\text{fcs}} \in \mathbb{R} \quad (4.47)$$

$$p = T_{\text{st}} \in \mathbb{R} \quad (4.48)$$

To avoid high gradients in the power values $P_{\text{fcs,bus}}$ and $P_{\text{bat,bus}}$, we choose to penalize their derivatives in the cost function. One way to achieve this is by adopting the delta formulation of the prediction model in (4.44). The approach extends the state vector, denoted as x_{aug} in the augmented state space, to include the actuated variables, while the input vector u_{aug} comprises their rates of change. The implementation of this augmentation can vary. In this work, we choose to augment the system dynamics prior to discretization. The resulting prediction model is:

$$\begin{aligned} \dot{x}_{\text{aug}}(t) = f(x_{\text{aug}}(t), u_{\text{aug}}(t), p(t)) &= \begin{bmatrix} \dot{x}(t) \\ \dot{u}(t) \end{bmatrix} = \begin{bmatrix} f(x(t), u(t - \tau) + \int_{t-\tau}^t u_{\text{aug}}(t) dt, p(t)) \\ u_{\text{aug}}(t) \end{bmatrix} \\ &= \begin{bmatrix} f(x(t), u(t - \tau) + u_{\text{aug}}(t) \cdot \tau, p(t)) \\ u_{\text{aug}}(t) \end{bmatrix} \end{aligned}$$

where the augmented state $x_{\text{aug}}(t)$ combines the current state $x(t)$ with the input from the previous timestep $u(t - \tau)$, while $u_{\text{aug}}(t)$ denotes the input's rate of change $\dot{u}(t)$:

$$x_{\text{aug}}(t) = \begin{bmatrix} x(t) \\ u(t - \tau) \end{bmatrix} \quad (4.49)$$

$$u_{\text{aug}}(t) = \dot{u}(t) \quad (4.50)$$

This approach allows us to penalize and constrain the input derivatives while maintaining a sparse structure of the resulting OCP. The state space variables now comprise the following system variables:

$$x = [P_{\text{fcs,bus}} \quad P_{\text{bat,bus}} \quad x_{\text{SoC}}]^\top \in \mathbb{R}^3 \quad (4.51)$$

$$u = [\dot{P}_{\text{fcs,bus}} \quad \dot{P}_{\text{bat,bus}}]^\top \in \mathbb{R}^2 \quad (4.52)$$

Optimal control problem

The cost function of the OCP of the power-split NMPC can be stated in terms of the following three objectives:

1. maximize FCS efficiency
2. be charge sustaining, i.e. the battery SoC at the end of the cycle should correspond to its initial value $x_{\text{SoC,ref}}$
3. avoid high gradients in FCS power and battery system power

These objectives can be formalized in terms of the following stage cost:

$$\begin{aligned} \ell(x(t), u(t)) \triangleq & Q_{\eta_{\text{fcs}}} \cdot \log(0.6 - \eta_{\text{fcs}}(t)) \\ & + Q_{\text{SoC}} \cdot (x_{\text{SoC,ref}}(t) - x_{\text{SoC}}(t))^2 \\ & + \|u(t)\|_R, \end{aligned}$$

where $Q_{\eta_{\text{fcs}}} \geq 0$ and $Q_{\text{SoC}} \geq 0$ are positive semidefinite weighting coefficients while $R > 0$ is a positive weighting matrix of appropriate dimension with

$$R \triangleq \begin{bmatrix} R_{\dot{P}_{\text{fcs,bus}}} & 0 \\ 0 & R_{\dot{P}_{\text{bat,bus}}} \end{bmatrix}.$$

The economic objective (1) penalizes small FCS efficiencies relative to a threshold of 60% using a logarithmic function while the charge sustainability objective (2) refers to a tracking objective and is thus stated as a quadratic term. Objective (3) is expressed by weighting the derivatives of the corresponding input variables quadratically as well. To force the battery SoC at the end of the prediction horizon, i.e., at time t_f , to be equal to its reference, we additionally introduce the terminal tracking objective with a different weighting. Objective (1) is also included in the terminal cost, with the same weighting coefficient as in the stage cost, to ensure that all optimization variables are penalized.

$$\begin{aligned} \ell_f(x(t_f)) \triangleq & Q_{\eta_{\text{fcs}}} \cdot \log(0.6 - \eta_{\text{fcs}}(t_f)) \\ & + Q_{\text{SoC},t_f} \cdot (x_{\text{SoC,ref}}(t_f) - x_{\text{SoC}}(t_f))^2 \end{aligned}$$

with semipositive weighting coefficient $Q_{\text{SoC},t_f} \geq 0$. As the power delivered by the FCS and the battery system should satisfy the total power demand, we enforce a relaxed equality constraint as given below.

$$-\varepsilon_1 \leq P_{\text{fcs,bus}} + P_{\text{bat,bus}} - P_{\text{drive,ref}} \leq \varepsilon_1$$

Note that this constraint is enforced for the entire prediction horizon except for the first time step t_0 , as this can conflict with the initial state constraint. Furthermore, input and state box constraints are implemented, where the upper bound for the FCS power $P_{\text{fcs,bus}}$ is dependent on the stack temperature T_{st} , as mentioned in Section 4.3.1.

$$x_{\min} \leq x \leq x_{\max}$$

$$u_{\min} \leq u \leq u_{\max}$$

In addition to the hard box constraints on the battery SoC enforced in (4.3.2), we implement soft box constraints on the battery SoC with tighter bounds.

$$x_{\text{SoC,min,soft}} - \varepsilon_2 \leq x_{\text{SoC}} \leq x_{\text{SoC,max,soft}} + \varepsilon_2$$

Needless to say, all slack variables are constrained to be greater than zero.

$$\boldsymbol{\varepsilon} = [\boldsymbol{\varepsilon}_1 \quad \boldsymbol{\varepsilon}_2]^\top \in \mathbb{R}^2 \geq 0$$

All in all, this yields the following OCP

$$\begin{aligned} & \underset{x(\cdot), u(\cdot), \boldsymbol{\varepsilon}}{\text{minimize}} && \int_{t_0}^{t_f} \ell(x(t), u(t), p) dt + \ell_f(x(t_f), p) + \|\boldsymbol{\varepsilon}\|_{S_q} + |\boldsymbol{\varepsilon}|_{S_l} \\ & \text{subject to} && x(t_0) = \hat{x}_0, \\ & && \dot{x}(t) = f(x(t), u(t)), \quad \forall t \in [t_0, t_f], \\ & && h(x(t), u(t), \boldsymbol{\varepsilon}) \leq 0, \quad \forall t \in [t_0, t_f], \end{aligned} \tag{4.53}$$

where t_k refers to the current time step k and t_f to the end of the optimization horizon. In addition to the stage and terminal cost, the slack variables are penalized with a quadratic and linear penalty norm to ensure exact penalties. The constraints comprise of the initial state constraint and the system dynamics as equality constraints and the inequality constraints elaborated earlier.

Transcription and solution of the nonlinear program

In direct optimal control methods, the continuous-time OCP in (4.53) is discretized, resulting in a finite-dimensional NLP. Here, multiple shooting is applied and the NLP reads:

$$\begin{aligned} & \underset{x_{0:N}, u_{0:N-1}, \boldsymbol{\varepsilon}}{\text{minimize}} && \sum_{k=0}^{N-1} \ell(x_k, u_k, p) + \ell_f(x_N, p) + \|\boldsymbol{\varepsilon}\|_{S_q} + |\boldsymbol{\varepsilon}|_{S_l} \\ & \text{subject to} && x_0 = \hat{x}_0, \\ & && x_{k+1} = f_d(x_k, u_k), \quad k = 0, \dots, N-1, \\ & && h(x_k, u_k, \boldsymbol{\varepsilon}) \leq 0, \quad k = 0, \dots, N-1, \end{aligned}$$

where control and prediction horizon are both set to $N = 15$, with the sampling time T_s aligned to the interval of power demand changes, which is 1 s. The system dynamics are discretized using a fourth-order explicit Runge-Kutta method with one integration step per shooting interval. Throughout this thesis, the transcribed optimization problem of the power-split NMPC is then interfaced with and solved by the open-source solver IPOPT.

4.4 Low-level control

4.4.1 Prediction model formulation

The controller-internal model used by the tracking MPC consists of the differential and algebraic equations that capture the dynamics of the air path, the terminal voltage, and the power generation for both the FCS and the battery system, as detailed in Section 4.2. The model comprises six states and four inputs. The run-time parameters of the model are the stack temperature, T_{st} , and the SoC of the battery, x_{SoC} .

$$x = [p_{ca} \ p_{im} \ p_{om} \ p_{O_2} \ m_{im} \ \omega_{cp}]^T \in \mathbb{R}^6 \quad (4.54)$$

$$u = [I_{st} \ U_{cm} \ h_{om} \ I_{bat}]^T \in \mathbb{R}^4 \quad (4.55)$$

$$p = [T_{st} \ x_{SoC}]^T \in \mathbb{R}^2 \quad (4.56)$$

The relevant ODEs are summarized below:

$$\frac{dp_{ca}}{dt} = \frac{R_a \cdot T_{st}}{V_{ca}} \cdot (\dot{m}_{ca}^{in} - \dot{m}_{O_2}^{react} - \dot{m}_{ca}^{out}) \quad (4.57)$$

$$\frac{dp_{im}}{dt} = \frac{\gamma \cdot R_a}{V_{im}} \cdot (\dot{m}_{cp}^{out} \cdot T_{cp}^{out} - \dot{m}_{im}^{out} \cdot T_{im}) \quad (4.58)$$

$$\frac{dp_{om}}{dt} = \frac{R_a \cdot T_{om}}{V_{om}} \cdot (\dot{m}_{ca}^{out} - \dot{m}_{om}^{out}) \quad (4.59)$$

$$\frac{dp_{O_2}}{dt} = \frac{R \cdot T_{st}}{M_{O_2} \cdot V_{ca}} \cdot (\dot{m}_{O_2}^{in} - \dot{m}_{O_2}^{react} - \dot{m}_{O_2}^{out}) \quad (4.60)$$

$$\frac{dm_{im}}{dt} = \dot{m}_{cp}^{out} - \dot{m}_{im}^{out} \quad (4.61)$$

$$\frac{d\omega_{cp}}{dt} = \frac{1}{J_{cp}} \cdot (\tau_{cm} - \tau_{cp}) \quad (4.62)$$

They cover the dynamic states of the FCS, while the battery system is considered through a static relationship between the battery current I_{bat} , the battery SoC x_{SoC} and the output battery power $P_{bat,bus}$, see Section 4.2.5 and Section 4.2.6. For output tracking, we choose the following five variables:

$$y = [P_{drive} \ P_{fcs,bus} \ P_{bat,bus} \ \dot{m}_{cr} \ \Pi]^T \in \mathbb{R}^5 \quad (4.63)$$

The relevant equations for the output variables are given in Section 4.2.6 and Section 4.2.1. Notably, without the output variable P_{drive} , which is the sum of $P_{fcs,bus}$ and $P_{bat,bus}$, the prediction model has an equal number of output variables and input variables. We include the total power P_{drive} to allow for separate weighting of this reference tracking objective from tracking of the individual powers. Dynamically, this can enhance total power tracking by, for example,

encouraging one of the individual powers to overshoot to better meet the overall power demand. For brevity, we will refer to the prediction model of the tracking MPC as follows:

$$\begin{aligned}\dot{x}(t) &= f(x(t), u(t), p) \\ y(t) &= \phi(x(t), u(t), p)\end{aligned}\tag{4.64}$$

Similar to the prediction model formulation of the power-split NMPC elaborated in Section 4.3.2, we also adopt the delta formulation of (4.64) for the design of the tracking MPC in the low-level control. This approach allows for penalizing and constraining the rates of change of the inputs and, in combination with multiple shooting, results in a sparse OCP structure of the ensuing NLP and QP subproblems. This structure is well-suited for the employment of efficient, structure-exploiting solvers. Accordingly, we form the following augmented system based on the delta formulation:

$$\begin{aligned}\dot{\tilde{x}}(t) &= f(\tilde{x}(t), \tilde{u}(t), p) \\ y(t) &= \phi(\tilde{x}(t), \tilde{u}(t), p)\end{aligned}$$

with

$$\begin{aligned}\tilde{x} &= [p_{ca} \ p_{im} \ p_{om} \ p_{O_2} \ m_{im} \ \omega_{cp} \ I_{st} \ U_{cm} \ h_{om} \ I_{bat}]^T \in \mathbb{R}^{10} \\ \tilde{u} &= [\dot{I}_{st} \ \dot{U}_{cm} \ \dot{h}_{om} \ \dot{I}_{bat}]^T \in \mathbb{R}^4\end{aligned}$$

Note that, from this point forward, the augmented state \tilde{x} and input vector \tilde{u} will be referred to simply as x and u , respectively.

4.4.2 Optimal control problem

The tracking MPC is defined by the OCP in (4.65). The objective function regulates the output variables y to their reference values $y_{ref} \in \mathbb{R}^5$ and penalizes the input changes u over the prediction time t_p . The last two terms in the objective function minimize the violation of soft constraints, where $\varepsilon \in \mathbb{R}^3$ is the vector of slack variables. These are penalized both quadratically and linearly to allow for exact penalties.

$$\begin{aligned}\text{minimize}_{x(\cdot), u(\cdot), \varepsilon} & \frac{1}{2} \cdot \int_{t_0=t_k}^{t_f} (\|y(t) - y_{ref}(t)\|_Q^2 + \|u(t)\|_R^2) dt + \|\varepsilon\|_{S_q}^2 + |\varepsilon|_{S_l} \\ \text{subject to} & \quad 0 = x(t_0) - \hat{x}_0, \\ & \quad \dot{x} = f(x(t), u(t), p), \quad \forall t \in [t_0, t_f], \\ & \quad 0 \geq h(x(t), u(t), p, \varepsilon), \quad \forall t \in [t_0, t_f]\end{aligned}\tag{4.65}$$

The optimized state and input trajectories must adhere to the initial state constraint and the prediction model dynamics. To account for actuator limitations, the inequality constraints $h(x(t), u(t), p, \varepsilon)$ impose specific input constraints as part of the OCP:

$$\begin{aligned} I_{\text{st}}^{\min} &\leq I_{\text{st}}, \\ h_{\text{om}}^{\min} &\leq h_{\text{om}} \leq h_{\text{om}}^{\max}, \end{aligned}$$

where h_{om} is the opening position of the outlet manifold valve. The lower bound on the stack current, I_{st}^{\min} , ensures operation of the FC stack away from the OCV, which is a catalyst and membrane stressor [45]. To accommodate actuator dynamics that are not explicitly modeled, we also constrain the rate of input changes, recognizing that this notably simplifies the typically more complex behavior of actuators:

$$u_{\min} \leq u \leq u_{\max}$$

The above constraints on the actuators are all enforced as hard constraints, i.e. no constraint violation is allowed. In contrast, the safety and degradation constraints considered in this work are imposed as soft constraints to ensure feasibility of the problem. Safe operation of the compressor is guaranteed within the choke and surge boundary:

$$\begin{aligned} 13.125 \cdot \dot{m}_{\text{cr}} + 0.82 - \Pi &\leq \varepsilon_{\text{choke}} \\ \Pi - 278.6852 \cdot \dot{m}_{\text{cr}}^2 - 17.8817 \cdot \dot{m}_{\text{cr}} - 0.7888 &\leq \varepsilon_{\text{surge}} \end{aligned}$$

Moreover, to ensure sufficient delivery of oxygen to the cathode channels, the OER λ_{O_2} is constrained to stay above a minimum value, $\lambda_{\text{O}_2}^{\min}$. Since $\lambda_{\text{O}_2} = f(I_{\text{st}}^{-1})$ is a nonlinear function dependent on the inverse of the stack current I_{st} (see (4.17)), we implement an equivalent formulation of this constraint that is linear in I_{st} :

$$\lambda_{\text{O}_2}^{\min} \cdot \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2, \text{im}}^{\text{out}} \leq \dot{m}_{\text{O}_2}^{\text{react}} \cdot \varepsilon_{\lambda_{\text{O}_2}}$$

We noticed, that if we choose the right-hand side to be just $\varepsilon_{\lambda_{\text{O}_2}}$, the optimizer tends to violate the constraint more than we find acceptable. Lastly, all slack variables are required to be non-negative:

$$\varepsilon = \begin{bmatrix} \varepsilon_{\lambda_{\text{O}_2}} & \varepsilon_{\text{choke}} & \varepsilon_{\text{surge}} \end{bmatrix}^{\top} \in \mathbb{R}^3 \geq 0$$

4.4.3 Transcription to nonlinear program

Using the so-called direct method, the continuous-time OCP in (4.65) is discretized into a finite-dimensional NLP. Here, multiple shooting is applied and the resulting NLP reads:

$$\begin{aligned}
& \underset{x_{0:N-1}, u_{0:N-1}, \varepsilon}{\text{minimize}} && \frac{1}{2} \cdot \sum_{k=0}^{N-1} (\|y_k - y_{\text{ref},k}\|_Q^2 + \|u_k\|_R^2) + \|\varepsilon\|_{S_q}^2 + |\varepsilon|_{S_1} \\
& \text{subject to} && 0 = x_0 - \hat{x}_0, \\
& && x_{k+1} = f_d(x_k, u_k, p), \quad k = 0, \dots, N-2, \\
& && 0 \geq h(x_k, u_k, p, \varepsilon), \quad k = 0, \dots, N-1
\end{aligned} \tag{4.66}$$

Here, the discrete-time outputs are given by

$$y_k = \phi(x_k, u_k, p), \quad k = 0, \dots, N-1.$$

Since no terminal cost or terminal constraints are considered, the state at the final grid point is not included as an optimization variable, and the system dynamics are enforced only over the first $N-1$ discretization intervals, corresponding to the optimization variables $x_{0:N-1}$.

When transcribing the OCP, several design aspects influence the numerical size of the NLP, which affects both the computational cost of solving it and the resulting control performance. In this work, we analyze the choice of prediction horizon N , integration scheme, and sampling time $T_s = (t_f - t_0)/N$.

Sampling time T_s

The sampling time determines the upper bound of achievable control bandwidth. Ideally, we seek the sampling time to be as fast as the fastest time constant of the prediction model. To analyze its time constants, we simulate the model in an open-loop manner with inputs corresponding to a closed-loop FCS power trajectory covering different operating points between 8.5 kW and 74.2 kW, see Fig. 4.8. At each operating point, the system dynamics are linearized and the eigenvalues λ_{OP} of the system matrix $A_{\text{OP}} = \partial f / \partial x|_{\text{OP}}$ are evaluated.

We assume that the reciprocal of the absolute value of the real part of each eigenvalue, specifically $1/|\text{Re}(\lambda_{\text{OP},i})|$, is a good representation of the autonomous response time of the corresponding system mode i at the respective operating point. Furthermore, we assume that the resulting eigenvalues λ_{OP} provide a good approximation of the dynamic behavior of the controller-internal model, since the FCS power simulation trajectory encompasses a wide range of different operating points.

The real parts of the eigenvalues are plotted for each system mode and operating point in Fig. 4.9. Note that all eigenvalues were real-valued except at the operating points at 41.5 s

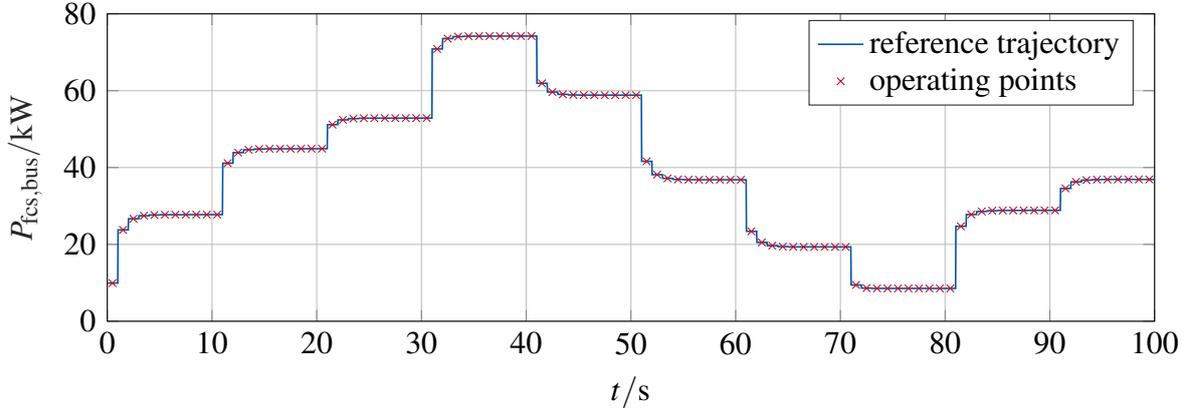


Fig. 4.8: FCS power reference trajectory and operating points, at which the system dynamics are linearized.

and 42.5 s, following a significant step down from a high power level. Upon evaluating the time constants, we find the fastest time constant to be 0.0065 s, while the system’s slowest time constant is 1.7288 s. The former provides a guide in selecting an appropriate order of magnitude for the sampling time. As the fastest mode is of order $\mathcal{O}(10^{-2})$ seconds, we choose $T_s = 10$ ms.

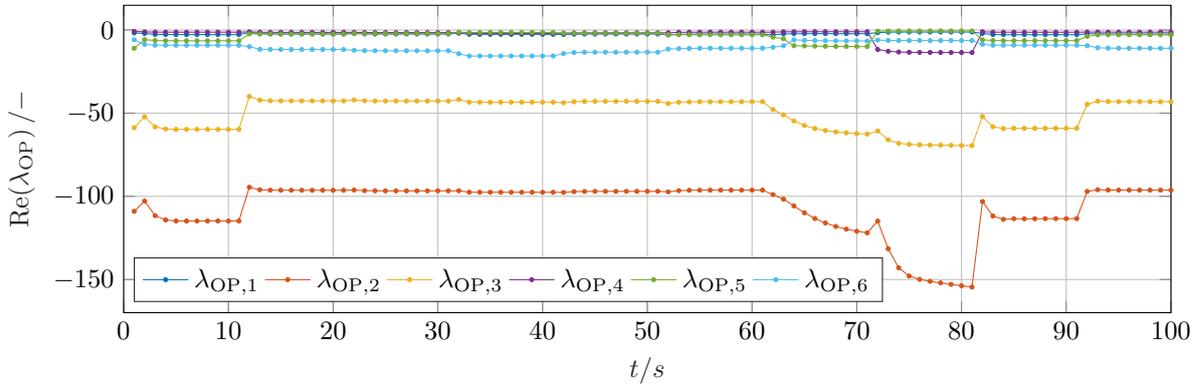


Fig. 4.9: Real parts of eigenvalues for each system mode and operating point.

Numerical integration scheme and step size

Every numerical integration scheme applied to the prediction model corresponds to a discretization error and a computational cost, both of which we aim to minimize. To select an appropriate scheme, we investigate accuracy and computational cost for ERK and IRK methods of different

stages s . For the latter, we utilize Gauss-Legendre (GL) collocation methods of order $p = 2 \cdot s$ and Radau IIa (RDIa) methods of order $p = 2 \cdot s - 1$. This approach is adopted from [98].

All mentioned integration schemes are implemented in MATLAB using `acados` [119] and each scheme is referred to with its abbreviation and the number of stages s , i.e. ERK with four stages is denoted as ERK4.

As a measure for accuracy, we define the relative discretization error e_{rel} for a simulation of N_t time steps starting from the initial state x_0 and with a constant input u as follows:

$$e_{\text{rel}} = \frac{1}{n_x} \sum_{i=1}^{n_x} \frac{1}{N_t} \sum_{k=0}^{N_t-1} \frac{|x_{i,k+1} - x_{i,k+1}^{\text{gt}}|}{|x_{i,k+1}^{\text{gt}}|},$$

with

$$\begin{aligned} x_{k+1} &= f_d(x_k, u, p), \\ x_{k+1}^{\text{gt}} &= f_d^{\text{gt}}(x_k, u, p), \end{aligned}$$

Here, $x_{i,k}$ denotes the i -th component of the discrete-time state vector with n_x states at time step k , and f_d denotes the discrete-time state transition induced by the respective integration scheme. This definition quantifies the relative deviation between state trajectories induced by different discretization schemes. The parameters p are assumed constant over the simulation horizon, and the subscript “gt” refers to the ground truth discretization, which is numerically calculated using GL3 with 400 integration steps.

We evaluate e_{rel} for each integration scheme and a range of integration steps from 1 to 10 (see Fig. 4.10a), where one integration step corresponds to a step size of 10ms. Additionally, we determine the corresponding mean CPU time out of 5000 simulation repetitions, conducted on a Lenovo T480s laptop, with an Intel Core i7-8550U at 1.8GHz with 16 GB RAM. The resulting relationship between e_{rel} and mean CPU time is shown in Fig. 4.10b.

Depending on the required accuracy, different schemes provide the most computationally efficient approximation of the state trajectories, as illustrated by the Pareto front in Fig. 4.10b. For accuracies of $e_{\text{rel}} > 10^{-7}$, the integration schemes ERK1, ERK2, ERK3 and ERK4, each with one integration step, dominate, ranking from least to most accurate. For accuracies between $e_{\text{rel}} = 10^{-7}$ and $e_{\text{rel}} = 10^{-10}$ the IRK method GL2 with one integration step and the ERK method ERK4 with two integration steps dominate. The latter is more accurate but also more computationally expensive. For higher precision ($e_{\text{rel}} < 10^{-10}$), GL3, which has the highest order tested, is the most efficient method. For our purposes, we decide that an accuracy around $e_{\text{rel}} = 10^{-7}$ is sufficient, and therefore, choose the standard ERK4 with one integration step.

With a given prediction horizon, the NLP is fully defined. For controller development purposes, the problem can be directly solved with interior-point NLP solvers, such as the robust open-source solver IPOPT. Its solution can serve as a benchmark, but is not suitable for embedded

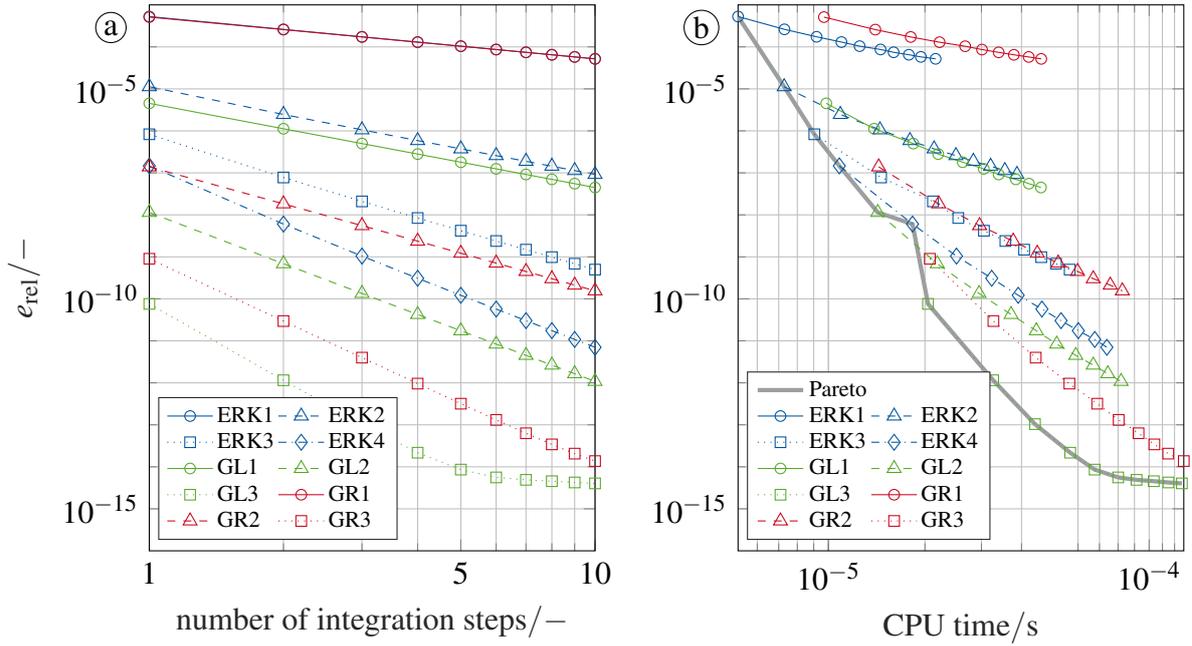


Fig. 4.10: (a) Relative discretization error vs. number of integration steps for tested integration methods, (b) relative discretization error vs. mean CPU time for tested integration schemes.

implementation. With tight limits on computation time, gradient-based optimization methods are particularly suitable as they converge within comparatively few iterations, given sufficient smoothness and continuity in problem functions. To exploit the fact, that we solve a similar problem in each sampling instant, we apply SQP due to its superior warm-starting capabilities. Unlike interior point methods, the algorithm can be initialized with the solution from the previous time step and, if close to the previous optimum, significantly reduce the required number of iterations. Moreover, there exist proven-in-use software packages for efficiently solving the quadratic subproblems emerging in SQP on resource-constrained systems.

4.4.4 Sequential quadratic programming formulation

An SQP method, when applied to the NLP described in (4.66), proceeds in each iteration i by solving the following structured QP subproblem:

$$\begin{aligned}
 & \underset{\Delta x_{0:N-1}, \Delta u_{0:N-1}}{\text{minimize}} && \sum_{k=0}^{N-1} \left(\frac{1}{2} \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix}^\top \cdot H_{k,i} \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} + q_{k,i}^\top \cdot \begin{bmatrix} \Delta x_k \\ \Delta u_k \end{bmatrix} \right) \\
 & \text{subject to} && 0 = \Delta x_0 - \Delta \hat{x}_0, \\
 & && \Delta x_{k+1} = A_{k,i} \cdot \Delta x_k + B_{k,i} \cdot \Delta u_k + b_{k,i}, \quad k = 0, \dots, N-2, \\
 & && 0 \geq E_{k,i} \cdot \Delta x_k + F_{k,i} \cdot \Delta u_k + c_{k,i}, \quad k = 0, \dots, N-1,
 \end{aligned}$$

where $H_{k,i} = \nabla_{ww}^2 \ell_{k,i}$ is the Hessian and $q_{k,i} = \nabla_w \ell_{k,i}$ is the gradient of the stage cost $\ell_{k,i} = \ell(w_{k,i})$ with respect to the current iterate $w_{k,i} = [x_{k,i} \ u_{k,i}]^\top$, i.e., the state and control input.

Given the following stage cost $\ell_{k,i}$:

$$\begin{aligned} \ell_{k,i} &= \frac{1}{2} \cdot \|y_{k,i} - y_{\text{ref},k,i}\|_Q^2 + \frac{1}{2} \cdot \|u_{k,i}\|_R^2 \\ &= \begin{pmatrix} \sqrt{Q} \cdot (y_{k,i} - y_{\text{ref},k,i}) \\ \sqrt{R} \cdot u_{k,i} \end{pmatrix}^\top \cdot \begin{pmatrix} \sqrt{Q} \cdot (y_{k,i} - y_{\text{ref},k,i}) \\ \sqrt{R} \cdot u_{k,i} \end{pmatrix} \\ &= r(x_{k,i}, u_{k,i})^\top \cdot r(x_{k,i}, u_{k,i}), \end{aligned}$$

a more specific expression for $H_{k,i}$ and $q_{k,i}$ can be derived:

$$H_{k,i} = \begin{bmatrix} H_{k,i}^{xx} & H_{k,i}^{xu} \\ H_{k,i}^{ux} & H_{k,i}^{uu} \end{bmatrix}, \quad q_{k,i} = \begin{bmatrix} q_{k,i}^x \\ q_{k,i}^u \end{bmatrix}$$

with

$$\begin{aligned} H^{xx} &= \nabla_{xx}^2 \ell_k(x_{k,i}, u_{k,i}) \\ H^{xu} &= H^{ux} = \nabla_{xu}^2 \ell_k(x_{k,i}, u_{k,i}) \\ H^{uu} &= \nabla_{uu}^2 \ell_k(x_{k,i}, u_{k,i}) \\ q_{k,i}^x &= \nabla_x \ell_k(x_{k,i}, u_{k,i}) = C_{k,i}^\top \cdot Q \cdot (y_{k,i} - y_{\text{ref},k,i}) \\ q_{k,i}^u &= \nabla_u \ell_k(x_{k,i}, u_{k,i}) = D_{k,i}^\top \cdot Q \cdot (y_{k,i} - y_{\text{ref},k,i}) + R \cdot u_{k,i} \end{aligned}$$

The Jacobian matrices $C_{k,i}$ and $D_{k,i}$ are obtained from linearization of the output equation

$$y_k = \phi(x_k, u_k, p),$$

around the current iterate $(x_{k,i}, u_{k,i})$, yielding

$$y_k \approx y_{k,i} + C_{k,i} \cdot (x_k - x_{k,i}) + D_{k,i} \cdot (u_k - u_{k,i})$$

with

$$C_{k,i} = \frac{\partial \phi}{\partial x}(x_{k,i}, u_{k,i}, p), \quad D_{k,i} = \frac{\partial \phi}{\partial u}(x_{k,i}, u_{k,i}, p)$$

The constraint Jacobian matrices and constraint violation vectors, used to update the problem at the current iterate $w_{k,i}$, are defined as follows:

$$\begin{aligned} A_{k,i} &= \nabla_x f_d(x_{k,i}, u_{k,i}, p) \\ B_{k,i} &= \nabla_u f_d(x_{k,i}, u_{k,i}, p) \end{aligned}$$

$$\begin{aligned}
 E_{k,i} &= \nabla_x h(x_{k,i}, u_{k,i}, p) \\
 F_{k,i} &= \nabla_u h(x_{k,i}, u_{k,i}, p) \\
 b_{k,i} &= f_d(x_{k,i}, u_{k,i}, p) - x_{k+1,i} \\
 c_{k,i} &= h(x_{k,i}, u_{k,i}, p)
 \end{aligned}$$

At its core, the SQP algorithm successively linearizes the NLP around the current iterate $w_{k,i}$, solves the resulting QP subproblem and updates the current iterate:

$$w_{k,i+1} = \begin{bmatrix} x_{k,i+1} \\ u_{k,i+1} \end{bmatrix} = \begin{bmatrix} x_{k,i} + \alpha \cdot \Delta x_k \\ u_{k,i} + \alpha \cdot \Delta u_k \end{bmatrix},$$

where α is the step size. Note that, two simplifications are made here: (1) slack variables are not implemented and (2) the Hessian is based on the stage cost $\ell_{k,i}$ instead of the stage Lagrangian. In addition, we apply the SQP method that employs the GN Hessian, i.e.:

$$H_{k,i} \approx J(w_{k,i})^\top \cdot J(w_{k,i})$$

with

$$\begin{aligned}
 J(w_{k,i}) &= \nabla_w r(w_{k,i}) \\
 &= \frac{\partial r(x_{k,i}, u_{k,i})}{\partial (x_{k,i}, u_{k,i})^\top},
 \end{aligned}$$

where $r(x_{k,i}, u_{k,i})$ is the residual vector. The selected Hessian approximation reduces computational costs as it only requires the evaluation of first-order derivatives. Moreover, the product of a matrix and its transpose is always at least positive semidefinite. In the case of a positive semidefinite matrix, a single operation suffices to make the matrix positive definite, i.e., the addition of a small regularization factor on all diagonal elements. These advantages are purchased at the cost of linear convergence instead of quadratic convergence in the case of an exact Hessian matrix. However, if the target can be reached, the omitted terms of the exact Hessian become small in the vicinity of the solution, mitigating the decreased convergence rate. In both cases, we perform one SQP iteration per sampling instant and the solution is shifted to obtain an initial guess for the next instant. The QPs are solved using the structure-exploiting QP solver from HPIPM [37], which exploits the sparsity of the QP subproblems induced by the multiple shooting approach.

For the following analysis on the choice of prediction horizon, a maximum number of 20 QP iterations is allowed where the step size α is set to 1. Ideally, we take a full step $\alpha = 1$ to quickly progress in the computed search direction. However, this can lead to convergence problems in specific cases, necessitating a reduction in step size. The precise choice of step size and number of iterations cannot be predicted in advance and is determined through simulation studies. Likewise, the number of SQP iterations is determined in simulation. Generally, the

real-time iteration scheme allows for a single iteration per time step [32]. With a sufficiently high control rate, the physical system changes only slightly from one time step to the next, so that the new optimum remains close to the previous solution. Under these conditions, a single iteration step can suffice to meet the convergence criteria. In the case of sufficient computation time, the number of steps can also be increased. The algorithm is implemented in MATLAB using the symbolic framework CasADi [5].

Prediction horizon N

Ideally, the choice of the prediction time $t_f = T_s \cdot N$, corresponding to a prediction horizon of N steps, ensures that the MPC can sufficiently predict how the manipulated variables may affect the cost or outputs of interest. In this work, we do not consider future references to be available in the prediction horizon. For constant references, the slowest time constant of the prediction model could serve as a benchmark for selecting the prediction horizon. However, we consider dynamically changing references with a characteristic time scale of approximately one second. This motivates a shorter prediction horizon than what would be suggested by the slowest system dynamics alone, since future reference values are not assumed to be available. Instead, the reference known at the current time step is held constant over the prediction horizon. Otherwise, the optimization would implicitly rely on an invalid assumption about the persistence of the current reference far into the future. At the same time, large horizons lead to more decision variables and constraints, resulting in a larger optimization problem with increased execution time and memory footprint. To quantify this trade-off, we evaluate the closed-loop performance and computational effort for varying prediction horizons.

Fig. 4.11a (blue) shows the evaluation of the average computation time for one control update in a closed-loop simulation conducted with a sampling time of $T_s = 10$ ms and varying prediction horizons. Additionally, we examine the relative cost difference ΔJ_{rel} for each N as defined in (4.67), where $J(N) = \sum_{i=1}^{n_c} J_i(N)/N$ denotes the average stage cost over one prediction horizon, accumulated over all n_c control updates. Among the prediction horizons tested, $J(N = 175)$ has the smallest value and is therefore used as the reference J_{ref} .

$$\Delta J_{\text{rel}}(N) = \frac{J(N) - J_{\text{ref}}}{J_{\text{ref}}} \quad (4.67)$$

Fig. 4.11a (red) shows that the smaller N , the larger ΔJ_{rel} . Overall, the CPU time increases with larger N with a linear trend, while the cost improvement decreases with larger N with an exponential trend. Plotting these two measures against each other, we observe a trade-off curve, depicted in Fig. 4.11b, guiding the control designer to balance between computation time and cost. Here, we choose $N = 35$.

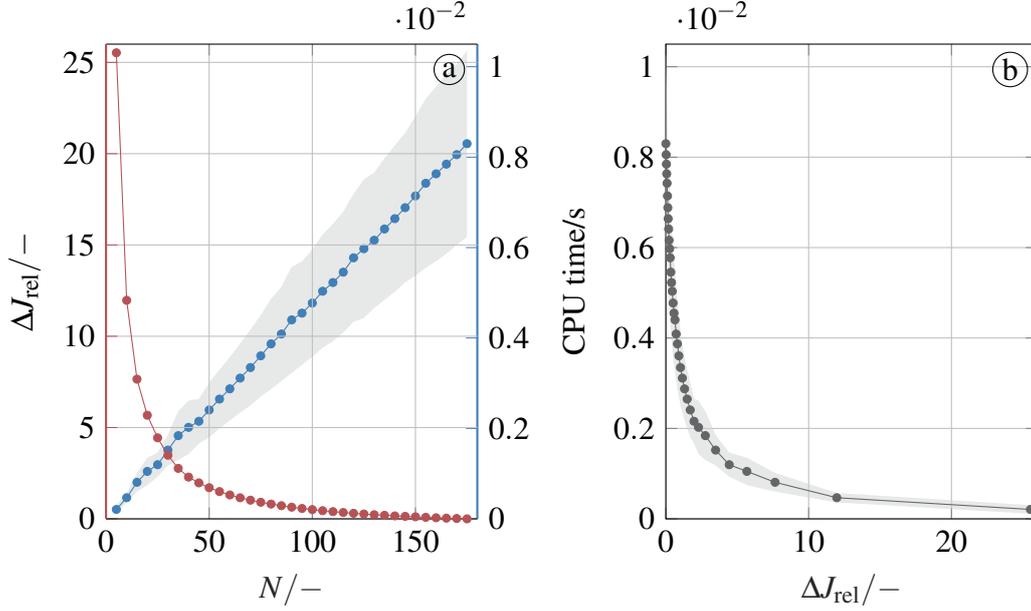


Fig. 4.11: (a) Relative cost difference and CPU time vs. prediction horizon, (b) CPU time vs. relative cost difference.

4.5 State and disturbance estimation

Due to the enclosed design of the FC stack, direct measurement of all relevant states of the process model specified in (4.54) is not feasible. Consequently, state estimation is necessary for implementing state-feedback controllers, such as MPC. Specifically, full state information is required by the tracking MPC at every time step to define the initial state in (4.65). To address this, we employ the EKF to handle the nonlinear system dynamics defined in (4.57)–(4.62). The output measurement vector y_m includes measurable state variables and the stack voltage U_{st} in (4.9):

$$y_m = [p_{\text{im}} \quad p_{\text{om}} \quad \omega_{\text{cp}} \quad U_{\text{st}}] \in \mathbb{R}^4 \quad (4.68)$$

Including the stack voltage is crucial for ensuring observability of the state vector [83]. For the implementation of the EKF, we employ an adapted version of the predictor-corrector algorithm outlined in [103], in which the corrector step precedes the predictor step to account for direct feedthrough in the measurement equation:

$$\begin{aligned} K_{k-1} &= P_{k-1}^- \cdot C_{m,k-1}^\top \cdot \left(C_{m,k-1} \cdot P_{k-1}^- \cdot C_{m,k-1}^\top + R_e \right)^{-1} \\ \hat{x}_{k-1} &= \hat{x}_{k-1}^- + K_{k-1} \cdot \left(y_{m,k-1} - \hat{y}_{m,k-1}^- \right) \\ P_{k-1} &= \left(I - K_{k-1} \cdot C_{m,k-1} \right) \cdot P_{k-1}^- \\ \hat{x}_k^- &= f_d(\hat{x}_{k-1}, u_{k-1}, p_{k-1}) \end{aligned}$$

$$\begin{aligned}\hat{y}_{m,k}^- &= \phi_m(\hat{x}_k^-, u_k, p_k) \\ P_k^- &= A_{k-1} \cdot P_{k-1} \cdot A_{k-1}^\top + Q_e\end{aligned}$$

with

$$\begin{aligned}C_{m,k-1} &= \frac{\partial \phi_m}{\partial x}(\hat{x}_{k-1}^-, u_{k-1}, p_{k-1}) \\ A_{k-1} &= \frac{\partial f_d}{\partial x}(\hat{x}_{k-1}^-, u_{k-1}, p_{k-1})\end{aligned}$$

where $\hat{y}_{m,k}^-$ is evaluated using the optimized input u_k . Here, the noise influence matrices are implicitly assumed to be identity matrices, such that the effect of process and measurement noise is fully represented by Q_e and R_e . The parameter vector p_k denotes run-time parameters that may vary between sampling instants, but are assumed constant within one sampling interval and over the prediction horizon of the MPC. Since we apply the discrete-time EKF, the nonlinear system dynamics are discretized using the fourth-order Runge-Kutta method with one integration step. The sensitivities are also calculated from the discrete system dynamics via algorithmic differentiation. An alternative approach commonly taken is to calculate the sensitivities using the continuous system dynamics and then discretize the corresponding continuous-time Jacobian matrices.

Assuming a system with feedthrough, i.e., $D_m \neq 0$, the measurement output $y_{m,k}$ depends on the input u_k and is therefore only available after the optimized input has acted on the system, becoming measurable at time step $k+1$. However, the tracking MPC requires an estimate of the current state to compute u_k . Consequently, when updating the control input at time step k , only the previous measurement $y_{m,k-1}$ is available. We therefore first perform the correction step using $y_{m,k-1}$ to obtain \hat{x}_{k-1} , followed by the prediction step to compute the a-priori estimate \hat{x}_k^- , which is then used by the MPC. Fig. 4.12b illustrates the chronological sequence of computation steps, which also applies to the covariance matrix P_k . This adopted approach is inspired by [76].

In a second step, we augment the dynamic model used within the EKF with constant output disturbances d :

$$\begin{aligned}x_{\text{aug},k+1} &= \begin{bmatrix} x_{k+1} \\ d_{k+1} \end{bmatrix} = \begin{bmatrix} f_d(x_k, u_k, p_k) \\ d_k \end{bmatrix} \\ y_{m,\text{aug},k} &= \phi_{m,\text{aug}}(x_k, u_k, p_k) + d_k\end{aligned}$$

In addition to estimating the system's current state x_k , this EKF also estimates the output disturbance d_k . The purpose of this approach is to additionally estimate the static model mismatch d_k and enable offset-free output tracking of the low-level control. Since the output variables we desire to track in the low-level MPC, defined in (4.63), are not the same as the measurement output y_m defined in (4.68), we further augment the measurement output vector to:

$$y_{m,\text{aug},k} = [p_{\text{im}} \quad p_{\text{om}} \quad \omega_{\text{cp}} \quad U_{\text{st}} \quad P_{\text{fcs,bus}} \quad P_{\text{bat,bus}} \quad \dot{m}_{\text{cp}}^{\text{out}}]^\top \in \mathbb{R}^7$$

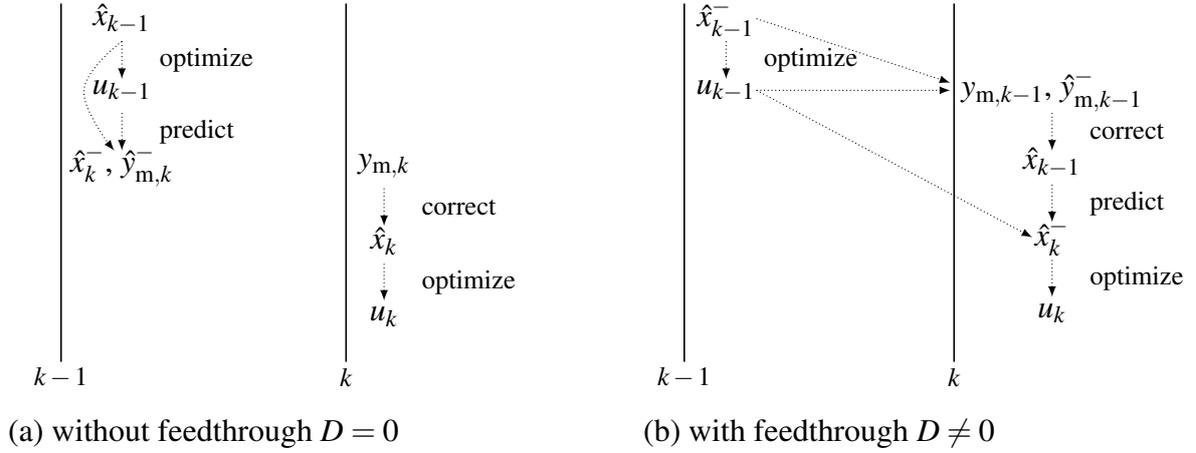


Fig. 4.12: Illustration of chronological sequence of computation steps for implementation of an EKF for a system without feedthrough (a) and with feedthrough (b).

We add the FCS power $P_{\text{fcs,bus}}$ and the battery system power $P_{\text{bat,bus}}$, whose collective static model mismatch corresponds to the static model mismatch of P_{drive} . We also add $\dot{m}_{\text{cp}}^{\text{out}}$, which is directly related to \dot{m}_{cr} . The compressor ratio Π is already covered as it directly relates to p_{im} . Consequently, $d_k \in \mathbb{R}^7$. Note that in order to achieve no offset, we augment the model with as many disturbances (and integrators) as we have measurements, $n_d = n_m$. This ensures that each measured output can be independently corrected for steady-state model mismatch. For that reason, the design procedure here requires the addition of 7 integrators even if we wish to control only a subset of measured variables. Although this is not necessary, the design procedure for the case where $n_d < n_m$ is more involved. [16] Finally, the estimated disturbances in \hat{d} are mapped to the disturbance vector required by the tracking MPC d_{MPC} based on the structure of the prediction model and the definitions of the controlled outputs:

$$\hat{d}_{\text{MPC}} = \begin{bmatrix} \hat{d}_5 + \hat{d}_6 \\ \hat{d}_5 \\ \hat{d}_6 \\ \hat{d}_7 \cdot \theta / \delta \\ \hat{d}_1 / p_a \end{bmatrix} \in \mathbb{R}^5$$

The control structure, including the state and disturbance estimator, is depicted in Fig. 4.13. We opted not to use the control structure depicted in Fig. 3.3 for offset-free steady-state tracking. Instead, we subtract the estimated disturbance \hat{d}_{MPC} from the reference for the tracking MPC. This method adapts the references but does not account for disturbances in the constraints on tracking variables. Despite this drawback, we chose this approach because the target selector is optimized every second, while the low-level control layer updates every 10 ms. Following the procedure in Section 3.3.4 is not feasible under these conditions as the estimated disturbances evolve faster than the target selector can compensate for.

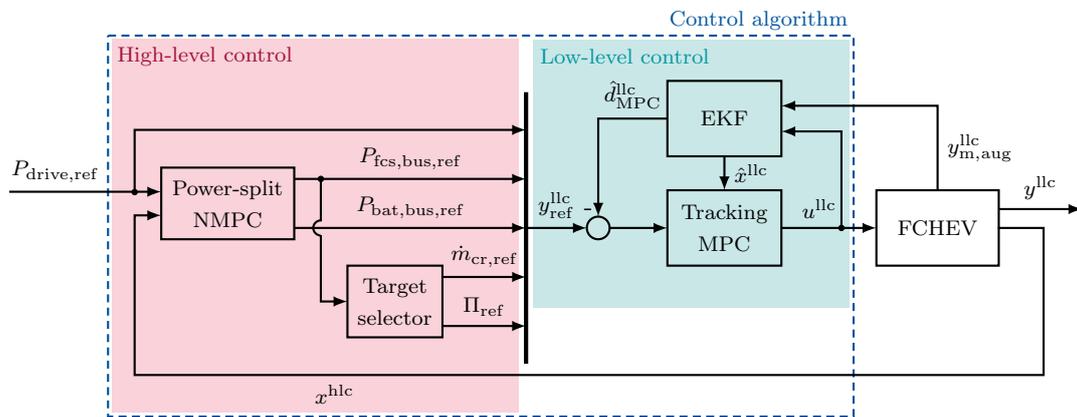


Fig. 4.13: Simplified schematic of the control structure with state and disturbance estimation.

5 Expansion to include humidity

This chapter builds upon the MPC framework introduced in Chapter 4 by incorporating humidity effects into the air-path dynamics. To this end, a water vapor balance is introduced on the cathode side, enabling a simplified representation of the membrane water content and extending the state-space model by one additional state. The resulting modifications to the control-oriented model, relative to those presented in Section 4.2, are detailed in Section 5.1. Subsequently, Sections 5.2 and 5.3 describe the corresponding adaptations of the high-level and low-level control formulations, respectively.

5.1 Control-oriented modeling

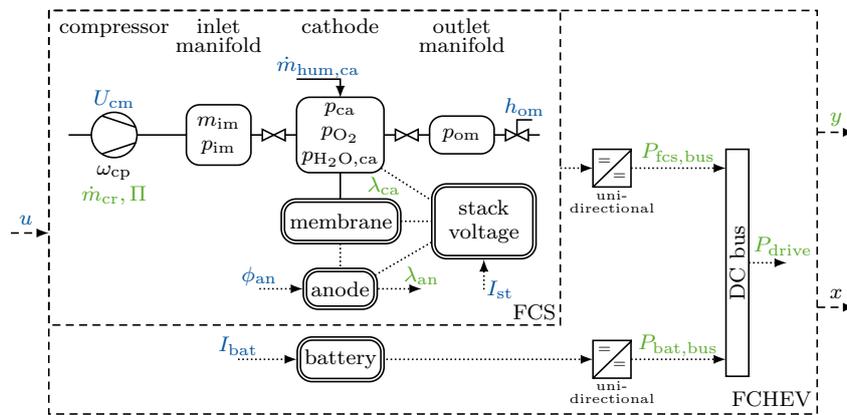


Fig. 5.1: Schematic overview of the model structure used for control-oriented modeling the FCHEV, focusing on air-path dynamics with humidity considerations.

Just like the model described in Section 4.2, the control-oriented model of the FCHEV considered here consists of a dynamic FCS model and a static battery model, as illustrated in Fig. 5.1. The battery model and the associated DCDC converter loss models remain unchanged from those used for the low-level tracking controller in Chapter 4 (see Sections 4.2.5 and 4.2.6). Accordingly, this section focuses on the FCS model, which comprises the air-path dynamics as well as static relationships for stack voltage, membrane humidity, and the anode. It is based on the following assumptions:

1. All gases are treated as ideal gases.
2. The temperature of the FC stack is perfectly controlled by the cooling system, ensuring a constant and uniform stack temperature of $T_{st} = 80^\circ\text{C}$.
3. The properties of the flow exiting the cathode (temperature, pressure, and humidity) are assumed to be identical to those inside the cathode and to dominate the electrochemical reactions at the CLs.
4. Water inside the cathode exists solely in vapor form. Once the relative humidity exceeds 100%, it instantly condenses to liquid form and is removed. Thus, flooding is not modeled.
5. Spatial variations are neglected.

Since the FCS model presented in Section 4.2, comprising the air-path dynamics and the stack voltage model, is extended here to account for humidity effects, the following focuses exclusively on the submodels that are added or modified with respect to Section 4.2.

5.1.1 Cathode flow model

For the cathode airflow dynamics, we apply mass continuity and the ideal gas law to oxygen, nitrogen and water vapor inside the cathode volume V_{ca} . This yields the following differential equations for the partial pressure of oxygen, p_{O_2} , nitrogen, p_{N_2} , and vapor, $p_{\text{H}_2\text{O},ca}$:

$$\frac{dp_{\text{O}_2}}{dt} = \frac{R \cdot T_{st}}{M_{\text{O}_2} \cdot V_{ca}} \cdot (\dot{m}_{\text{O}_2,im}^{\text{out}} - \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2,ca}^{\text{out}}), \quad (5.1)$$

$$\frac{dp_{\text{N}_2}}{dt} = \frac{R \cdot T_{st}}{M_{\text{N}_2} \cdot V_{ca}} \cdot (\dot{m}_{\text{N}_2,im}^{\text{out}} - \dot{m}_{\text{N}_2,ca}^{\text{out}}),$$

$$\frac{dp_{\text{H}_2\text{O},ca}}{dt} = \frac{R \cdot T_{st}}{M_{\text{H}_2\text{O}} \cdot V_{ca}} \left(\dot{m}_{\text{H}_2\text{O},im}^{\text{out}} + \dot{m}_{\text{H}_2\text{O}}^{\text{gen}} + \dot{m}_{\text{H}_2\text{O},mem} + \dot{m}_{\text{H}_2\text{O}}^{\text{hum},ca} - \dot{m}_{\text{H}_2\text{O},ca}^{\text{out}} - \dot{m}_{\text{H}_2\text{O}}^{\text{cond}} \right), \quad (5.2)$$

where R is the universal gas constant and M_i the molar mass of the respective species $i \in \{\text{O}_2, \text{N}_2, \text{H}_2\text{O}\}$. The partial pressure of water vapor in the cathode, $p_{\text{H}_2\text{O},ca}$, constitutes an additional dynamic state introduced by the humidity extension of the air-path model. The inlet and outlet mass flow rates, $\dot{m}_{i,im}^{\text{out}}$ and $\dot{m}_{i,ca}^{\text{out}}$ in (5.1) to (5.2), are calculated using thermodynamic properties. The former are:

$$\begin{aligned} \dot{m}_{\text{O}_2,im}^{\text{out}} &= \frac{x_{\text{O}_2,atm}}{1 + w_{im}} \cdot \dot{m}_{im}^{\text{out}}, \\ \dot{m}_{\text{N}_2,im}^{\text{out}} &= \frac{1 - x_{\text{O}_2,atm}}{1 + w_{im}} \cdot \dot{m}_{im}^{\text{out}}, \\ \dot{m}_{\text{H}_2\text{O},im}^{\text{out}} &= \frac{w_{im}}{1 + w_{im}} \cdot \dot{m}_{im}^{\text{out}}, \end{aligned}$$

where $\dot{m}_{\text{im}}^{\text{out}}$ is the inlet manifold exit flow, $x_{\text{O}_2,\text{atm}}$ the molar fraction of oxygen in ambient air, and w_{im} the humidity ratio in the inlet manifold. The latter is calculated as follows:

$$w_{\text{im}} = \frac{M_{\text{H}_2\text{O}}}{x_{\text{O}_2,\text{atm}} \cdot M_{\text{O}_2} + (1 - x_{\text{O}_2,\text{atm}}) \cdot M_{\text{N}_2}} \cdot \frac{\text{RH}_{\text{atm}} \cdot p_{\text{sat}}}{p_{\text{atm}} - \text{RH}_{\text{atm}} \cdot p_{\text{sat}}},$$

where p_{sat} is the saturation pressure at T_{st} , and p_{atm} and RH_{atm} are the ambient pressure and relative humidity, respectively. A linearized nozzle equation is used to calculate the cathode exit flow rate, $\dot{m}_{\text{ca,out}}$:

$$\dot{m}_{\text{ca,out}} = k_{\text{ca,out}} \cdot (p_{\text{ca}} - p_{\text{om}}), \quad (5.3)$$

where $p_{\text{ca}} = p_{\text{O}_2} + p_{\text{N}_2} + p_{\text{H}_2\text{O,ca}}$ is the cathode total pressure, p_{om} the outlet manifold pressure, and $k_{\text{ca,out}}$ the nozzle constant. With $\dot{m}_{\text{ca,out}}$ defined in (5.3), the mass flow rate of each species at the cathode exit, $\dot{m}_{i,\text{ca}}^{\text{out}}$, is given by:

$$\dot{m}_{i,\text{ca}}^{\text{out}} = \frac{p_i \cdot M_i}{p_{\text{O}_2} \cdot M_{\text{O}_2} + p_{\text{N}_2} \cdot M_{\text{N}_2} + p_{\text{H}_2\text{O,ca}} \cdot M_{\text{H}_2\text{O}}} \cdot \dot{m}_{\text{ca}}^{\text{out}}$$

Electrochemistry principles are used to calculate the rates of oxygen consumption, $\dot{m}_{\text{O}_2}^{\text{react}}$, and water vapor generation, $\dot{m}_{\text{H}_2\text{O}}^{\text{gen}}$, from the stack current I_{st} :

$$\dot{m}_{\text{O}_2}^{\text{react}} = \frac{M_{\text{O}_2} \cdot n_{\text{fc}} \cdot I_{\text{st}}}{4 \cdot F}, \quad \dot{m}_{\text{H}_2\text{O}}^{\text{gen}} = \frac{M_{\text{H}_2\text{O}} \cdot n_{\text{fc}} \cdot I_{\text{st}}}{2 \cdot F},$$

where F is the Faraday constant and n_{fc} the number of FCs in the stack. The net vapor mass flow across the membrane, $\dot{m}_{\text{H}_2\text{O,mem}}$, is calculated in the membrane hydration model, and the vapor injected by the humidifier, $\dot{m}_{\text{H}_2\text{O}}^{\text{hum,ca}}$, is a control input. Lastly, the vapor mass flow that instantaneously condenses once $p_{\text{H}_2\text{O,ca}}$ exceeds the saturation pressure p_{sat} , denoted by $\dot{m}_{\text{H}_2\text{O}}^{\text{cond}}$, is modeled as

$$\dot{m}_{\text{H}_2\text{O}}^{\text{cond}} = \Psi \cdot k_{\text{cond}} \cdot \frac{M_{\text{H}_2\text{O}} \cdot V_{\text{ca}}}{R \cdot T_{\text{st}}} \cdot (p_{\text{H}_2\text{O,ca}} - p_{\text{sat}}),$$

where k_{cond} is the condensation rate constant and Ψ the transition function between condensation ($p_{\text{H}_2\text{O,ca}} > p_{\text{sat}}$) and no condensation ($p_{\text{H}_2\text{O,ca}} \leq p_{\text{sat}}$). To avoid hard switching between 0 and 1, we use an approximation with a sigmoid function [95]:

$$\Psi = \frac{1}{\pi} \cdot \arctan(k_{\text{sig}} \cdot [p_{\text{H}_2\text{O,ca}} - p_{\text{sat}}]) + \frac{1}{2}, \quad (5.4)$$

The constant k_{sig} determines the width of the transition area, thus balancing approximation error and numerical stiffness.

5.1.2 Membrane hydration model

The membrane hydration model based on [84] calculates the water flow across the membrane and its water content. Both are assumed to be uniform across the membrane's surface. The

5 Expansion to include humidity

water transport occurs via two distinct mechanisms. The first, electro-osmotic drag, involves water molecules being dragged from anode to cathode by the protons, where the amount of water is proportional to the electro-osmotic drag coefficient, n_d . Secondly, back-diffusion, driven by the gradient of water concentration across the membrane, typically moves water from cathode to anode, where we assume the water concentration to change linearly over the membrane thickness, t_{mem} . Combining the two mechanisms, the total molar flux from anode to cathode is

$$\dot{n}_{\text{H}_2\text{O},\text{mem},\text{fc}} = n_d \cdot \frac{I_{\text{st}}}{A_{\text{fc}} \cdot F} - D_w \cdot \frac{c_{\text{H}_2\text{O},\text{ca}} - c_{\text{H}_2\text{O},\text{an}}}{t_{\text{mem}}},$$

where A_{fc} denotes the cell area and $c_{\text{H}_2\text{O},i}$ with $i \in \{\text{an}, \text{ca}\}$ the water concentration at the anode and cathode. The coefficient n_d and the diffusion coefficient of water through the membrane D_w vary with the membrane water content Λ_{mem} :

$$n_d = 0.0029 \cdot \Lambda_{\text{mem}}^2 + 0.05 \cdot \Lambda_{\text{mem}} - 3.4 \cdot 10^{-19}$$

$$D_w = D_\Lambda \cdot \exp\left(2416 \cdot \left(\frac{1}{303} - \frac{1}{T_{\text{st}}}\right)\right)$$

$$D_\Lambda = \begin{cases} 10^{-6} & , \Lambda_{\text{mem}} < 2 \\ 10^{-6} \cdot (1 + 2 \cdot (\Lambda_{\text{mem}} - 2)) & , 2 \leq \Lambda_{\text{mem}} \leq 3 \\ 10^{-6} \cdot (3 - 1.67 \cdot (\Lambda_{\text{mem}} - 3)) & , 3 < \Lambda_{\text{mem}} < 4.5 \\ 1.25 \cdot 10^{-6} & , \Lambda_{\text{mem}} \geq 4.5 \end{cases}$$

The membrane water content, as well as the water content of the anode, Λ_{an} , and cathode, Λ_{ca} , are calculated from the respective water activities a_i , $i \in \{\text{an}, \text{ca}, \text{mem}\}$:

$$\Lambda_i = \begin{cases} 0.043 + 17.81 \cdot a_i - 39.85 \cdot a_i^2 + 36.0 \cdot a_i^3 & , 0 < a_i \leq 1 \\ 14 + 1.4 \cdot (a_i - 1) & , 1 < a_i \leq 3, \end{cases}$$

with the water activities given by

$$a_i = \frac{p_{\text{H}_2\text{O},i}}{p_{\text{sat}}}, \quad i \in \{\text{an}, \text{ca}\}, \quad a_{\text{mem}} = \frac{1}{2} \cdot (a_{\text{an}} + a_{\text{ca}})$$

The vapor concentrations are calculated as follows:

$$c_{\text{H}_2\text{O},i} = \frac{\rho_{\text{mem,dry}} \cdot \Lambda_i}{M_{\text{mem,dry}}}, \quad i \in \{\text{an}, \text{ca}\},$$

where $\rho_{\text{mem,dry}}$ denotes the membrane's dry density and $M_{\text{mem,dry}}$ the membrane dry equivalent weight. Finally, the total vapor mass flow considered for the stack is

$$\dot{m}_{\text{H}_2\text{O},\text{mem}} = M_{\text{H}_2\text{O}} \cdot A_{\text{fc}} \cdot n_{\text{fc}} \cdot \dot{n}_{\text{H}_2\text{O},\text{mem},\text{fc}}$$

Note that the expressions for D_Λ and Λ_i are only piecewise continuously differentiable. To ensure numerical robustness and compatibility with gradient-based optimization methods, the hard switching is smoothed using a sigmoid approximation analogous to (5.4). The resulting membrane water transport model introduces an additional dynamic coupling between the electrochemical reaction, gas-phase humidity, and membrane hydration, which is taken into account in the extended state-space model used for control.

5.1.3 Other static relations

We do not consider a dynamic model of the anode. Instead, we assume perfect control of the anode pressure p_{an} , such that no pressure difference exists across the membrane, i.e.,

$$p_{an} = p_{ca}.$$

We further assume the relative humidity in the anode, RH_{an} , to be a control input. Under these assumptions, the hydrogen partial pressure, p_{H_2} , is given by

$$p_{H_2} = p_{an} - RH_{an} \cdot p_{sat}. \quad (5.5)$$

In addition to the power consumed by the compressor motor P_{cm} , the total stack power P_{st} is further reduced by the power consumption of the humidification subsystems on the cathode and anode sides, $P_{hum,ca}$ and $P_{hum,an}$. The constant term P_{aux} simply combines the power consumption of all remaining actuators.

$$P_{fcs,net} = P_{st} - P_{cm} - P_{hum,ca} - P_{hum,an} - P_{aux} \quad (5.6)$$

5.2 High-level control

The adapted control-oriented model of the FCS primarily impacts the high-level control within the target selector. The power-split NMPC incorporates the FCS through a static map of the FCS efficiency. This static map is derived from solving the nominal static optimization problem in the target selector. The following section presents the changes to the control problem formulation of the target selector and the resulting FCS efficiency curves.

Target selector

As previously explained, the target selector calculates the optimal steady-state states x_s and inputs u_s for the FCS with respect to an economic cost function, while ensuring compliance with the FCS power reference $P_{fcs,bus,ref}$ optimized by the power-split NMPC and the system

constraints. The relevant system equations are the ODEs describing the FCS given in (4.6), (5.1)–(5.2), (4.18), (4.5), and (4.1). For clarity and readability, the equations are listed below with their left-hand sides set to zero.

$$\begin{aligned}
 0 &= \frac{\gamma \cdot R_a}{V_{\text{im}}} \cdot (\dot{m}_{\text{cp}}^{\text{out}} \cdot T_{\text{cp}}^{\text{out}} - \dot{m}_{\text{im}}^{\text{out}} \cdot T_{\text{im}}) \\
 0 &= \frac{R \cdot T_{\text{st}}}{M_{\text{O}_2} \cdot V_{\text{ca}}} \cdot (\dot{m}_{\text{O}_2, \text{im}}^{\text{out}} - \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2, \text{ca}}^{\text{out}}) \\
 0 &= \frac{R \cdot T_{\text{st}}}{M_{\text{N}_2} \cdot V_{\text{ca}}} \cdot (\dot{m}_{\text{N}_2, \text{im}}^{\text{out}} - \dot{m}_{\text{N}_2, \text{ca}}^{\text{out}}) \\
 0 &= \frac{R \cdot T_{\text{st}}}{M_{\text{H}_2\text{O}} \cdot V_{\text{ca}}} \left(\dot{m}_{\text{H}_2\text{O}, \text{im}}^{\text{out}} + \dot{m}_{\text{H}_2\text{O}}^{\text{gen}} + \dot{m}_{\text{H}_2\text{O}, \text{mem}} + \dot{m}_{\text{H}_2\text{O}}^{\text{hum, ca}} - \dot{m}_{\text{H}_2\text{O}, \text{ca}}^{\text{out}} - \dot{m}_{\text{H}_2\text{O}}^{\text{cond}} \right) \\
 0 &= \frac{R_a \cdot T_{\text{om}}}{V_{\text{om}}} \cdot (\dot{m}_{\text{ca}}^{\text{out}} - \dot{m}_{\text{om}}^{\text{out}}) \\
 0 &= \dot{m}_{\text{cp}}^{\text{out}} - \dot{m}_{\text{im}}^{\text{out}} \\
 0 &= \frac{1}{J_{\text{cp}}} \cdot (\tau_{\text{cm}} - \tau_{\text{cp}})
 \end{aligned}$$

For brevity, we will refer to the steady-state FCS dynamics as

$$0 = f(x_s, u_s, p) \quad (5.7)$$

with the following states and inputs:

$$x = [p_{\text{im}} \ p_{\text{O}_2} \ p_{\text{N}_2} \ p_{\text{H}_2\text{O}, \text{ca}} \ p_{\text{om}} \ m_{\text{im}} \ \omega_{\text{cp}}]^\top \in \mathbb{R}^7 \quad (5.8)$$

$$u = [I_{\text{st}} \ U_{\text{cm}} \ h_{\text{om}} \ \dot{m}_{\text{hum, ca}} \ \text{RH}_{\text{an}}]^\top \in \mathbb{R}^5 \quad (5.9)$$

The relevant parameter for the FCS model is the stack temperature T_{st} . The optimization problem aims to minimize the same objective function as stated in Section 4.3.1:

$$J_{\text{economic}} = P_{\text{lhv}}$$

This is done considering the following inequality constraints $h(x_s, u_s, p, \varepsilon_s)$. Firstly, the steady-state dynamics in (5.7) are implemented as soft constraints due to potential infeasibility issues:

$$-\varepsilon_{\text{s, sys}} \leq f(x_s, u_s, T_{\text{st}}) \leq \varepsilon_{\text{s, sys}}$$

Another equality constraint, implemented with slack variables, ensures that the FCS output power $P_{\text{fcs, bus}}$ matches the FCS power optimized by the power-split NMPC, referred to as $P_{\text{fcs, bus, ref}}$:

$$-\varepsilon_{\text{s, power}} \leq P_{\text{fcs, bus}} - P_{\text{fcs, bus, ref}} \leq \varepsilon_{\text{s, power}}$$

Furthermore, the steady-state target problem includes the safety-critical constraints of the air path, which comprise the surge and choke compressor boundaries as well as a lower bound on the OER:

$$h_{\text{choke}} = 13.125 \cdot \dot{m}_{\text{cr}} + 0.82 - \Pi \leq 0 \quad (5.10)$$

$$h_{\text{surge}} = \Pi - 278.6852 \cdot \dot{m}_{\text{cr}}^2 - 17.8817 \cdot \dot{m}_{\text{cr}} - 0.7888 \leq 0 \quad (5.11)$$

$$h_{\text{oer}} = \lambda_{\text{O}_2, \text{min}} \cdot \dot{m}_{\text{O}_2}^{\text{react}} - \dot{m}_{\text{O}_2, \text{im}}^{\text{out}} \leq 0 \quad (5.12)$$

In contrast to Section 4.3.1, we do not implement back-off parameters here, as only the nominal target selector is relevant for this section. In addition to these air-path specific constraints, we now also constrain the water content in the anode, cathode and the membrane:

$$\Lambda_{i, \text{min}} \leq \Lambda_i \leq \Lambda_{i, \text{max}}, \quad i \in \{\text{an, ca, mem}\}$$

The steady-state target problem also includes box constraints on the states and inputs:

$$x_{\text{min}} \leq x_s \leq x_{\text{max}}$$

$$u_{\text{min}} \leq u_s \leq u_{\text{max}}$$

Lastly, all slack variables are demanded to be greater or equal to zero:

$$\varepsilon_s \geq 0$$

They are penalized with a quadratic and linear penalty norm in the objective function to allow for exact penalties. In summary, the steady-state target problem is formulated as follows:

$$\begin{aligned} & \underset{x_s, u_s, \varepsilon_s}{\text{minimize}} && P_{\text{lhv}} + \|\varepsilon_s\|_{\mathcal{S}_q} + |\varepsilon_s|_{\mathcal{S}_l} \\ & \text{subject to} && h(x_s, u_s, p, \varepsilon_s) \leq 0. \end{aligned}$$

The steady-state target problem is solved at the same frequency as the sampling time of the power-split NMPC, i.e. every 1 s. This frequency represents the maximum possible rate of change for the power reference $P_{\text{fcs, bus, ref}}$.

With 7 states and 5 inputs, the optimization problem has 12 degrees of freedom. Given 8 (relaxed) equality constraints from the steady-state system dynamics and the required FCS power delivery, 4 degrees of freedom remain. As in Section 4.3.1, 2 of those are used to determine the statically optimal operating point in the compressor map. The remaining 2 degrees of freedom could be utilized to achieve statically efficient water contents in the anode and cathode, fully avoiding overactuation in the low-level control of the FCS.

As detailed in Section 4.3.1, we solve the optimization problem offline for a range of FCS output power and stack temperature combinations. This process generates a look-up table that provides a closed-form solution of the FCS efficiency, which is utilized in the power-split NMPC. Selected FCS efficiency curves are plotted in Fig. 5.2.

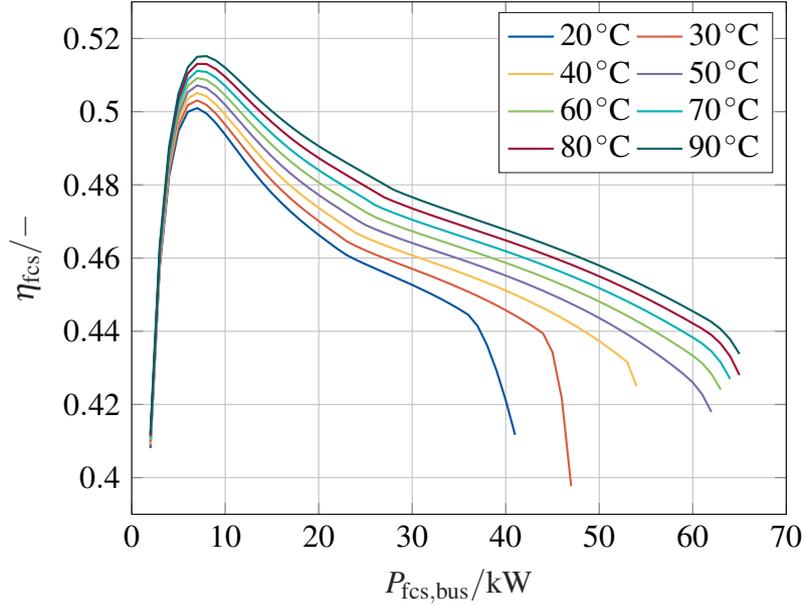


Fig. 5.2: Resulting FCS efficiency curves η_{fcs} from the nominal steady-state target problem with consideration of humidity, plotted over corresponding FCS output powers $P_{\text{fcs,bus}}$ for different stack temperatures T_{st} .

It is noteworthy that the efficiency curves differ significantly from those in Fig. 4.7. Most notably, the range of achievable FCS power decreased across all temperatures, with a particularly significant reduction at lower temperatures. Additionally, the maximum efficiency has declined across all temperatures, primarily because the net FCS power now accounts for the power consumption by the humidifiers.

5.3 Low-level control

5.3.1 Prediction model formulation

The controller-internal model used by the tracking MPC consists of the differential and algebraic equations that capture the dynamics of the air path with consideration of humidity, the membrane hydration, the terminal voltage, and the power generation for both the FCS and the battery system, as detailed in Section 4.2 and Section 5.1. The model comprises seven states and six inputs. The run-time parameters of the model are the stack temperature, T_{st} , and the SoC of the battery, x_{SoC} .

$$x = [p_{\text{im}} \quad p_{\text{O}_2} \quad p_{\text{N}_2} \quad p_{\text{H}_2\text{O,ca}} \quad p_{\text{om}} \quad m_{\text{im}} \quad \omega_{\text{cp}}]^{\top} \in \mathbb{R}^7 \quad (5.13)$$

$$u = [I_{\text{st}} \quad U_{\text{cm}} \quad h_{\text{om}} \quad \dot{m}_{\text{hum,ca}} \quad \text{RH}_{\text{an}} \quad I_{\text{bat}}]^{\top} \in \mathbb{R}^6 \quad (5.14)$$

$$p = [T_{st} \quad x_{SoC}]^T \in \mathbb{R}^2 \quad (5.15)$$

The relevant ODEs are summarized below:

$$\begin{aligned} \frac{dp_{im}}{dt} &= \frac{\gamma \cdot R_a}{V_{im}} \cdot (\dot{m}_{cp}^{out} \cdot T_{cp} - \dot{m}_{im}^{out} \cdot T_{im}) \\ \frac{dp_{O_2}}{dt} &= \frac{R \cdot T_{st}}{M_{O_2} \cdot V_{ca}} \cdot (\dot{m}_{O_2,im}^{out} - \dot{m}_{O_2}^{react} - \dot{m}_{O_2,ca}^{out}) \\ \frac{dp_{N_2}}{dt} &= \frac{R \cdot T_{st}}{M_{N_2} \cdot V_{ca}} \cdot (\dot{m}_{N_2,im}^{out} - \dot{m}_{N_2,ca}^{out}) \\ \frac{dp_{H_2O,ca}}{dt} &= \frac{R \cdot T_{st}}{M_{H_2O} \cdot V_{ca}} \cdot (\dot{m}_{H_2O,im}^{out} + \dot{m}_{H_2O}^{gen} + \dot{m}_{H_2O,mem} + \dot{m}_{H_2O}^{hum,ca} - \dot{m}_{H_2O,ca}^{out} - \dot{m}_{H_2O}^{cond}) \\ \frac{dp_{om}}{dt} &= \frac{R_a \cdot T_{om}}{V_{om}} \cdot (\dot{m}_{ca}^{out} - \dot{m}_{om}^{out}) \\ \frac{dm_{im}}{dt} &= \dot{m}_{cp}^{out} - \dot{m}_{im}^{out} \\ \frac{d\omega_{cp}}{dt} &= \frac{1}{J_{cp}} \cdot (\tau_{cm} - \tau_{cp}) \end{aligned}$$

The main change compared to the ODEs in Section 4.4.1 is that we now account for the presence of vapor in the FC stack. This adjustment does not impact the differential equations resulting for the inlet or outlet manifold or the compressor, as these components do not have species-specific source and sink terms. However, this does affect the cathode flow model. Consequently, we now have differential equations for the oxygen, nitrogen, and vapor in the cathode. By assuming perfect control of the pressure difference across the membrane and using a simplified static consideration of the anode side (see Section 5.1.3), we can further calculate the hydration levels of the anode, cathode, and the membrane. Here, we choose the following seven variables for output tracking:

$$y = [P_{del} \quad P_{fcs,bus} \quad P_{bat,bus} \quad \dot{m}_{cr} \quad \Pi \quad \Lambda_{ca} \quad \Lambda_{an}]^T \in \mathbb{R}^7 \quad (5.16)$$

The delta formulation is adopted for MPC prediction model formulation and the state and input vector of the augmented state space are the following:

$$\tilde{x} = [p_{im} \quad p_{O_2} \quad p_{N_2} \quad p_{H_2O,ca} \quad p_{om} \quad m_{im} \quad \omega_{cp} \quad \dots \quad I_{st} \quad U_{cm} \quad h_{om} \quad \dot{m}_{hum,ca} \quad RH_{an} \quad I_{bat}]^T \in \mathbb{R}^{13} \quad (5.17)$$

$$\tilde{u} = [\dot{I}_{st} \quad \dot{U}_{cm} \quad \dot{h}_{om} \quad \dot{m}_{hum,ca} \quad \dot{RH}_{an} \quad \dot{I}_{bat}]^T \in \mathbb{R}^6 \quad (5.18)$$

Note that from this point forward, the augmented state \tilde{x} and input vector \tilde{u} will simply be referred to as x and u .

5.3.2 Optimal control problem

The structure of the OCP presented in (4.65) remains unchanged. The inclusion of humidity in the prediction model automatically affects the tracking objective and the constraints. Additionally, we incorporate specific inequality constraints to account for humidity-related parameters. These constraints include bounds on the water content in the anode, cathode and membrane, as well as on the relative humidity in the anode:

$$\Lambda_i^{\min} - \varepsilon_{\Lambda_i}^{\min} \leq \Lambda_i \leq \Lambda_i^{\max} + \varepsilon_{\Lambda_i}^{\max}$$

$$0 \leq \text{RH}_{\text{an}} \leq 1$$

All in all, the OCP reads:

$$\begin{aligned} \underset{x(\cdot), u(\cdot), \varepsilon}{\text{minimize}} \quad & \frac{1}{2} \cdot \int_{t_0=t_k}^{t_f} (\|y(t) - y_{\text{ref}}(t)\|_Q^2 + \|u(t)\|_R^2) dt + \|\varepsilon\|_{S_q}^2 + |\varepsilon|_{S_l} \\ \text{subject to} \quad & 0 = x(t_0) - \hat{x}_0, \\ & \dot{x} = f(x(t), u(t), p), \quad \forall t \in [t_0, t_f], \\ & I_{\text{st}}^{\min} \leq I_{\text{st}}(t), \quad \forall t \in [t_0, t_f], \\ & h_{\text{om}}^{\min} \leq h_{\text{om}}(t) \leq h_{\text{om}}^{\max}, \quad \forall t \in [t_0, t_f], \\ & \text{RH}_{\text{an}}^{\min} \leq \text{RH}_{\text{an}}(t) \leq \text{RH}_{\text{an}}^{\max}, \quad \forall t \in [t_0, t_f], \\ & u_{\min} \leq u(t) \leq u_{\max}, \quad \forall t \in [t_0, t_f], \\ & \varepsilon_{\text{choke}} \geq 13.125 \cdot \dot{m}_{\text{cr}}(t) + 0.82 - \Pi(t), \quad \forall t \in [t_0, t_f], \\ & \varepsilon_{\text{surge}} \geq \Pi(t) - 278.6852 \cdot (\dot{m}_{\text{cr}}(t))^2 - 17.8817 \cdot \dot{m}_{\text{cr}}(t) - 0.7888, \quad \forall t \in [t_0, t_f], \\ & \varepsilon_{\lambda_{\text{O}_2}} \cdot \dot{m}_{\text{O}_2}^{\text{react}}(t) \geq \lambda_{\text{O}_2}^{\min} \cdot \dot{m}_{\text{O}_2}^{\text{react}}(t) - \dot{m}_{\text{O}_2, \text{min}}^{\text{out}}(t), \quad \forall t \in [t_0, t_f], \\ & \Lambda_{\text{an}}^{\min} - \varepsilon_{\Lambda_{\text{an}}}^{\min} \leq \Lambda_{\text{an}}(t) \leq \Lambda_{\text{an}}^{\max} + \varepsilon_{\Lambda_{\text{an}}}^{\max}, \quad \forall t \in [t_0, t_f], \\ & \Lambda_{\text{mem}}^{\min} - \varepsilon_{\Lambda_{\text{mem}}}^{\min} \leq \Lambda_{\text{mem}}(t) \leq \Lambda_{\text{mem}}^{\max} + \varepsilon_{\Lambda_{\text{mem}}}^{\max}, \quad \forall t \in [t_0, t_f], \\ & \Lambda_{\text{ca}}^{\min} - \varepsilon_{\Lambda_{\text{ca}}}^{\min} \leq \Lambda_{\text{ca}}(t) \leq \Lambda_{\text{ca}}^{\max} + \varepsilon_{\Lambda_{\text{ca}}}^{\max}, \quad \forall t \in [t_0, t_f]. \end{aligned} \tag{5.19}$$

5.3.3 Transcription to nonlinear program

For the transcription of the OCP in (5.19) into an NLP, we follow the same approach as in Sections 4.4.3 and 4.4.4 to select the sampling time T_s , numerical integration scheme, and prediction horizon N .

For the **sampling time**, we simulate the model along a FCS power trajectory covering different operating points, as described in Section 4.4.3. The power range for this analysis is slightly

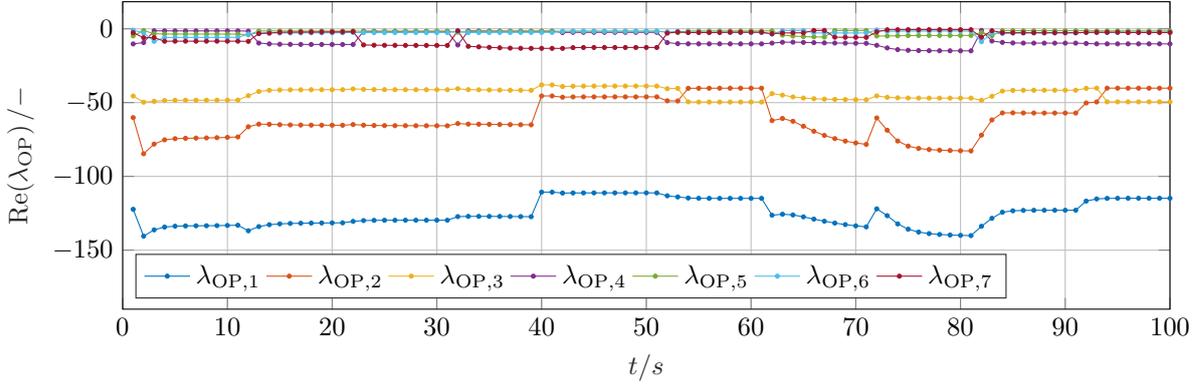


Fig. 5.3: Real parts of eigenvalues for each system mode and operating point, with humidity consideration.

different, ranging from 7.6kW to 65kW. Upon evaluating the time constants, we obtain the fastest time to be 0.0071 s, while the slowest is 1.67 s. Thus, the range is similar to that found in Section 4.4.3, with the fastest mode being of order $\mathcal{O}(10^{-2})$ seconds.

Comparing Fig. 4.9 to Fig. 5.3, which display the real parts of the eigenvalues over time, we observe that in Fig. 5.3, the fastest mode is consistently faster than 10ms over the considered operating range, unlike the fastest mode in Fig. 4.9. Based on this, we choose $T_s = 8$ ms.

The trends observed for the **numerical integration schemes** are very similar to those in Section 4.4.3. The primary difference is that numerical integration of the model with humidity consideration results in a larger relative integration error e_{rel} and is computationally more expensive. Consequently, the Pareto front in Fig. 5.4 shifts diagonally to the upper right compared to Fig. 4.10. Using an ERK4 scheme and one integration step still yields an accuracy of $\mathcal{O}(10^{-6})$, which is deemed acceptable.

Upon implementing the SQP algorithm as detailed in Section 4.4.4, we analyze the choice of the **prediction horizon**. Due to the increased model complexity introduced by humidity, the same prediction horizon results in higher computational effort compared to the case without humidity. Comparing Fig. 4.11 and Fig. 5.5, it can be observed that the CPU time increases more rapidly with growing N . The relative cost difference ΔJ_{rel} cannot be directly compared, as the reference value $J_{\text{ref}} = J(N = 95)$ differs from that used in Chapter 4. Nevertheless, the qualitative trends remain consistent. Overall, the results indicate a clear trade-off between computational effort and performance improvement. Based on this trade-off, a prediction horizon of $N \leq 25$ is selected to ensure real-time feasibility while maintaining satisfactory closed-loop performance.

5 Expansion to include humidity

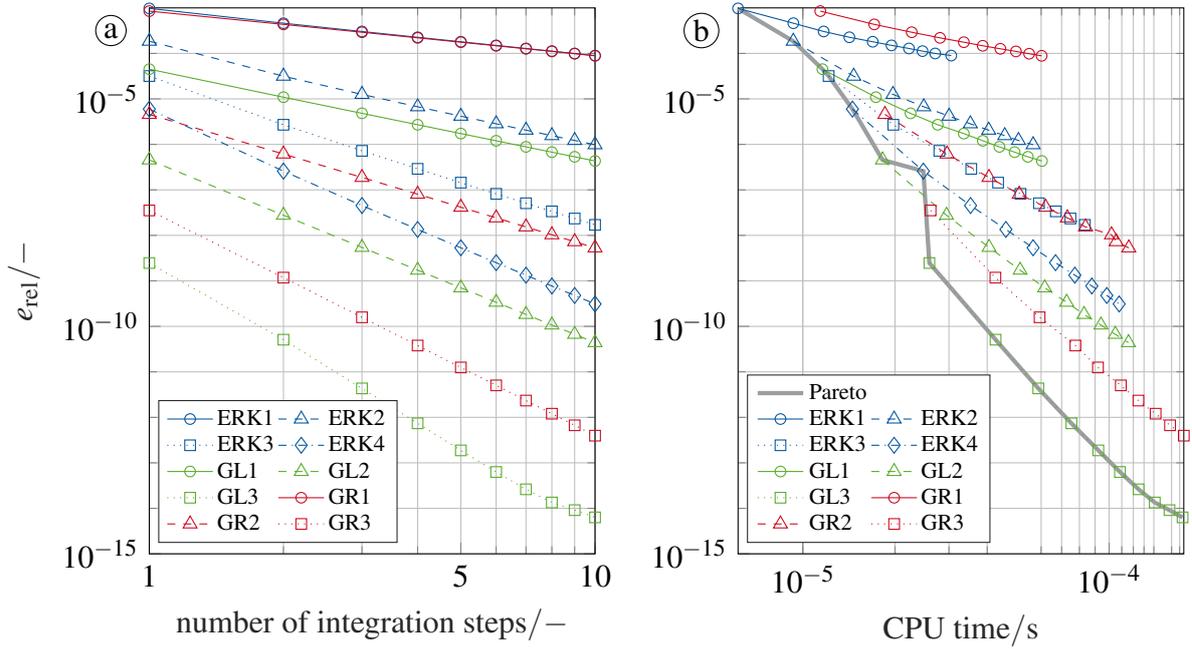


Fig. 5.4: (a) Relative discretization error vs. number of integration steps for tested integration methods, (b) relative discretization error vs. mean CPU time for tested integration schemes.

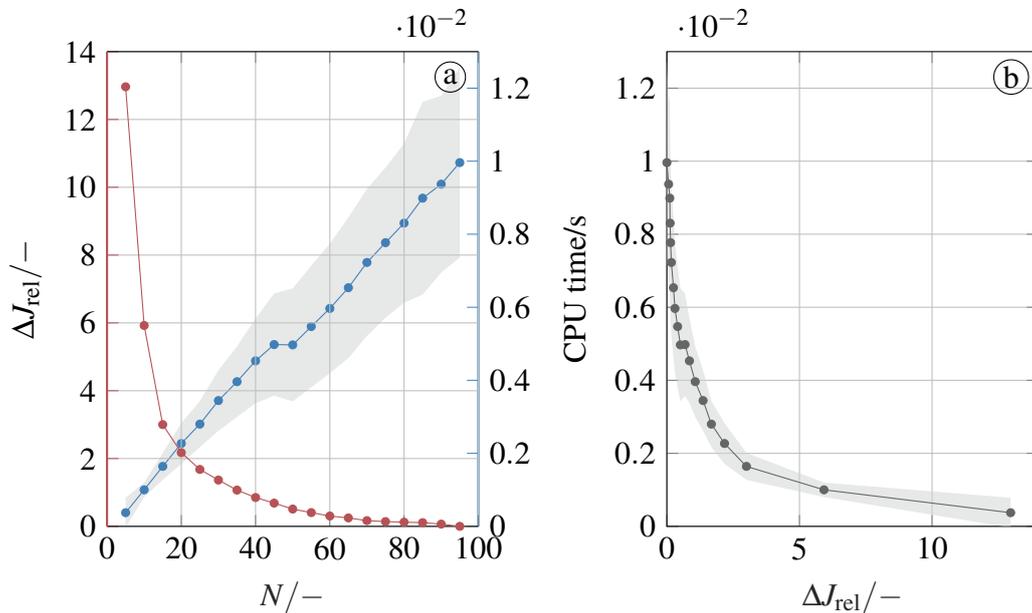


Fig. 5.5: (a) Relative cost difference and CPU time vs. prediction horizon, (b) CPU time vs. relative cost difference.

6 Simulative studies and results

In this chapter, simulative studies are conducted to answer research questions 2 to 5. The focus is on evaluating the performance, feasibility, and robustness of the proposed hierarchical MPC framework for automotive FCSs.

First, the tracking MPC in the low-level control layer is investigated under ideal conditions and independently of the high-level control layer in Sections 6.1 and 6.2. This allows for an isolated assessment of the prediction model and the employed optimization algorithms. Subsequently, Section 6.3 presents closed-loop simulation results of the complete hierarchical MPC scheme interacting with a comprehensive plant model. Finally, the impact of humidity consideration on low-level control performance is analyzed in Section 6.4.

For the evaluation of control performance, the performance metrics introduced in Section 2.4 are employed. All studies are conducted using a demanding step power reference scenario with a total duration of 15 s. The load cycle consists of a power step from 2 kW to 40 kW at $t = 6$ s, followed by a step back to 2 kW at $t = 9$ s. This scenario is chosen because it is highly challenging for fuel-cell-dominant hybrid electric vehicles and exposes the fundamental trade-off between fast dynamic power delivery and adherence to safety-critical air-path constraints.

During a rapid power increase, sufficient OER requires a fast rise in cathode mass flow. However, the admissible mass flow is limited by the compressor choke boundary, while the associated pressure dynamics are comparatively slow. Conversely, during a rapid power decrease, the surge boundary becomes the dominant limiting factor. As such, the step scenario is particularly well suited to assess the dynamic capabilities and constraint-handling properties of the investigated control approaches.

6.1 Comparison of tracking LTV MPC and tracking NMPC

The tracking MPC in the low-level control layer directly actuates the fuel-cell-dominant hybrid electric vehicle and relies on a physics-based prediction model derived in Section 4.2. Since the underlying system dynamics are nonlinear, this nonlinearity must be addressed in the controller design.

In this section, two alternative realizations of the tracking controller are investigated. First, a nonlinear MPC based on a sequential quadratic programming approach is considered. This controller is referred to as SQP NMPC and is implemented by applying the SQP algorithm in a real-time iteration (RTI) scheme, i.e., one SQP iteration is performed per sampling instant. Second, the nonlinear prediction model is linearized at each operating point, resulting in a linear time-varying prediction model and a corresponding QP-based controller, referred to as LTV MPC.

While both approaches solve one quadratic programming per sampling instant, they fundamentally differ in how the prediction model is treated. In the LTV MPC, the system matrices remain constant over the entire prediction horizon. In contrast, the SQP NMPC updates the linearized system matrices at each prediction step, effectively going beyond linearization at a single operating point. This section investigates whether such a nonlinear prediction is required to meet the control requirements of automotive FCSs.

Research question 2:

Is a linearized prediction sufficient to meet the control requirements for automotive FCSs when applying MPC?

If the linear MPC, which relies on successive linearization of the prediction model at each operating point, proves sufficient, it can reduce the computational burden.

To answer this question, we test both controllers in the step evaluation scenario with pre-optimized references. This ensures that both controllers receive identical references, excluding the high-level control layer from this evaluation. The controllers are implemented in MATLAB, using CasADi for sensitivity calculations. For both approaches, the resulting QPs are solved using the sparsity-exploiting solver from the HPIPM framework. MPC weighting matrices and HPIPM solver options are chosen to be the same for both controllers. The exact Hessian is used for the RTI SQP NMPC algorithm. Furthermore, we test the controllers under ideal conditions, i.e., they run against the nonlinear controller-internal model and we assume full state measurement. Consequently, there is no plant model mismatch apart from potential linearization errors. All closed-loop simulations for this evaluation are performed in MATLAB.

Since the focus of this evaluation is the low-level control, the relevant performance metrics considered here are dynamic power delivery and adherence to safety-critical constraints, quantified by the performance metrics $t_{90,P}$ and h_{\max} . In this case, if $h_{\max} \leq 0$ then all critical air-path constraints defined in Section 2.4 are satisfied.

Upon evaluating the performance metrics, we see that they are nearly identical for the SQP NMPC and the LTV MPC, as shown in Table 6.1. The values slightly favor for LTV MPC, but the differences are minimal. Both controllers achieve rapid power delivery well below one second and exhibit slight violations of the safety-critical constraints. These observations are confirmed by the closed-loop trajectories in Fig. 6.1 and Fig. 6.2. However, a closer examination reveals that while the SQP NMPC enables a smooth power increase, the LTV MPC

Table 6.1: Performance metrics for dynamic power delivery and constraint satisfaction, evaluated for SQP NMPC with one iteration and LTV MPC in the step evaluation scenario.

	$t_{90,P_{\text{drive}}}$	$t_{90,P_{\text{fcs}}}$	h_{max}
NMPC SQP	0.287	0.4	$4.253 \cdot 10^{-4}$
LTV MPC	0.28	0.4	$3.05 \cdot 10^{-4}$

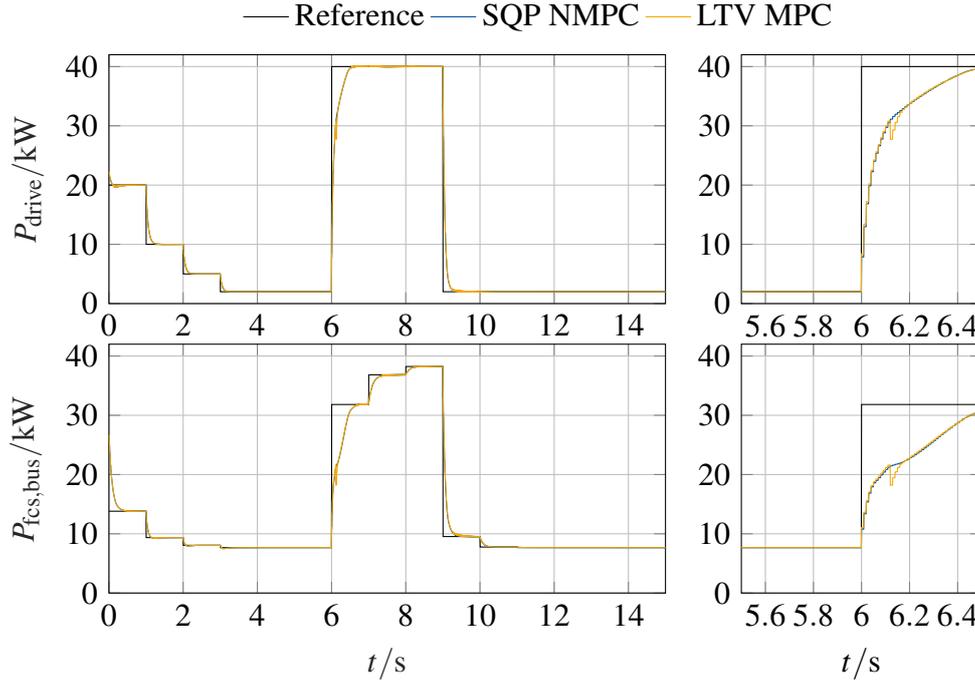


Fig. 6.1: Closed-loop tracking results of total power P_{drive} and FCS power $P_{\text{fcs,bus}}$ for SQP NMPC with one iteration and LTV MPC in the step evaluation scenario.

suffers from a sudden power decrease during acceleration, as shown in Fig. 6.1. Specifically, there is a power loss of around 3 kW, equating to around 10% with respect to the first power step-up reference for the FCS.

Examining the safety-critical constraints in Fig. 6.2, we observe the impact of this sudden power decrease on the OER trajectory and the choke constraint trajectory. Additionally, we see that the compressor frequently operates at its boundaries. The minor constraint violations, evaluated in Table 6.1, occur at the choke boundary. During the power step-up, we see that both the OER and the choke constraint trajectories quickly approach their respective boundaries. The compressor ramps up, as shown by U_{cm} in Fig. 6.3, to supply more oxygen. Simultaneously, the back-pressure valve h_{om} quickly closes for a short amount of time to aid in pressure increase and therefore avoid choking before it opens again to avoid suppressing mass flow. Nevertheless,

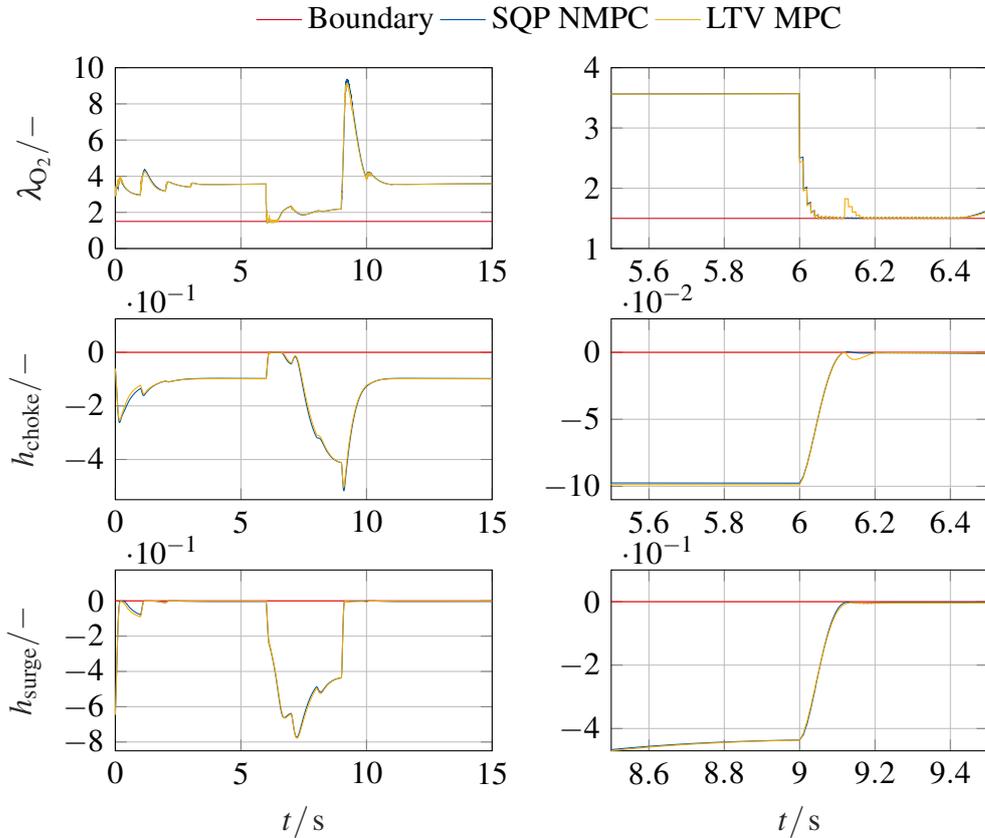


Fig. 6.2: Constraint trajectories for OER and compressor choke and surge, resulting from closed-loop control with SQP NMPC using one iteration and LTV MPC in the step evaluation scenario.

the sudden and substantial power increase leads to minimal choke constraint violations.

The reason for the sudden power decrease is related to the choke boundary, as becomes apparent when examining the controller-internal predictions around the time the decrease occurs. Fig. 6.4 shows selected output and constraint trajectories (in black) and their prediction trajectories over the entire prediction horizon (in blue) for the relevant time window. The top row presents the trajectories and predictions based on the nonlinear prediction model, while the bottom row shows those based on the linearized prediction model. Both sets are derived from the state and control input prediction trajectories, x_k and u_k with $k = 0, \dots, N - 1$, resulting from the closed-loop simulation with the LTV MPC.

The most significant difference between linear and nonlinear prediction can be seen in the choke constraint prediction trajectories. While the linearized system predicts the system will stay below or on the boundary, the nonlinear system, based on the same state and input values, shows that these optimized values would lead to constraint violation. Although the interior-point algorithm and the tolerances set for HPIPM allow slight constraint violation in the SQP NMPC,

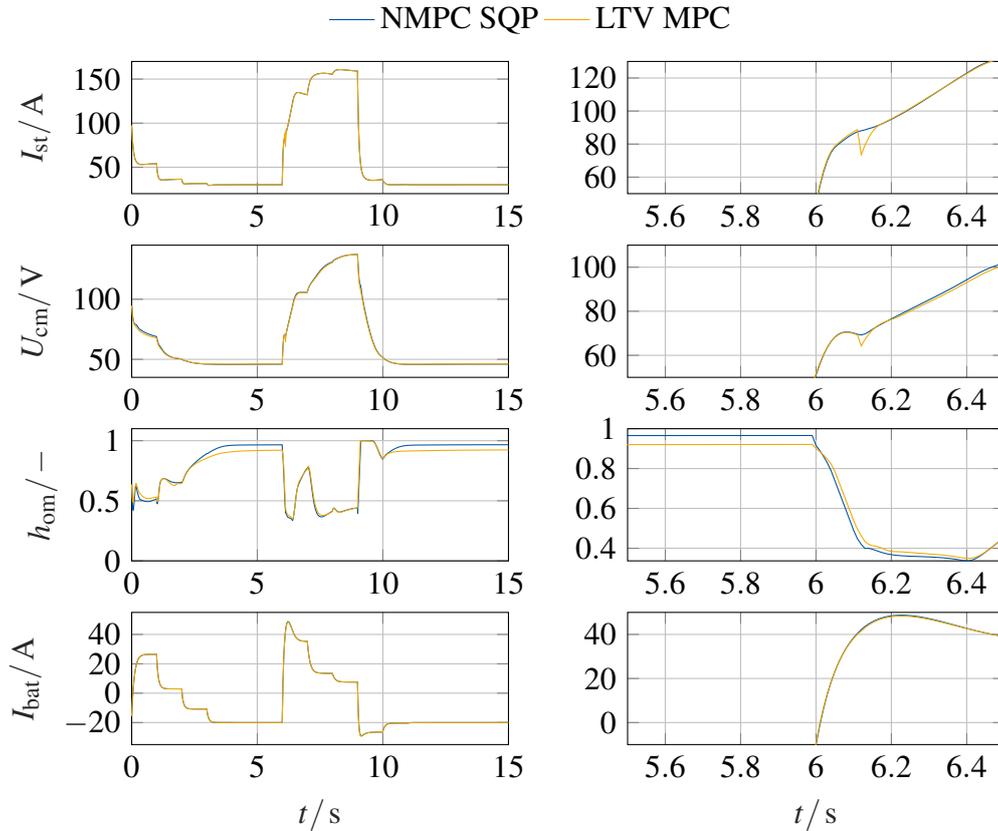


Fig. 6.3: Control input trajectories resulting from closed-loop control with SQP NMPC using one iteration and LTV MPC in the step evaluation scenario.

the nonlinear predictions correctly anticipate the system behavior. Conversely, the LTV MPC does not expect the system to violate the constraint; when it does, it changes the inputs abruptly to mitigate the issue. Additionally, we observe that the predictions for the surge constraint trajectory also differ significantly between linear and nonlinear predictions.

Based on the conducted simulations, a linearized prediction is not sufficient to robustly meet the control requirements of automotive FCSs when highly dynamic power tracking is required. Although the LTV MPC achieves comparable rise times and only minor constraint violations in terms of aggregated performance metrics, it exhibits fundamentally different closed-loop behavior under demanding operating conditions.

In particular, inaccuracies in the linearized prediction of the choke and surge constraints lead to abrupt corrective control actions once constraint violations occur. In the investigated scenario, this manifests as an approximately 10% sudden power reduction during acceleration. The nonlinear prediction employed by the SQP NMPC, on the other hand, anticipates the constraint activation more accurately and enables smoother control actions without such artifacts.

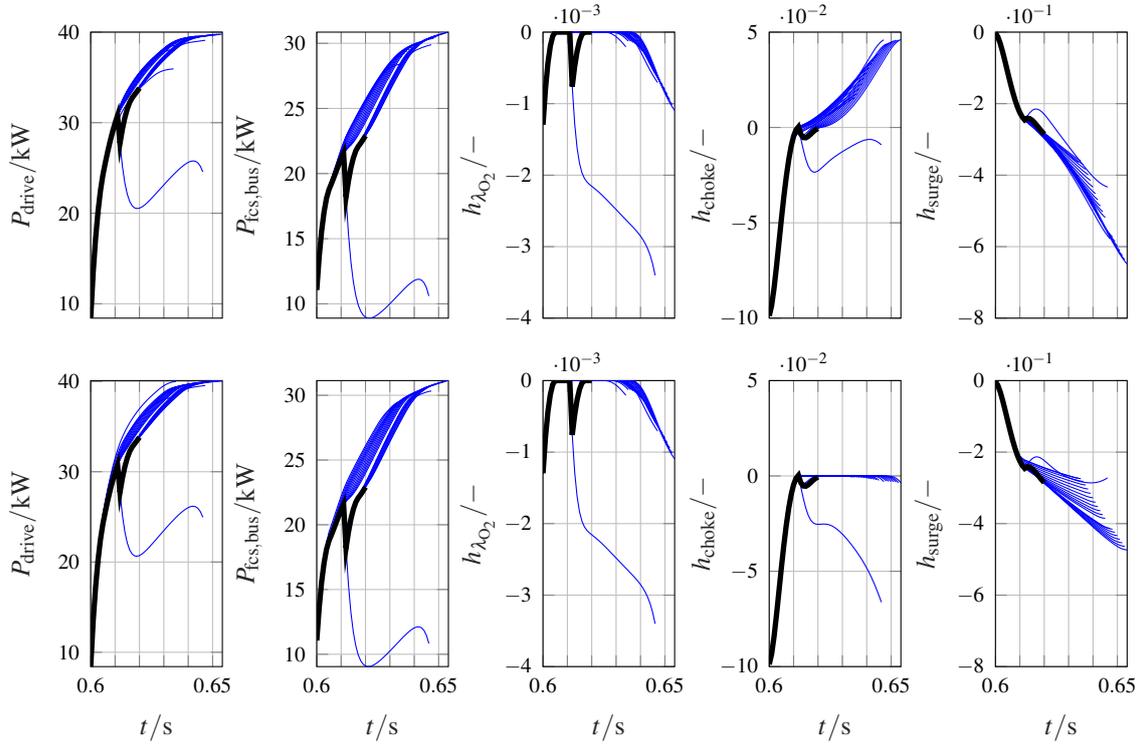


Fig. 6.4: Selected output and constraint trajectories and their prediction trajectories based on control results from LTV MPC. Top row: Predictions are calculated with nonlinear prediction model; bottom row: Predictions are calculated with linearized prediction model.

These results indicate that while LTV MPC can be a viable solution for moderately dynamic operating scenarios, nonlinear prediction is advantageous—and in highly dynamic cases essential—for achieving smooth power delivery while respecting safety-critical constraints in fuel-cell-dominant hybrid electric vehicles.

6.2 Real-time feasibility of tracking NMPC

Based on the findings in Section 6.1, an NMPC based on an SQP approach is selected for the low-level control layer. In contrast to the high-level control, the tracking NMPC directly provides control inputs to the fuel-cell-dominant hybrid electric vehicle and therefore operates at a significantly higher update rate.

As a consequence, the tracking controller must explicitly account for the fastest relevant system dynamics while simultaneously meeting strict real-time requirements. Within the proposed hierarchical MPC framework, the real-time feasibility of the SQP-based NMPC implementation on embedded hardware therefore represents a critical aspect for practical applicability.

This section evaluates whether an effective tracking NMPC can be implemented in a real-time capable manner on embedded hardware, despite the computational complexity associated with nonlinear optimization. Accordingly, it addresses the following research question:

Research question 3:

Is it possible to design an effective and real-time feasible MPC algorithm for automotive FCSs?

To answer this question, we again use the step evaluation scenario. Due to the sudden and substantial changes in the operating point in both directions, this scenario presents a challenging evaluation case. It is therefore not only suitable for quantifying the controller’s potential with respect to control performance but also its computational resource requirements.

As in Section 6.1, the control algorithm is evaluated under idealized conditions, i.e., assuming full state measurement and no plant-model mismatch. While this setup does not reflect a fully realistic deployment scenario, it allows for an explicit quantification of performance losses solely induced by enforcing real-time constraints on embedded hardware. As such, the results obtained in this section can serve as a baseline for further testing under more realistic conditions, including plant-model mismatch, measurement noise, and disturbance estimation.

Based on the analysis in Section 4.4.3, a sampling time of $T_s = 10$ ms is selected, which aligns with the fastest relevant dynamics of the air-path subsystem of the FCS. To ensure real-time capability, the maximum turnaround time of the controller must therefore not exceed the specified sampling time.

To evaluate the controller’s effectiveness, we again use the $t_{90,P}$ and h_{\max} performance metrics to quantify dynamic power delivery and potential constraint violations. Additionally, a closed-loop solution using the solver IPOPT to directly solve the NLP defined in (4.66) is computed as a baseline solution to quantify for the potential control performance losses due to real-time implementation.

Three different solution strategies are evaluated:

1. Nonlinear MPC solved using the interior-point NLP solver IPOPT
2. SQP-based NMPC using an exact second-order Hessian and solved with the structure-exploiting QP solver HPiPM OCP
3. SQP-based NMPC using a Gauss-Newton Hessian approximation and solved with HPiPM OCP

The baseline solution obtained from IPOPT is computed in MATLAB using the ‘Opti stack’ from CasADi with a prediction horizon of $N = 35$. Based on the analysis in Section 4.4.3, we decided that $N = 35$ strikes a good balance between computational cost and control performance. The closed-loop control simulation results for output tracking and the corresponding safety-critical constraint trajectories are provided in blue in Fig. 6.5 and Fig. 6.6, respectively.

6 Simulative studies and results

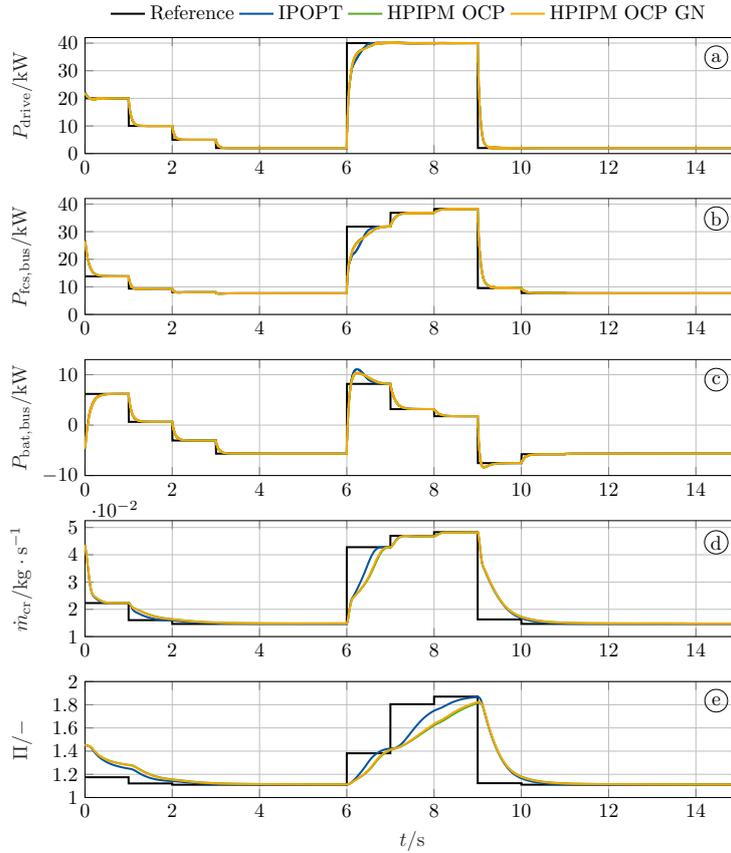


Fig. 6.5: Closed-loop output tracking results for all tested NMPC algorithms.

For implementation of the SQP algorithm detailed in Section 4.4.4, the tunable control parameters we consider here, apart from the prediction horizon N , are the step length α and the maximum number of QP iterations we allow per SQP iteration. According to the RTI scheme, the number of SQP iterations is fixed to one. While the step length α affects the controller performance and is likely to be re-tuned in closed loop with a more realistic plant model or on a test bench, the maximum admissible QP iterations and the prediction horizon N significantly impact the computational requirements.

For this evaluation, we take a full step, i.e. $\alpha = 1$, as this provides an effective closed-loop control solution.

For the other two parameters, we follow a simple, sequential procedure: First, we perform closed-loop simulations where we reduce the maximum number of admissible QP iterations for both SQP algorithms until the maximum relative steady-state tracking error among all outputs exceeds 2%. These closed-loop simulations are performed in MATLAB using CasADi for sensitivity calculations and with a prediction horizon of $N = 35$. In this validation scenario, the minimum maximum amount of QP iterations is 9 for both SQP algorithms with exact Hessian

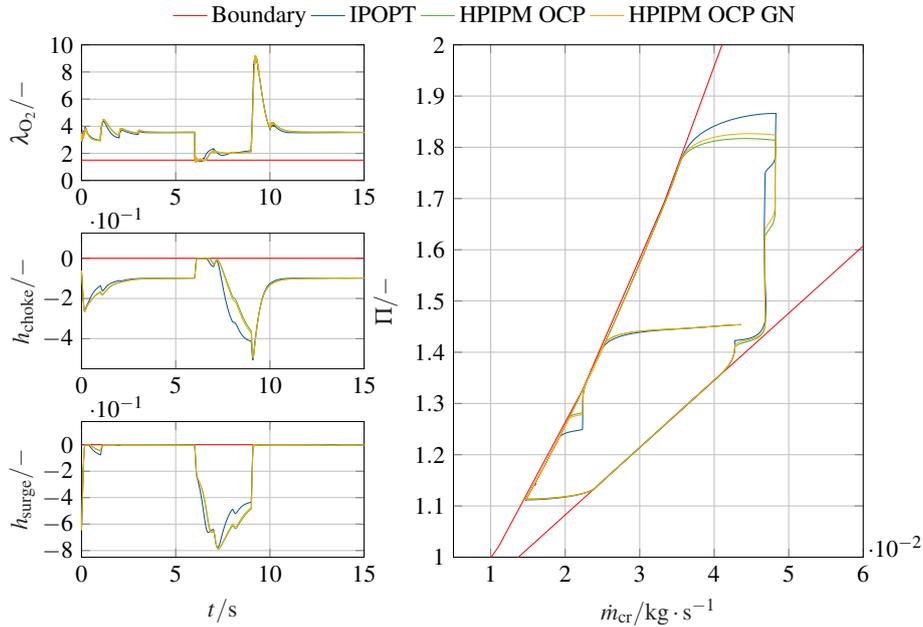


Fig. 6.6: Closed-loop constraint trajectories for all NMPC algorithms.

and GN Hessian. With 9 admissible QP iterations per SQP iteration, the largest relative steady-state tracking error is slightly above 1.5%. However, with 8 QP iterations, the largest relative steady-state tracking error is almost 4%. Usually, we observe the biggest relative steady-state tracking error for the output pressure ratio Π due to the slow compressor dynamics.

Subsequently, we implement the algorithms in SIMULINK using S-functions obtained via C-code generation from CasADi's MATLAB interface. The sparse QP solver from HPIPM is interfaced via S-functions as well. For experiments on the embedded hardware, the C-code implementation of the control algorithm is generated using the built-in MATLAB toolchain and compiled along with the code of the solver. The embedded hardware used is the dSPACE MicroLabBox, which features the DS1202 base board, equipped with an NXP (Freescale) QorIQ P5020 dual-core processor running at 2 GHz and 128 MB of flash memory.

In simulative experiments on the MicroLabBox, the maximum controller turnaround time t_{MLB} , defined as the worst-case computation time per control update, is evaluated for the SQP algorithms. Real-time feasibility is considered achieved if $t_{MLB} \leq T_s$. As real-time feasibility for a prediction horizon of $N = 35$ is not achieved on the hardware, we successively reduce the prediction horizon until real-time feasibility is given. The results show that the SQP algorithm using the sparse HPIPM OCP solver is real-time feasible for $N = 21$ with exact Hessian and $N = 22$ with the GN Hessian, with a maximum turnaround time of $t_{MLB} = 9.7$ ms.

The relevant closed-loop control results for the real-time capable tracking NMPC are shown in Fig. 6.5 and Fig. 6.6, with exact Hessian results in green and GN Hessian results in orange. The

results between the two SQP algorithms are very similar, but they differ from the baseline IPOPT solution. Tracking of the compressor setpoints is slower for the SQP algorithms compared to the baseline solution, and the setpoints are not always reached, as shown in Fig. 6.5d and Fig. 6.5e. Consequently, there is also a noticeable difference in power tracking performance; however it is not necessarily disadvantageous. The power trajectories P_{drive} and $P_{\text{fcs,bus}}$, resulting from the real-time capable NMPCs, reach their respective references later for the power step up at 6s, but the acceleration is smoother and at times also faster compared to the baseline solution, particularly visible for the FCS (Fig. 6.5b). Regarding the safety critical constraints, the time spent on the OER boundary and choke boundary is minimally longer for the SQP algorithms. The performance metrics evaluated in Tab. 6.2 show minimal violation of the choke boundary,

Table 6.2: Maximum turnaround time and performance metrics for dynamic power delivery and constraint satisfaction, evaluated for all tested NMPC algorithms in the step evaluation scenario.

	$t_{90,P_{\text{drive}}}$	$t_{90,P_{\text{fcs}}}$	h_{max}	t_{MLB}
NMPC IPOPT	0.289	0.404	$4.549 \cdot 10^{-4}$	–
SQP NMPC HPIPM	0.219	0.45	$4.747 \cdot 10^{-4}$	9.7 ms
SQP NMPC HPIPM GN	0.22	0.459	$4.701 \cdot 10^{-4}$	9.7 ms

similar across all tested algorithms. Similar to Section 6.1, the dynamic power delivery metric does not reflect the observations made from Fig. 6.5. While, according to the $t_{90,P}$ values, the dynamic power delivery is slower for the SQP algorithms in terms of the FCS and faster for the total power, Fig. 6.5 shows that the real-time capable controllers reach both total power and FCS power references later but with a more consistent increase during acceleration. Consequently, the battery power peak during the acceleration process is less pronounced than in the baseline solution.

Overall, the tracking NMPC implemented using an SQP algorithm and executed on the dSPACE MicroLabBox achieves real-time feasibility for the considered control task. For a prediction horizon of $N = 21\text{--}22$, the maximum turnaround time remains below the sampling time of $T_s = 10\text{ms}$, thereby satisfying the real-time requirement.

Despite the reduced prediction horizon and the resulting approximation of the baseline solution obtained with IPOPT, the closed-loop control performance remains comparable. All tested NMPC variants reach the relevant power references within one second and exhibit only minimal violations of the choke boundary during the demanding step evaluation scenario. These constraint violations remain small and are comparable across all tested algorithms.

Under the considered idealized conditions, the results demonstrate that an effective and real-time capable NMPC implementation for automotive FCSs can be achieved on embedded hardware. It is emphasized that these conclusions apply to the low-level control task focusing

on the air-path dynamics of the FCS and do not yet account for uncertainties introduced by plant-model mismatch or state estimation.

6.3 Simulative validation of the overall control concept

In this section, we present the control results of the hierarchical MPC scheme in closed loop with a more comprehensive plant model of the FCHEV. The simulation setup and the structure of the control algorithm are depicted in Fig. 4.13, where the tracking NMPC derived in Section 6.2 is employed as the low-level controller.

When closing the loop with a higher-fidelity plant model, plant-model mismatch inevitably arises. Despite the presence of a disturbance estimator, this mismatch can lead to noticeable degradation in control performance, in particular with respect to safety-critical constraint satisfaction during highly dynamic operation. First, the implemented disturbance estimator is designed to compensate steady-state offsets only. As demonstrated in previous sections, however, the safety-critical constraints are most restrictive during highly dynamic operation. Second, as discussed in Section 4.5, the estimated disturbances are only used for reference adaptation outside the low-level NMPC. Consequently, disturbances are not explicitly considered in the inequality constraints of the tracking NMPC, such as the surge and choke constraints formulated on the compressor mass flow \dot{m}_{cr} and pressure ratio Π .

Finally, humidity is considered in the plant model but not in the controller-internal prediction model. The additional vapor partial pressure alters the air-path dynamics and pressure levels. Although the state observer converges, its accuracy remains fundamentally limited by the model mismatch. As a result, constraint satisfaction, in particular with respect to the OER, cannot be guaranteed in closed loop.

To mitigate these effects, we make use of the back-off parameters introduced in Section 4.3.1 within the static optimization problem solved by the target selector. This effectively tightens the steady-state constraints and yields more conservative compressor operating points. The objective is to improve closed-loop constraint satisfaction in the presence of plant-model mismatch.

6.3.1 Closed-loop control results

In this section, we evaluate the closed-loop control performance with two parameter sets for the back-off parameters. The first set, denoted with NT for no tuning, sets all back-off parameters to zero, which corresponds to the parameterization of the nominal target selector. The closed-loop control results for the tracking performance and constraint satisfaction are given in Fig. 6.7 and Fig. 6.8, respectively. The results corresponding to the NT parameter set are depicted in

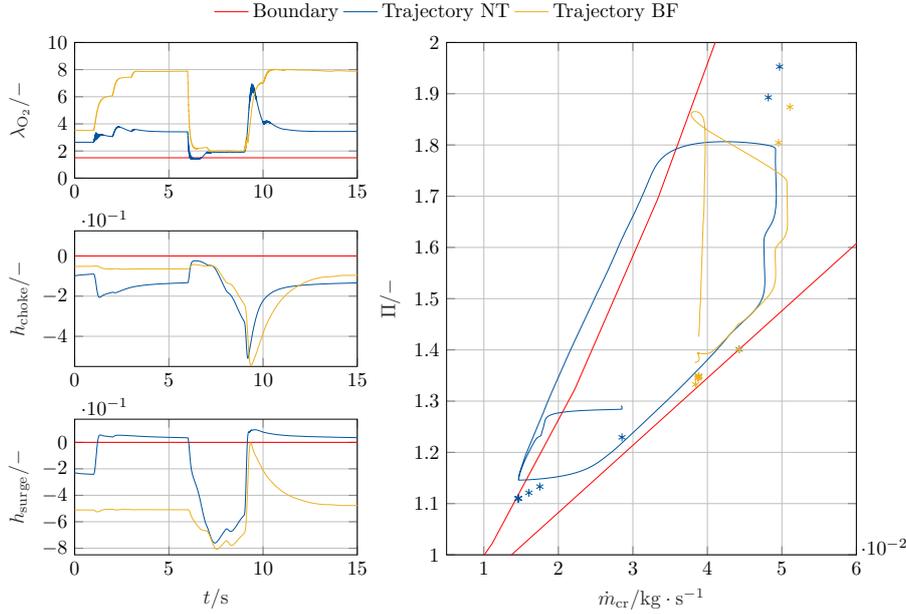


Fig. 6.7: Closed-loop constraint trajectories of control concept with plant model.

blue. In addition to the trajectories, we also plot the optimized operating points by the target selector in the compressor map. We see in Fig. 6.7 that both the surge and the OER constraint are violated. In particular, the surge constraint violation is major as the plant-model mismatch leads to a shift in the compressor trajectory to higher pressure ratios. This can be explained by the fact that humidity is considered in the plant but not in the prediction model. Consideration of humidity results in a partial pressure of vapour in the system which raises the pressure level which the controller is not aware of and actuates as if this extra pressure component is not present. The slight violation in the OER constraint also stems from a mismatch in the air-path dynamics, which cannot be eliminated by the state estimator which relies on the same prediction model as the NMPC. Given the constraint violations, the power tracking results for the NT back-off parameter set is slower compared to the results from the previous sections. This observation matches the $t_{90,P}$ values given in Table 6.3.

For the second parameter set, denoted as BF (brute force), the back-off parameters are manually adjusted until all constraints are satisfied in closed loop. This tuning procedure is purely heuristic and does not guarantee optimality.

Since the surge constraint is the most critical, the corresponding back-off parameter $p_{bo,surge}$ is iteratively increased until the surge constraint is respected. This adjustment also alleviates violations of the OER constraint without modifying the remaining back-off parameters. The resulting BF back-off parameters are:

$$p_{bo,surge} = 0.55$$

6.3 Simulative validation of the overall control concept

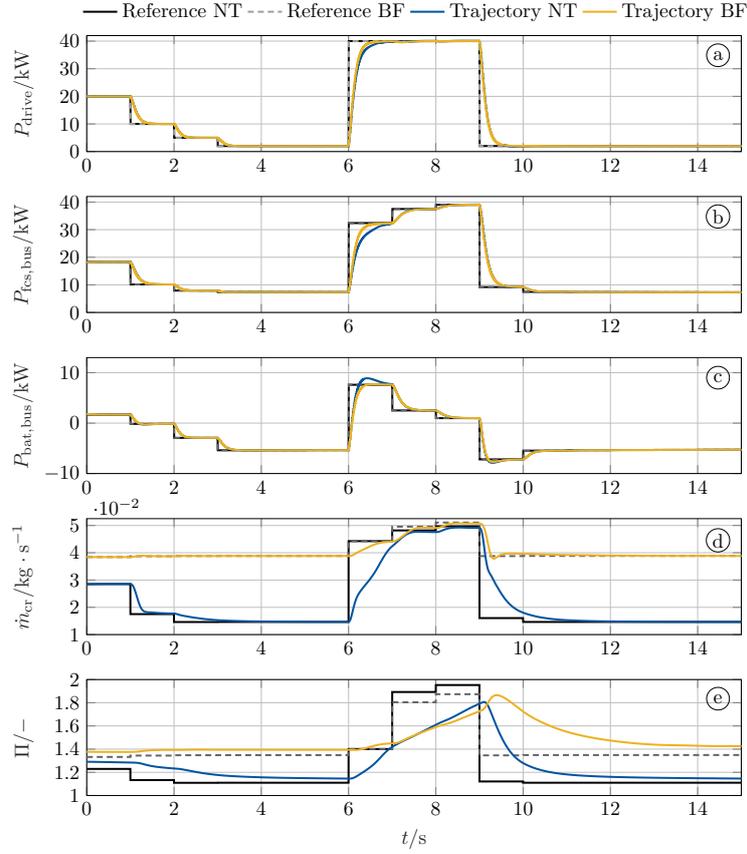


Fig. 6.8: Closed-loop constraint trajectories resulting from the overall control concept with plant model.

$$p_{bo,choke} = 0$$

$$p_{bo,over} = 0$$

The corresponding closed-loop control results are depicted in yellow. The trajectory in the compressor map reveals that the admissible area in the compressor is restricted considerably, so that the system is forced to higher mass flows, and also higher pressures. While this allows for more dynamic power tracking, see Fig. 6.8, it comes at the expense of efficiency. The observed efficiency degradation is primarily caused by the enforced surplus of air mass flow required to ensure constraint satisfaction. This operating regime shifts the compressor to higher pressure ratios, which significantly increases compressor power consumption and therefore reduces overall system efficiency. This is confirmed by Table 6.3, as both $t_{90,P}$ are better for the total power and FCS power but both the total efficiency and the FCS efficiency decrease by 2% which is considerable for such a short load cycle. The table further confirms that all constraints are satisfied in the BF parameter set while some constraint is violated for the NT parameter set.

Table 6.3: Performance metrics for evaluation of the overall control concept with plant model.

	$t_{90,drive}$	$t_{90,fcs}$	h_{max}	η_{fchev}	η_{fcs}
no tuning	0.34	0.48	0.0953	0.46	0.51
brute force	0.27	0.28	-0.0035	0.44	0.49

6.3.2 Closed loop control results with MOBO tuning

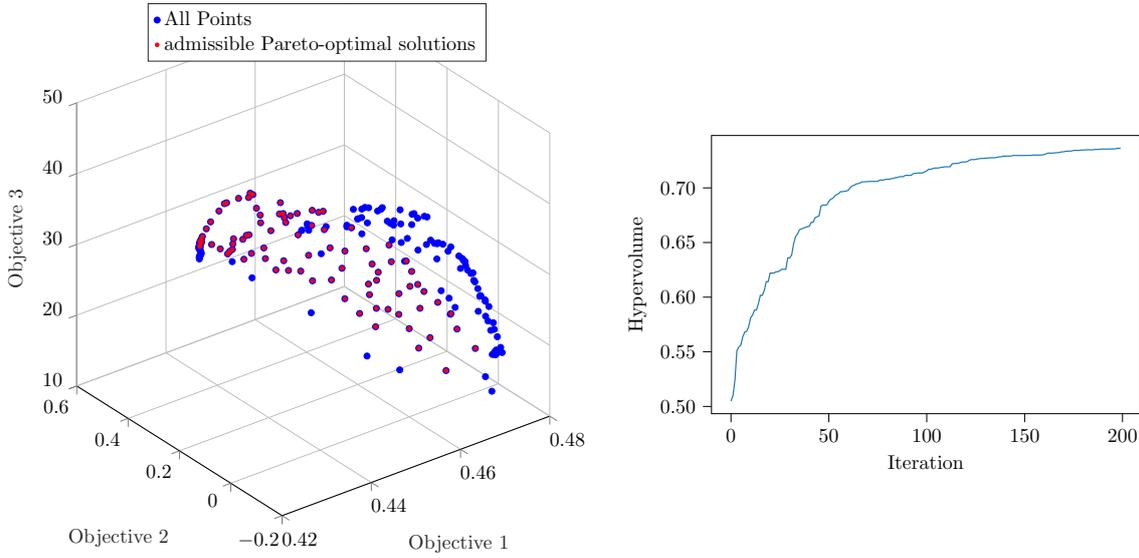


Fig. 6.9: Left: All observed points and observed Pareto front from MOBO optimization, right: Hypervolume evolution over MOBO iterations.

This section applies MOBO for tuning of some of the control design parameters of the power-split NMPC and two of the three back-off parameters. The MOBO optimization is conducted with the same step scenario which we use for evaluation. The parameter space is defined as follows:

$$10^{-2} \leq Q_{\eta_{fcs}} \leq 100 \quad (6.1)$$

$$10^{-2} \leq Q_{SoC} \leq 100 \quad (6.2)$$

$$10^{-2} \leq Q_{SoC,tf} \leq 100 \quad (6.3)$$

$$10^{-3} \leq R_{P_{fcs,bus}} \leq 1 \quad (6.4)$$

$$10^{-3} \leq R_{P_{bat,bus}} \leq 1 \quad (6.5)$$

$$0 \leq p_{bo,surge} \leq 0.54 \quad (6.6)$$

$$0 \leq p_{bo,oer} \leq 1 \quad (6.7)$$

The three objectives are the vehicle efficiency (objective 1), constraint satisfaction with respect to the surge and the OER constraint (objective 2), and the maximum delivered power by the FCS during the step cycle (objective 3). The reference point is chosen as $(0.42, -0.13, 6)$, such that constraint violations are explicitly permitted during the MOBO process. We run the MOBO for 200 iterations and initialize the algorithm with a Sobol sequence of 8 parameterizations.

Fig. 6.9 (right) shows that the hypervolume improvement becomes less after around 60 iterations. The sampled points in the objective space are depicted in Fig. 6.9 (left), with red points being admissible Pareto-optimal solutions among the observations made excluding those with a negative value for objective 2, as this indicates constraint violation.

From the admissible Pareto-optimal solutions, we select a parameter set that achieves

- (i) a maximum FCS power demand comparable to previous evaluations (approximately 32kW),
- (ii) high overall efficiency, and
- (iii) strict constraint satisfaction, indicated by $h_{\max} \leq 0$.

The closed-loop control trajectories for this parameter set are depicted in Fig. 6.10 and Fig. 6.11. The dynamic power tracking of the total power is satisfactory with $t_{90, P_{\text{drive}}} = 0.28 \text{ s} < 1$. The power tracking results for the FCS are now not entirely comparable because the references are now considerably different.

Notably, the resulting references from this parameterization and many other Pareto-optimal solutions from the MOBO algorithm have the highest FCS power demand occurring at the power step up. The FCS power reference then gradually decreases, which is the opposite in previous evaluations. This makes sense because the surge boundary is the most critical in closed-loop with the plant and reducing the power change during power step down helps to avoid operation close to the surge boundary. Fig. 6.11 shows that no constraint violations occur.

Table 6.4: Performance metrics for evaluation of the overall control concept with plant model and MOBO tuning.

	$t_{90, \text{drive}}$	$t_{90, \text{fcs}}$	h_{\max}	η_{fchev}	η_{fcs}
no tuning	0.34	0.48	0.0953	0.46	0.51
brute force	0.27	0.28	-0.0035	0.44	0.49
MOBO tuning	0.28	0.27	$-6.66 \cdot 10^{-4}$	0.46	0.51

Table 6.4 further reveals that despite no constraint violation and good total power dynamic tracking, the efficiency values are higher than brute force tuning results.

6 Simulative studies and results

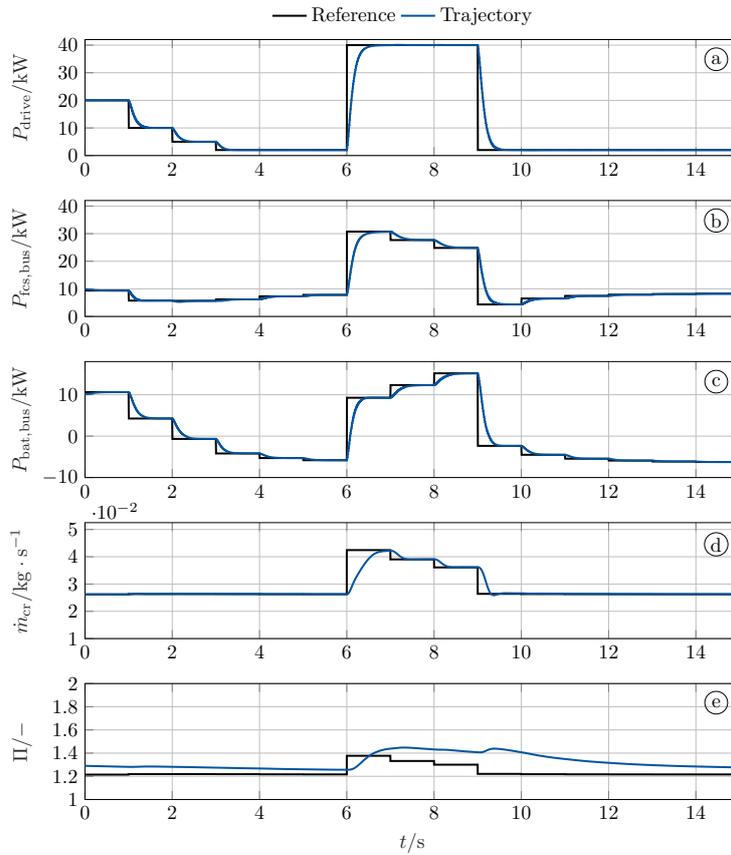


Fig. 6.10: Closed-loop output tracking results of the overall control concept with plant model and MOBO parameterization.

Research question 4:

How can the proposed hierarchical MPC scheme achieve an optimal trade-off between dynamic power delivery, constraint satisfaction, and efficiency?

The results demonstrate that the closed-loop performance of the proposed hierarchical MPC scheme is strongly influenced by the chosen control design parameters. In particular, the tuning of the high-level control layer, including the weighting of efficiency objectives and the selection of back-off parameters, plays a decisive role in balancing dynamic power delivery, constraint satisfaction, and efficiency.

While manual tuning of back-off parameters can enforce constraint satisfaction, it typically leads to overly conservative operation and unnecessary efficiency losses. In contrast, the application of multi-objective Bayesian optimization (MOBO) enables a systematic exploration of the trade-off space and yields parameterizations that simultaneously ensure constraint satisfaction, high efficiency, and satisfactory dynamic performance.

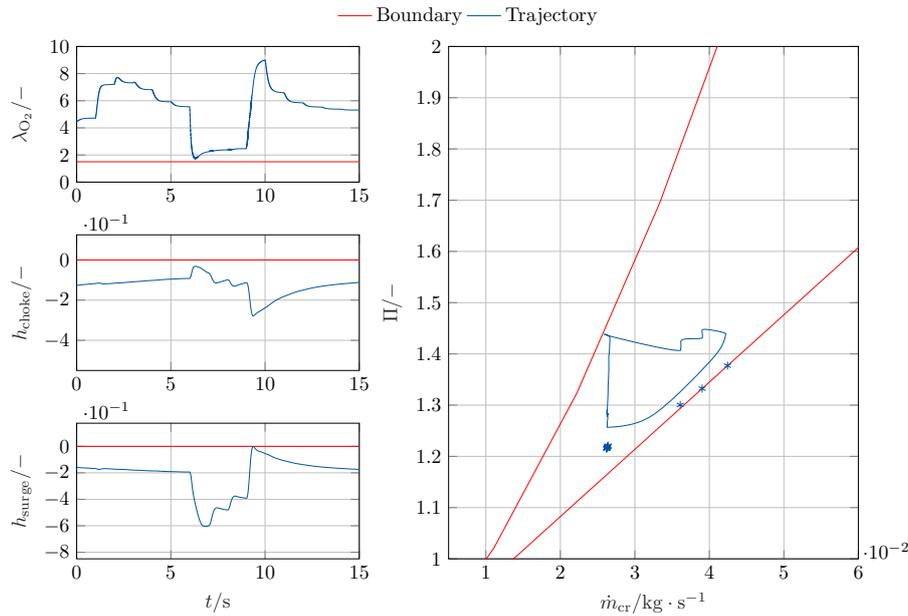


Fig. 6.11: Closed-loop constraint trajectories of control concept with plant model and MOBO parameterization.

In this work, MOBO is applied to tune the high-level control layer and selected back-off parameters. As long as the number of tuning parameters remains moderate (here fewer than 20), MOBO proves to be an effective and practical tool. A prerequisite for its successful application is a clear understanding of which control design parameters significantly affect the control objectives, which can be supported by sensitivity analysis. In our case, the choice of parameters is rather straightforward since the task of the tracking MPC is to track the references supplied to it and tuning its weights will just determine how well these are followed. The reference generation part done by the high-level control layer, however, considerably affects the behavior of the system.

6.4 Evaluation of humidity consideration

In this section, we present simulation results of the low-level tracking NMPC with explicit consideration of humidity, as described in Section 5.3, under ideal conditions. The controller is evaluated in closed loop with its prediction model, assuming full state measurement. Compared to the analysis in Section 6.2, this evaluation differs only by the inclusion of humidity effects in the prediction model and constraints.

The simulations are carried out in MATLAB with a sampling time of 8 ms and ERK4 as integration scheme. For the prediction horizon, we choose $N = 25$. With these parameters, the SQP

algorithm is implemented using a Gauss–Newton Hessian approximation, a step length of $\alpha = 0.1$, and a maximum of 13 QP iterations per SQP iteration. In the following we show the closed-loop simulation results for the step evaluation scenario, which are depicted in Fig. 6.12, Fig. 6.13 and Fig. 6.14 in blue. The corresponding algorithm is referred to as SQP NMPC GN.

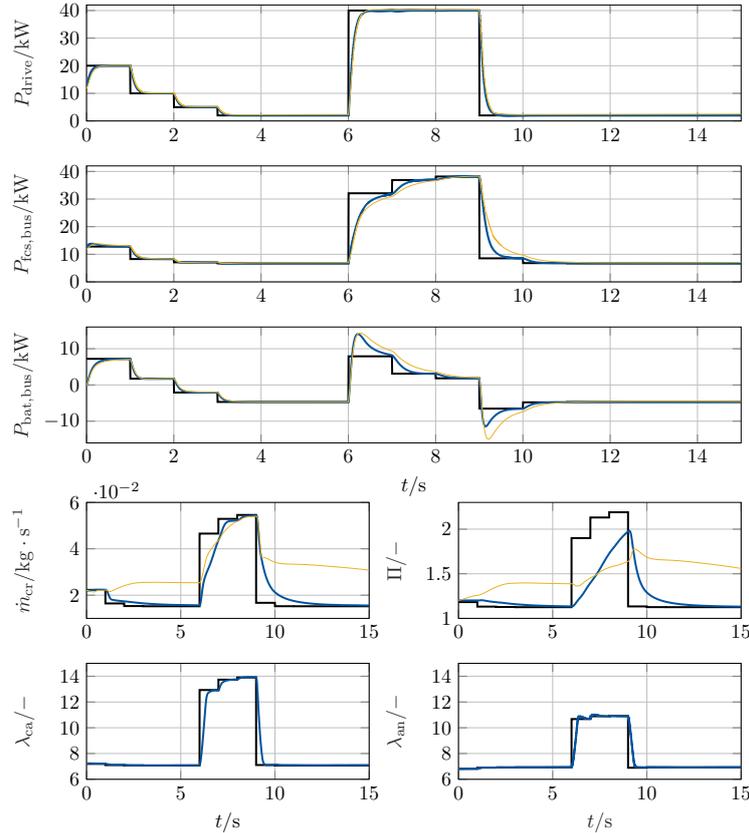


Fig. 6.12: Closed-loop output tracking results for low-level control with humidity consideration.

From the power tracking results we observe that incorporating humidity in the air path deteriorates tracking performance of the automotive FCS. In fact, the FCS power trajectory in Fig. 6.12b does not reach its commanded reference within one second after the power step up occurs at 6 s. This slower dynamic response is also reflected by the $t_{90,P_{fcs}}$ -value in Table 6.5. As a consequence, the battery system needs to compensate more compared to the results in Section 6.2 to ensure rapid total power delivery during the power step up.

We further notice that in comparison to the algorithm without humidity consideration in Section 6.2, the pre-optimized references for the compressor are larger with humidity consideration. That means, to achieve similar FCS power levels as in Fig. 6.5b, the target selector with humidity consideration optimizes higher compressor mass flows and pressure ratios, see Fig. 6.12d and Fig. 6.12e. This behavior is expected, since considering water vapor increases the total

gas mass and partial pressures in the air path, which requires higher compressor mass flows and pressure ratios to achieve comparable oxygen supply levels. However, these references are dynamically not reachable due to the slow pressure dynamics in the air path and the choke boundary limitation. We can see in the compressor map in Fig. 6.13 that the system operates very close to the choke boundary. This is during the power step up, where given the power dynamics of the air path the compressor mass flow is limited by the choke boundary. Similar observations can be made during power step down, in particular in terms of slow dynamic behavior of the FCS power, the compressor mass flow and the pressure ratio.

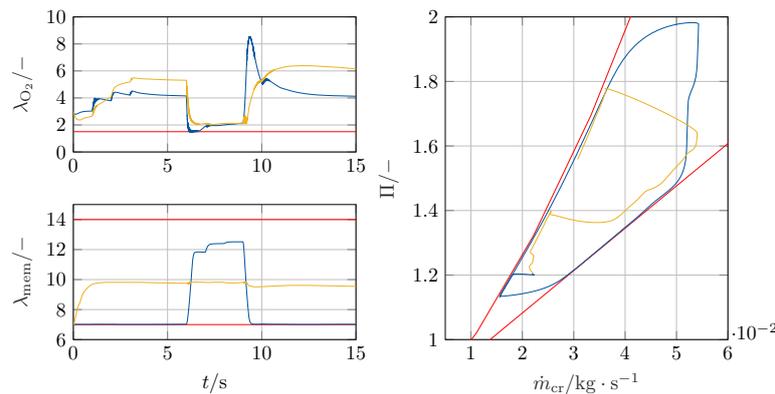


Fig. 6.13: Closed-loop constraint trajectories for low-level control with humidity consideration.

Overall, no constraint violations occur, see Fig. 6.13. This is confirmed by the negative h_{\max} -value in Table 6.5, where the constraint on the membrane humidity is also incorporated. At the same time, the chosen evaluation scenario is not ideally suited to demonstrate the controller's capability to prevent membrane drying, which predominantly occurs during prolonged low-power operation. This effect occurs mostly during low FCS power operation, which is not the case here. More specifically, we see that the membrane humidity is high during high FCS power generation due to the water produced in the reaction. Before and after the step in power demand, the membrane humidity is operated on its lower boundary because this is optimized by the references for the membrane water content on the anode and cathode side. We will see later that if the references on these values are not enforced, the membrane water content does not operate on its lower boundary.

Nevertheless, we see in Fig. 6.14 that when the membrane humidity operates on its lower boundary, the controller commands a relatively high relative humidity RH_{an} for the anode. It should be noted that modeling the anode relative humidity as a directly actuated control input is a simplification. In reality, changes in anode humidification are subject to actuator dynamics and transport delays, which are neglected in the present control-oriented model. The fact that the control algorithm demands high anode relative humidities throughout the evaluation scenario shows us that humidification of the anode is important for proper water management.

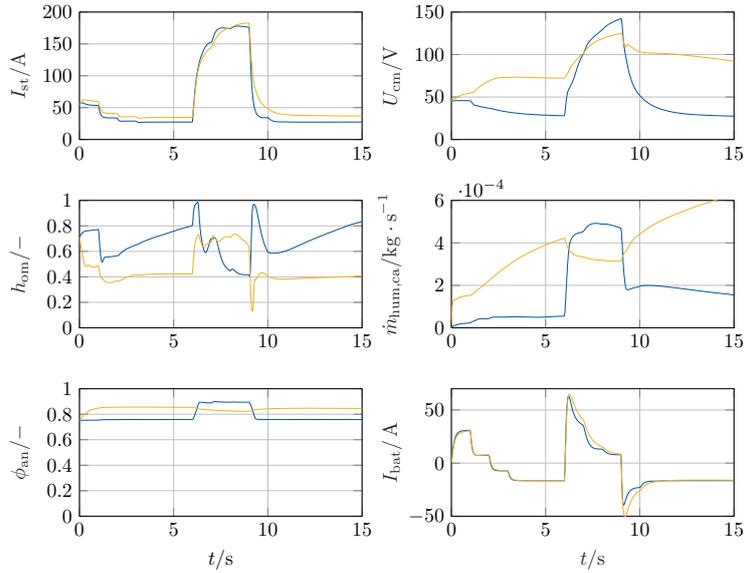


Fig. 6.14: Closed-loop input trajectories for low-level control with humidity consideration.

Table 6.5: Performance metrics for dynamic power delivery and constraint satisfaction, evaluated for the tracking NMPC with humidity consideration in the step evaluation scenario.

	$t_{90, P_{\text{drive}}}$	$t_{90, P_{\text{ics}}}$	h_{max}	t_{MLB}
SQP NMPC GN	0.21	0.472	$-7.37 \cdot 10^{-6}$	—
SQP NMPC GN RT	0.239	0.631	-0.0026	8 ms

Real-time feasibility

Based on the analysis in Chapter 5, it is already evident that achieving real-time feasibility with explicit humidity consideration is more challenging, as the system dimension increases by one additional state and stricter sampling time requirements arise. The latter can be relaxed, though we stick to 8 ms here.

As explained in Section 6.2, for testing the control algorithm on the embedded hardware, our workflow involves implementing the control algorithm in SIMULINK from where the code is generated for the MicroLabBox. When transferring the control algorithm from MATLAB to SIMULINK, the results should in principle stay the same, since we use the same solver with the same settings. However, we notice a difference which is most likely due to a difference of the HPIPM toolbox for MATLAB and SIMULINK. For the previous analysis in Section 6.2, the differences were negligible but more profound for the evaluation in this section in terms of a tracking offset that we did not see before. We notice that the problem is less profound when we choose to not track the water content in the cathode and anode. This does not interfere with

our additional control objective of primarily constraining the membrane humidity. Leaving out tracking of the two variables gives the controller more freedom at the cost of possible nonunique solutions and worse FCS efficiency. The results for a real-time capable controller with humidity consideration shown here are for the control problem adapted to track only 5 output variables, without Λ_{ca} and Λ_{an} .

After iterative manual tuning of the prediction horizon and the maximum number of allowed QP iterations, we found a real-time capable version of the controller with $N = 16$ and a maximum number of 5 QP iterations. The mean/max turnaround time for the step evaluation scenario is 7.7 ms/8 ms.

The results are depicted in yellow in the figures above and the performance metrics are given in Table 6.5 (SQP NMPC GN RT). As expected, the control performance deteriorates quite a bit as the settings only allow very suboptimal solutions, especially in terms of tracking of the compressor setpoint. We further notice that without tracking the water content in the cathode and anode, the membrane humidity does not operate on its boundary. In fact, the membrane is kept more stable which is beneficial for the membrane's health. However, this observation is only valid for this specific evaluation scenario.

Research question 5:

How does the consideration of humidity affect control performance?

The results show that neglecting humidity in the air path and FC modeling constitutes a major simplification that can significantly distort the predicted dynamic behavior of the system. In particular, the inclusion of humidity leads to noticeably slower dynamic power response due to increased gas mass and tighter compressor constraints. Therefore, humidity effects must be considered when assessing the dynamic capability of automotive FCSs. Furthermore, the results highlight the importance of anode humidification for proper water management and membrane health.

6.5 Summary of results

This chapter presented an extensive simulative evaluation of the proposed control concepts for automotive FCSs, with a particular focus on tracking performance, constraint satisfaction, real-time feasibility, and the impact of humidity consideration. The results were used to answer research questions 2 to 5.

Linearized vs. nonlinear prediction (Research question 2). The comparison between tracking LTV MPC and tracking NMPC showed that, although both approaches achieve comparable performance metrics in terms of rise time and overall constraint violation, relevant

qualitative differences emerge under highly dynamic operating conditions. While the LTV MPC is able to deliver rapid power responses, its linearized prediction may fail to accurately anticipate nonlinear constraint behavior, particularly with respect to compressor surge and choke boundaries. This can result in abrupt corrective control actions and transient power drops during aggressive power demands. In contrast, the SQP-based NMPC with nonlinear prediction provides more accurate constraint anticipation and smoother closed-loop behavior. Therefore, a purely linearized prediction is not sufficient to robustly meet the control requirements of automotive FCSs when highly dynamic operation is required.

Real-time feasibility of tracking NMPC (Research question 3). The real-time feasibility analysis demonstrated that an SQP-based NMPC can be implemented in a real-time capable manner on embedded hardware for automotive FC applications. By employing the RTI scheme in combination with a sparsity-exploiting QP solver, the controller achieves turnaround times below the sampling time of 10ms on a dSPACE MicroLabBox. Although small performance degradations compared to an offline NLP solution are observed, the resulting closed-loop behavior remains satisfactory in terms of dynamic power delivery and constraint satisfaction. These results confirm that an effective and real-time feasible NMPC algorithm for automotive FCSs can be realized under idealized conditions.

Validation of the hierarchical control concept (Research question 4). The simulative validation of the full hierarchical MPC framework in closed loop with a comprehensive plant model revealed that plant-model mismatch significantly affects constraint satisfaction and control performance. In particular, neglecting disturbances in the constraint formulation and ignoring humidity effects in the prediction model leads to violations of surge and OER constraints. The introduction of back-off parameters in the target selector enables conservative steady-state reference generation and restores constraint satisfaction in closed loop. Furthermore, MOBO proved to be an effective tool for systematically tuning high-level control parameters and back-off values. The results show that the proposed hierarchical MPC scheme can achieve a meaningful trade-off between dynamic power delivery, constraint satisfaction, and efficiency, provided that the control design parameters are carefully selected.

Impact of humidity consideration (Research question 5). The evaluation of humidity consideration in the low-level control layer revealed a noticeable impact on system dynamics and control performance. Including humidity in the air-path model leads to increased compressor pressure levels and mass flow requirements, which slow down the dynamic power response of the FCS. As a consequence, the FC power exhibits reduced dynamic capability, and the battery system must compensate more during transient operation. At the same time, the consideration of humidity improves physical consistency and enables explicit control of membrane water content. The results indicate that neglecting humidity constitutes a major simplification that may lead to overly optimistic assessments of dynamic capability in automotive FCSs.

Overall conclusion of Chapter 6. Overall, the simulative studies confirm that NMPC, combined with a hierarchical control structure and systematic parameter tuning, provides a viable and effective solution for the control of automotive FCSs. The results highlight the importance of nonlinear prediction, real-time capable optimization, and physically consistent modeling when addressing demanding dynamic operation and strict safety constraints.

7 Summary and outlook

This thesis presented the design, analysis, and evaluation of a hierarchical model predictive control (MPC) framework for fuel-cell-dominant hybrid electric vehicles. The proposed control architecture consists of a high-level control layer that determines the power split between the fuel cell system and the battery and generates optimal operating-point references. A low-level control layer directly actuates the vehicle and ensures dynamic tracking of these references while respecting operational and safety-critical constraints.

Achieving fast power delivery in automotive fuel cell systems is particularly challenging due to nonlinear air-path dynamics that become especially relevant during highly dynamic operation, as well as the presence of competing objectives, namely dynamic performance, constraint satisfaction, and efficiency. This work addressed these challenges from a control-theoretic perspective by investigating the application of nonlinear, real-time capable MPC algorithms, and from a system-level perspective by explicitly accounting for plant-model mismatch, efficiency optimization, and the interaction between the high-level and low-level control layers.

Summary of contributions The contributions of this thesis can be summarized along three main dimensions.

Low-level control design and real-time feasibility First, the design of a real-time-capable low-level tracking controller for automotive fuel cell systems was investigated. A physics-based nonlinear prediction model was employed to represent the dominant air-path dynamics that are relevant during highly dynamic operation. By comparing a tracking nonlinear MPC (NMPC) based on a sequential quadratic programming (SQP) scheme with a linear time-varying MPC relying on successive linearization, it was shown that nonlinear prediction is advantageous when rapid and substantial power changes are required. While both approaches achieved comparable performance in terms of dynamic power delivery metrics and overall constraint violation measures, linearized predictions were shown to inadequately anticipate constraint activation under highly dynamic conditions. This resulted in abrupt corrective control actions and undesired transient behavior during acceleration.

Furthermore, this thesis demonstrated that, with a careful algorithmic design, the proposed SQP-based NMPC can be implemented in real time on embedded automotive hardware (dSPACE MicroLabBox). The controller achieved fast dynamic power tracking with only minimal

violations of safety-critical constraints, even in a demanding step-load scenario. These results establish that nonlinear MPC is not only beneficial from a control-performance perspective but also feasible for real-time automotive applications when combined with structure-exploiting solvers and tailored problem formulations. These findings are specific to an air-path-focused control formulation of the automotive fuel cell system and should be interpreted within this scope.

Vehicle-level evaluation and efficiency optimization Second, the hierarchical control concept was evaluated at vehicle level in closed loop with a plant model of higher complexity than the controller-internal model, thereby intentionally introducing plant-model mismatch. While state and disturbance estimation were employed within the control architecture, they were used exclusively to adapt the references tracked by the low-level controller and did not modify the high-level optimization. As a result, plant-model mismatch was shown to affect constraint satisfaction during highly dynamic operation, even when steady-state tracking performance remained acceptable.

To mitigate this effect, back-off parameters were introduced in the steady-state optimization to conservatively tighten critical constraints. This approach provides a practical means of improving closed-loop robustness in the presence of modeling uncertainties and directly motivated the subsequent systematic tuning of high-level control parameters.

To systematically explore the trade-off between dynamic performance, constraint satisfaction, and efficiency, this work applied multi-objective Bayesian optimization (MOBO) to tune selected high-level control parameters. The results showed that MOBO is a suitable tool for closed-loop controller tuning when the number of optimization variables remains moderate. For the considered evaluation scenario, efficiency improvements of approximately 2% were achieved compared to a baseline parameterization obtained without systematic tuning, while maintaining constraint satisfaction, highlighting the significant influence of high-level control design parameters on overall system behavior.

Humidity-aware modeling and control Third, the control-oriented fuel cell model was extended to explicitly account for humidity dynamics in the air path and membrane. The results revealed that neglecting humidity constitutes a major simplification that can lead to overly optimistic assessments of dynamic performance. When humidity was considered, dynamic power delivery was noticeably slower, since rapid power increases demand high air mass flows to maintain sufficient oxygen excess ratio, while the comparatively slow pressure dynamics restrict the achievable mass flow during transients, causing operation to be limited by the compressor choke boundary. This demonstrates that humidity effects must be included to realistically evaluate and design control strategies for automotive fuel cell systems, particularly when dynamic capability requirements are of interest.

Future work

The results of this thesis open several directions for future research. From a control perspective, further improvements could be achieved by incorporating available power demand predictions into the MPC framework. Such predictions would enable anticipatory control actions, potentially reducing constraint activation and improving dynamic performance, especially in highly transient driving scenarios.

In terms of computational efficiency, alternative problem formulations and advanced condensing or partial-condensing techniques could further reduce the computational burden of nonlinear MPC, enabling longer prediction horizons or additional system states to be considered within the same real-time constraints. The controller designs and real-time feasibility analyses presented in this work can serve as a benchmark for such investigations.

At system level, future work should focus on the development and validation of control-oriented models that capture additional system-level effects beyond the air path, including thermal dynamics, actuator limitations, and degradation phenomena. Such models are essential for transferring advanced control concepts from simulation to experimental test benches and real vehicles.

Finally, the estimation of internal fuel cell states, particularly humidity-related quantities, represents a critical extension of the proposed control framework. Reliable estimation of membrane and gas-phase humidity would enable more robust water management strategies and facilitate the application of the proposed control concepts to operating conditions that are critical for membrane drying and long-term durability. Evaluating the humidity-aware control strategy across a broader range of driving scenarios therefore constitutes an important step toward practical deployment.

A Appendix

A.1 Equivalent hydrogen consumption

They masses in 2.24 can be calculated as follows:

$$\dot{m}_{\text{H}_2,\text{fcs}} = \frac{M_{\text{H}_2} \cdot n \cdot I_{\text{st}}}{2 \cdot F} \quad (\text{A.1})$$

$$m_{\text{H}_2,\text{fcs}} = \int \dot{m}_{\text{H}_2,\text{fcs}} dt \quad (\text{A.2})$$

$$\dot{m}_{\text{H}_2,\text{bat}} = P_{\text{bat},\text{net}} \cdot \delta \cdot \frac{\dot{m}_{\text{H}_2,\text{fcs},\text{avg}}}{P_{\text{fcs},\text{net},\text{avg}}} \quad (\text{A.3})$$

$$m_{\text{H}_2,\text{bat}} = \int \dot{m}_{\text{H}_2,\text{bat}} dt \quad (\text{A.4})$$

where $\dot{m}_{\text{H}_2,\text{fcs}}$ and $\dot{m}_{\text{H}_2,\text{fcs},\text{avg}}$ are the FCS hydrogen consumption rate and average consumption rate, respectively, $P_{\text{fcs},\text{net},\text{avg}}$ is the net FCS average output power, M_{H_2} is the molar mass of hydrogen, \dot{m}_{bat} is the equivalent hydrogen consumption by the battery and δ can be calculated using (A.5):

$$\delta = \begin{cases} \frac{1}{\eta_{\text{cha},\text{avg}} \cdot \eta_{\text{dis}}} & P_{\text{bat},\text{net}} \geq 0 \\ \eta_{\text{dis},\text{avg}} \cdot \eta_{\text{cha}} & P_{\text{bat},\text{net}} < 0 \end{cases} \quad (\text{A.5})$$

With

$$\eta_{\text{dis}} = \frac{1 - \sqrt{1 - 4 \cdot R_{\text{dis}} \cdot P_{\text{bat},\text{net}} / U_{\text{ocv}}^2}}{2} \quad (\text{A.6})$$

$$\eta_{\text{cha}} = \frac{2}{1 - \sqrt{1 - 4 \cdot R_{\text{cha}} \cdot P_{\text{bat},\text{net}} / U_{\text{ocv}}^2}} \quad (\text{A.7})$$

where η_{dis} and η_{cha} is the discharging/charging efficiency of the battery, $\eta_{\text{dis},\text{avg}}$ and $\eta_{\text{cha},\text{avg}}$ are the average discharging/charging efficiency of the battery, R_{dis} and R_{cha} is the total discharging/charging resistance of the battery, respectively. U_{ocv} is the OCV of the battery. The correction coefficient denoted by k can be obtained by:

$$k = 1 - 2 \cdot \mu \cdot \frac{x_{\text{SoC}} - 1/2 \cdot (x_{\text{SoC},\text{max}} + x_{\text{SoC},\text{min}})}{x_{\text{SoC},\text{max}} + x_{\text{SoC},\text{min}}} \quad (\text{A.8})$$

where μ is the balance factor during the cycle, x_{SoC} is the state of charge of the battery and $x_{\text{SoC},\text{min}/\text{max}}$ denote its minimum/maximum, respectively.

A.2 Coefficients for stack voltage model

The coefficients in equations (4.11), (4.12) and (4.13) are determined using nonlinear regression with polarization data from an automotive propulsion-sized PEMFC stack [83]. By assuming that the data is obtained from the FC stack operating under a well-controlled environment, where cathode gas is fully humidified and OER (ratio of oxygen supplied to oxygen reacted) is regulated at 2, the pressure terms in the activation and concentration overvoltage terms can be related to oxygen partial pressure, p_{O_2} , and vapor saturation pressure, p_{sat} . The regression results are

$$E_0 = 0.279 - 8.5 \cdot 10^{-4} \cdot (T_{\text{fc}} - 298.15) + 4.3085 \cdot 10^{-5} \cdot T_{\text{fc}} \cdot \left[\ln \left(\frac{p_{\text{ca}} - p_{\text{sat}}}{1.01325} \right) + \frac{1}{2} \ln \left(\frac{0.1173 \cdot (p_{\text{ca}} - p_{\text{sat}})}{1.01325} \right) \right] \quad (\text{A.9})$$

$$E_a = \left(-1.618 \cdot 10^{-5} \cdot T_{\text{fc}} + 1.618 \cdot 10^{-2} \right) \cdot \left(\frac{p_{O_2}}{0.1173} + p_{\text{sat}} \right)^2 + \left(1.8 \cdot 10^{-4} \cdot T_{\text{fc}} - 0.166 \right) \cdot \left(\frac{p_{O_2}}{0.1173} + p_{\text{sat}} \right) + \left(-5.8 \cdot 10^{-4} T_{\text{fc}} + 0.5736 \right) \quad (\text{A.10})$$

$$c_1 = 10 \quad (\text{A.11})$$

$$b_{11} = 0.05139 \quad (\text{A.12})$$

$$b_{12} = 0.00326 \quad (\text{A.13})$$

$$b_2 = 350 \quad (\text{A.14})$$

$$c_2 = \begin{cases} \left(7.16 \cdot 10^{-4} \cdot T_{\text{fc}} - 0.622 \right) \cdot \left(\frac{p_{O_2}}{0.1173} + p_{\text{sat}} \right) + \left(-1.45 \cdot 10^{-3} \cdot T_{\text{fc}} + 1.68 \right), & \text{if } \frac{p_{O_2}}{0.1173} + p_{\text{sat}} < 2 \text{ atm} \\ \left(8.66 \cdot 10^{-4} \cdot T_{\text{fc}} - 0.068 \right) \cdot \left(\frac{p_{O_2}}{0.1173} + p_{\text{sat}} \right) + \left(-1.60 \cdot 10^{-4} \cdot T_{\text{fc}} + 0.54 \right), & \text{else} \end{cases} \quad (\text{A.15})$$

$$c_3 = 2 \quad (\text{A.16})$$

$$i_{\text{fc,max}} = 2.2 \quad (\text{A.17})$$

A.3 Approximation of square root function

This section presents the MATLAB function `sqrt`, which is designed to be a NaN-safe square root function. The function behaves like the standard square root for non-negative inputs and provides a smooth approximation for negative inputs, asymptotically converging to zero.

Mathematical Formulation

The `sqrt` function is formulated to ensure stability and smoothness. The key steps are as follows:

Define a small positive constant a to avoid division by zero and to stabilize the computation:

$$a = 1 \times 10^{-6} \tag{A.18}$$

Compute a smooth approximation of the absolute value of x :

$$x_{\text{abs_eps}} = \sqrt{x^2 + a} \tag{A.19}$$

Calculate the mean of x and its smooth absolute value:

$$x_{\text{mean}} = \frac{x_{\text{abs_eps}} + x}{2} \tag{A.20}$$

Compute the function output y as the square root of the mean:

$$y = \sqrt{x_{\text{mean}}} \tag{A.21}$$

Derive the derivative of y with respect to x :

$$\frac{dy}{dx} = \frac{1}{\sqrt{8(x_{\text{abs_eps}} + x)}} \left(\frac{x}{x_{\text{abs_eps}}} + 1 \right) \tag{A.22}$$

A.4 Modeling parameters

Table A.1: Thermodynamic constants used for modelling the FCS

Parameter	Symbol	SI unit	Value
Atmospheric pressure	p_{atm}	Pa	1.01325×10^5
Atmospheric temperature	T_{atm}	K	298.15
Ratio of specific heats of air	γ	–	1.4
Constant pressure specific heat of air	C_p	$\text{Jkg}^{-1} \text{K}^{-1}$	1004
Air density	ρ_a	kg m^{-3}	1.23
Universal gas constant	R	$\text{J mol}^{-1} \text{K}^{-1}$	8.3145
Air gas constant	R_a	$\text{Jkg}^{-1} \text{K}^{-1}$	286.9
Oxygen gas constant	R_{O_2}	$\text{Jkg}^{-1} \text{K}^{-1}$	259.8
Nitrogen gas constant	R_{N_2}	$\text{Jkg}^{-1} \text{K}^{-1}$	296.8
Vapor gas constant	R_v	$\text{Jkg}^{-1} \text{K}^{-1}$	461.5
Hydrogen gas constant	R_{H_2}	$\text{Jkg}^{-1} \text{K}^{-1}$	4124.3
Molar mass of air	M_a	kg mol^{-1}	28.97×10^{-3}
Molar mass of oxygen	M_{O_2}	kg mol^{-1}	32.00×10^{-3}
Molar mass of nitrogen	M_{N_2}	kg mol^{-1}	28.00×10^{-3}
Molar mass of vapor	M_v	kg mol^{-1}	18.02×10^{-3}
Molar mass of hydrogen	M_{H_2}	kg mol^{-1}	2.016×10^{-3}
Faraday constant	F	s A mol^{-1}	96485
Temperature of fuel cell	T_{fc}	K	353.15

Table A.2: Parameters used for modelling the FCS

Parameter	Symbol	SI unit	Value
Motor constant	k_t	N m A^{-1}	0.0153
Motor constant	R_{cm}	Ω	0.82
Motor constant	k_v	$\text{V rad}^{-1} \text{s}$	0.0153
Compressor efficiency	η_{cp}	—	0.8
Compressor motor mechanical efficiency	η_{cm}	—	0.98
Number of cells in fuel cell stack	n	—	381
Fuel cell active area	A_{fc}	m^2	280×10^{-4}
Inlet manifold volume	V_{im}	m^3	0.02
Cathode volume	V_{ca}	m^3	0.01
Anode volume	V_{an}	m^3	0.005
Outlet manifold volume	V_{om}	m^3	0.005
Inlet manifold outlet orifice constant	$k_{\text{im,out}}$	$\text{kg s}^{-1} \text{Pa}^{-1}$	0.3629×10^{-5}
Cathode outlet orifice constant	$k_{\text{ca,out}}$	$\text{kg s}^{-1} \text{Pa}^{-1}$	0.2177×10^{-5}
Membrane dry density	$\rho_{\text{mem,dry}}$	kg m^{-3}	2×10^3
Membrane dry equivalent weight	$M_{\text{mem,dry}}$	kg mol^{-1}	1.1
Membrane thickness	t_{mem}	m	1.275×10^{-4}
Compressor diameter	d_{cp}	m	0.2286
Compressor and motor inertia	J_{cp}	kg m^2	5×10^{-5}
Outlet manifold throttle discharge coefficient	C_D	—	0.0124
Outlet manifold throttle area	A_T	m^2	0.002
Average ambient air relative humidity	Φ_{atm}	—	0.5
Oxygen mole fraction at cathode inlet	$x_{\text{O}_2,\text{in}}$	—	0.21
Hydrogen mole fraction at anode inlet	$x_{\text{H}_2,\text{in}}$	—	1.0

Bibliography

- [1] ABDOLSHAH, M. ; SHILTON, A. ; RANA, S. ; GUPTA, S. ; VENKATESH, S.: Multi-objective Bayesian optimisation with preferences over objectives. In: *Advances in Neural Information Processing Systems*. Vol. 32. 2019.
- [2] ADAMS, J. A. ; YANG, W.-c. ; OGLESBY, K. A. ; OSBORNE, K. D.: The Development of Ford's P2000 Fuel Cell Vehicle. In: *SAE transactions* 109 (2000), pp. 1634–1645. DOI: 10.4271/2000-01-1061.
- [3] ALBERSMEYER, J. ; DIEHL, M. M.: The lifted Newton method and its application in optimization. In: *SIAM Journal on Optimization* 20. no. 3 (2010), pp. 1655–1684. DOI: 10.1137/080724885.
- [4] AMPHLETT, J. C. ; BAUMERT, R. M. ; MANN, R. F. ; PEPPLEY, B. A. ; ROBERGE, P. R. ; HARRIS, T. J.: Performance modeling of the Ballard Mark IV solid polymer electrolyte fuel cell: I. Mechanistic model development. In: *Journal of the Electrochemical Society* 142. no. 1 (1995), p. 1. DOI: 10.1149/1.2043866.
- [5] ANDERSSON, J. A. E. ; GILLIS, J. ; HORN, G. ; RAWLINGS, J. B. ; DIEHL, M.: CasADi: a software framework for nonlinear optimization and optimal control. In: *Mathematical Programming Computation* 11 (2019), pp. 1–36. DOI: 10.1007/s12532-018-0139-4.
- [6] AXEHILL, D.: Controlling the level of sparsity in MPC. In: *Systems & Control Letters* 76 (2015), pp. 1–7. DOI: 10.1016/j.sysconle.2014.12.002.
- [7] BACHER-CHONG, E. ; AYUBIRAD, M. A. ; QIU, Z. ; WANG, H. ; GOSHTASBI, A. ; OSSAREH, H. R.: Efficiency-aware and constraint-aware control of PEMFC air-path using a reference governor and MIMO internal model controller. In: *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 3050–3057. DOI: 10.23919/ACC53348.2022.9867418.
- [8] BACHER-CHONG, E. ; AYUBIRAD, M. A. ; QIU, Z. ; WANG, H. ; GOSHTASBI, A. ; OSSAREH, H. R.: Hierarchical fuel-cell airpath control: An efficiency-aware MIMO control approach combined with a novel constraint-enforcing reference governor. In: *IEEE Transactions on Control Systems Technology* 32. no. 2 (2023), pp. 534–549. DOI: 10.1109/TCST.2023.3329909.
- [9] BELLANTONI, J. F. ; DODGE, K. W.: A square root formulation of the Kalman-Schmidt filter. In: *AIAA Journal* 5. no. 7 (1967), pp. 1309–1314. DOI: 10.2514/3.4189.

- [10] BERTSEKAS, D. P.: *Dynamic programming and optimal control: Volume I*. 4th edition. Athena Scientific, 2012.
- [11] BERTSEKAS, D. P.: *Nonlinear programming*. 3rd edition. Athena Scientific, 2016.
- [12] BETTS, J. T.: *Practical methods for optimal control and estimation using nonlinear programming*. 2nd edition. Society for Industrial and Applied Mathematics, 2010. DOI: 10.1137/1.9780898718577.
- [13] BOCK, H. G.: Recent advances in parameter identification techniques for ODE. In: *Numerical treatment of inverse problems in differential and integral equations*. Vol. 2. Progress in Scientific Computing. Birkhäuser Boston, 1983, pp. 95–121. DOI: 10.1007/978-1-4684-7324-7{\textunderscore}7.
- [14] BOCK, H. G. ; PLITT, K. J.: A multiple shooting algorithm for direct solution of optimal control problems. In: *IFAC Proceedings Volumes* 17. no. 2 (1984), pp. 1603–1608. DOI: 10.1016/s1474-6670(17)61205-9.
- [15] BOCK, H. G.: Numerical solution of nonlinear multipoint boundary value problems with applications to optimal control. In: *Zeitschrift für Angewandte Mathematik und Mechanik* 58. no. 7 (1978), pp. 407–454.
- [16] BORRELLI, F. ; BEMPORAD, A. ; MORARI, M.: *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.
- [17] BORUP, R. ; MEYERS, J. ; PIVOVAR, B. ; KIM, Y. S. ; MUKUNDAN, R. ; GARLAND, N. ; MYERS, D. ; WILSON, M. ; GARZON, F. ; WOOD, D. ; ZELENAY, P. ; MORE, K. ; STROH, K. ; ZAWODZINSKI, T. ; BONCELLA, J. ; MCGRATH, J. E. ; INABA, M. ; MIYATAKE, K. ; HORI, M. ; OTA, K. ; OGUMI, Z. ; MIYATA, S. ; NISHIKATA, A. ; SIROMA, Z. ; UCHIMOTO, Y. ; YASUDA, K. ; KIMIJIMA, K.-i. ; IWASHITA, N.: Scientific aspects of polymer electrolyte fuel cell durability and degradation. In: *Chemical Reviews* 107. no. 10 (2007), pp. 3904–3951. DOI: 10.1021/cr050182l.
- [18] BREITINGER, J.: *Water management in highly dynamic PEM fuel cell system operation*. PhD thesis. Gottfried Wilhelm Leibniz Universität Hannover, 2024. DOI: 10.15488/16364.
- [19] BROCHU, E. ; CORA, V. M. ; FREITAS, N. de: A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. In: *arXiv preprint arXiv:1012.2599* (2010).
- [20] BRYSON, A. E. ; HO, Y.-C.: *Applied optimal control: Optimization, estimation and control*. Taylor & Francis Group, 1975. DOI: 10.1201/9781315137667.
- [21] BÜCHI, F. N. ; SRINIVASAN, S.: Operating proton exchange membrane fuel cells without external humidification of the reactant gases: Fundamental aspects. In: *Journal of The Electrochemical Society* 144. no. 8 (1997), pp. 2767–2772. DOI: 10.1149/1.1837893.

- [22] BUNDESAMT FÜR WIRTSCHAFT UND AUSFUHRKONTROLLE: *Elektromobilität (bis 18.12.2023) - Erhöhter Umweltbonus für E-Autos*. 2020.
- [23] CHEN, C. ; FULLER, T. F.: The effect of humidity on the degradation of Nafion® membrane. In: *Polymer Degradation and Stability* 94. no. 9 (2009), pp. 1436–1447. DOI: 10.1016/j.polyimdegradstab.2009.05.016.
- [24] CONN, A. R. ; GOULD, N. I. M. ; TOINT, P. L.: *Trust region methods*. Society for Industrial and Applied Mathematics, 2000. DOI: 10.1137/1.9780898719857.
- [25] CUI, J. ; CHEN, Q. ; LI, X. ; ZHANG, S.: Recent advances in non-precious metal electrocatalysts for oxygen reduction in acidic media and PEMFCs: an activity, stability and mechanism study. In: *Green Chemistry* 23. no. 18 (2021), pp. 6898–6925. DOI: 10.1039/D1GC01040A.
- [26] CZESLIK, C. ; SEEMANN, H. ; WINTER, R.: *Basiswissen Physikalische Chemie*. Springer-Verlag, 2010.
- [27] DAULTON, S. ; ERIKSSON, D. ; BALANDAT, M. ; BAKSHY, E.: Multi-objective Bayesian optimization over high-dimensional search spaces. In: *Uncertainty in Artificial Intelligence* 180 (2022), pp. 507–517.
- [28] DEB, K. ; PRATAP, A. ; AGARWAL, S. ; MEYARIVAN, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. In: *IEEE Transactions on Evolutionary Computation* 6. no. 2 (2002), pp. 182–197. DOI: 10.1109/4235.996017.
- [29] DESHWAL, A. ; SIMON, C. M. ; DOPPA, J. R.: Bayesian optimization of nanoporous materials. In: *Molecular Systems Design & Engineering* 6. no. 12 (2021), pp. 1066–1086. DOI: 10.1039/D1ME00093D.
- [30] DICKS, A. L. ; RAND, D. A. J.: *Fuel cell systems explained*. John Wiley & Sons, 2018.
- [31] DIEHL, M. ; BOCK, H. G. ; DIEDAM, H. ; WIEBER, P.-B.: Fast direct multiple shooting algorithms for optimal robot control. In: *Fast Motions in Biomechanics and Robotics*. Vol. 340. Springer Berlin Heidelberg, 2006, pp. 65–93. DOI: 10.1007/978-3-540-36119-0{\textunderscore}4.
- [32] DIEHL, M. ; BOCK, H. ; SCHLÖDER, J. P. ; FINDEISEN, R. ; NAGY, Z. ; ALLGÖWER, F.: Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations. In: *Journal of Process Control* 12. no. 4 (2002), pp. 577–585. DOI: 10.1016/s0959-1524(01)00023-3.
- [33] DIEHL, M. ; FERREAU, H. J. ; HAVERBEKE, N.: Efficient numerical methods for nonlinear MPC and moving horizon estimation. In: *Nonlinear Model Predictive Control*. Vol. 384. Springer Berlin Heidelberg, 2009, pp. 391–417. DOI: 10.1007/978-3-642-01094-1{\textunderscore}32.

- [34] DIEHL, M. ; USLU, I. ; FINDEISEN, R. ; SCHWARZKOPF, S. ; ALLGÖWER, F. ; BOCK, H. G. ; BÜRNER, T. ; GILLES, E. D. ; KIENLE, A. ; SCHLÖDER, J. P. ; STEIN, E.: Real-time optimization for large scale processes: Nonlinear model predictive control of a high purity distillation column. In: *Online Optimization of Large Scale Systems*. Springer Berlin Heidelberg, 2001, pp. 363–383. DOI: 10.1007/978-3-662-04331-8{\textunderscore}20.
- [35] FERREAU, H. J. ; KIRCHES, C. ; POTSCSKA, A. ; BOCK, H. G. ; DIEHL, M.: qpOASES: a parametric active-set algorithm for quadratic programming. In: *Mathematical Programming Computation* 6. no. 4 (2014), pp. 327–363. DOI: 10.1007/s12532-014-0071-1.
- [36] FORSGREN, A. ; GILL, P. E. ; WRIGHT, M. H.: Interior methods for nonlinear optimization. In: *SIAM Review* 44. no. 4 (2002), pp. 525–597. DOI: 10.1137/s0036144502414942.
- [37] FRISON, G. ; DIEHL, M.: HPIPM: a high-performance quadratic programming framework for model predictive control. In: *IFAC-PapersOnLine* 53. no. 2 (2020), pp. 6563–6569. DOI: 10.1016/j.ifacol.2020.12.073.
- [38] FRISON, G. ; KOUZOUPIS, D. ; SARTOR, T. ; ZANELLI, A. ; DIEHL, M.: Basic linear algebra subroutines for embedded optimization. In: *ACM Transactions on Mathematical Software (TOMS)* 44. no. 4 (2018), pp. 1–30. DOI: 10.1145/3210754.
- [39] FRISON, G. ; SORENSEN, H. H. B. ; DAMMANN, B. ; JORGENSEN, J. B.: High-performance small-scale solvers for linear model predictive control. In: *2014 European Control Conference (ECC)*. IEEE, 2014, pp. 128–133. DOI: 10.1109/ecc.2014.6862490.
- [40] GAUDRIE, D.: *High-dimensional Bayesian multi-objective optimization*. PhD thesis. Université de Lyon, 2019.
- [41] GERARD, M. ; POIROT-CROUVEZIER, J.-P. ; HISSEL, D. ; PERA, M.-C.: Oxygen starvation analysis during air feeding faults in PEMFC. In: *International Journal of Hydrogen Energy* 35. no. 22 (2010), pp. 12295–12307. DOI: 10.1016/j.ijhydene.2010.08.028.
- [42] GÖRKE, D., ed.: *Untersuchungen Zur Kraftstoffoptimalen Betriebsweise Von Parallelhybridfahrzeugen und Darauf Basierende Auslegung Regelbasierter Betriebsstrategien*. Wissenschaftliche Reihe Fahrzeugtechnik Universität Stuttgart Ser. Springer Fachmedien Wiesbaden GmbH, 2016.
- [43] GOSHTASBI, A.: *Modeling, parameter identification, and degradation-conscious control of polymer electrolyte membrane (PEM) fuel cells*. PhD thesis. University of Michigan, 2019.
- [44] GOSHTASBI, A. ; ERSAL, T.: LQ-MPC design for degradation-conscious control of PEM fuel cells. In: *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1555–1560. DOI: 10.23919/acc.2019.8814417.

- [45] GOSHTASBI, A. ; ERSAL, T.: Degradation-conscious control for enhanced lifetime of automotive polymer electrolyte membrane fuel cells. In: *Journal of Power Sources* 457 (2020), p. 227996. DOI: 10.1016/j.jpowsour.2020.227996.
- [46] GOSHTASBI, A. ; PENCE, B. L. ; CHEN, J. ; DEBOLT, M. A. ; WANG, C. ; WALDECKER, J. R. ; HIRANO, S. ; ERSAL, T.: A mathematical model toward real-time monitoring of automotive PEM fuel cells. In: *Journal of The Electrochemical Society* 167. no. 2 (2020), p. 024518. DOI: 10.1149/1945-7111/ab6dd1.
- [47] GROS, S. ; ZANON, M. ; QUIRYNEN, R. ; BEMPORAD, A. ; DIEHL, M.: From linear to nonlinear MPC: bridging the gap via the real-time iteration. In: *International Journal of Control* 93. no. 1 (2016), pp. 62–80. DOI: 10.1080/00207179.2016.1222553.
- [48] GUZZELLA, L. ; SCIARRETTA, A., eds.: *Vehicle propulsion systems: Introduction to modeling and optimization*. 3rd edition. Springer Berlin Heidelberg, 2013. DOI: 10.1007/978-3-642-35913-2.
- [49] HAHN, S.: *Modellbasierte Betriebsstrategie für PEM-Brennstoffzellensysteme*. Springer Fachmedien Wiesbaden, 2023. DOI: 10.1007/978-3-658-42086-4.
- [50] HAHN, S. ; BRAUN, J. ; KEMMER, H. ; REUSS, H.-C.: Adaptive operation strategy of a polymer electrolyte membrane fuel cell air system based on model predictive control. In: *International Journal of Hydrogen Energy* 46. no. 33 (2021), pp. 17306–17321. DOI: 10.1016/j.ijhydene.2021.02.135.
- [51] HÄHNEL, C. ; AUL, V. ; HORN, J.: Error handling approach of a PEM fuel cell system by nonlinear model predictive control. In: *2015 23rd Mediterranean Conference on Control and Automation (MED)*. IEEE, 2015, pp. 508–513. DOI: 10.1109/MED.2015.7158798.
- [52] HÄHNEL, C. ; AUL, V. ; HORN, J.: Power control for efficient operation of a PEM fuel cell system by nonlinear model predictive control. In: *IFAC-PapersOnLine* 48. no. 11 (2015), pp. 174–179. DOI: 10.1016/j.ifacol.2015.09.179.
- [53] HÄHNEL, C. ; AUL, V. ; HORN, J.: Power efficient operation of a PEM fuel cell system using cathode pressure and excess ratio by nonlinear model predictive control. In: *2015 European Control Conference (ECC)*. IEEE, 2015, pp. 3340–3345. DOI: 10.1109/ECC.2015.7331050.
- [54] HÄHNEL, C. ; AUL, V. ; HORN, J.: Online identification of an electric PEMFC model for power control by NMPC. In: *2015 20th International Conference on Methods and Models in Automation and Robotics (MMAR)*. IEEE, 2015, pp. 133–138. DOI: 10.1109/MMAR.2015.7283860.
- [55] HÄHNEL, C. ; CLOPPENBORG, A. ; HORN, J.: Offset-free nonlinear model predictive control of electrical power of a PEM fuel cell system using an Extended Kalman Filter. In: *2016 24th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2016, pp. 106–111. DOI: 10.1109/MED.2016.7535896.

- [56] HASEGAWA, T. ; IMANISHI, H. ; NADA, M. ; IKOGI, Y.: Development of the fuel cell system in the Mirai FCV. In: *SAE Technical Paper Series*. 2016. DOI: 10.4271/2016-01-1185.
- [57] INTERNATIONAL ENERGY AGENCY: *Global EV outlook 2024: Moving towards increased affordability*. 2024.
- [58] JAHNKE, T. ; FUTTER, G. ; LATZ, A. ; MALKOW, T. ; PAPAKONSTANTINO, G. ; TSOTRIDIS, G. ; SCHOTT, P. ; GÉRARD, M. ; QUINAUD, M. ; QUIROGA, M. ; FRANCO, A. A. ; MALEK, K. ; CALLE-VALLEJO, F. ; FERREIRA DE MORAIS, R. ; KERBER, T. ; SAUTET, P. ; LOFFREDA, D. ; STRAHL, S. ; SERRA, M. ; POLVERINO, P. ; PIANESE, C. ; MAYUR, M. ; BESSLER, W. G. ; KOMPIS, C.: Performance and degradation of Proton Exchange Membrane Fuel Cells: State of the art in modeling from atomistic to system scale. In: *Journal of Power Sources* 304 (2016), pp. 207–233. DOI: 10.1016/j.jpowsour.2015.11.041.
- [59] JOŠEVSKI, M.: *Predictive energy management of hybrid electric vehicles with uncertain torque demand forecast for on-road operation*. PhD thesis. RWTH Aachen University, 2018.
- [60] KALMAN, R. E.: A new approach to linear filtering and prediction problems. In: *Journal of Basic Engineering* 82. no. 1 (1960), pp. 35–45. DOI: 10.1115/1.3662552.
- [61] KARASUYAMA, M. ; KASUGAI, H. ; TAMURA, T. ; SHITARA, K.: Computational design of stable and highly ion-conductive materials using multi-objective bayesian optimization: Case studies on diffusion of oxygen and lithium. In: *Computational Materials Science* 184 (2020), p. 109927. DOI: 10.1016/j.commatsci.2020.109927.
- [62] KLELL, M. ; EICHLSEDER, H. ; TRATTNER, A.: *Wasserstoff in der Fahrzeugtechnik: Erzeugung, Speicherung, Anwendung*. 4. aktualisierte und erweiterte Auflage. ATZ/MTZ-Fachbuch. Springer Vieweg, 2018.
- [63] KONAKOVIC LUKOVIC, M. ; TIAN, Y. ; MATUSIK, W.: Diversity-guided multi-objective Bayesian optimization with batch evaluations. In: *Advances in Neural Information Processing Systems*. Vol. 33. Curran Associates, Inc., 2020, pp. 17708–17720.
- [64] KURZWEIL, P.: *Brennstoffzellentechnik: Grundlagen, Materialien, Anwendungen, Gaserzeugung*. Springer-Verlag, 2016.
- [65] KUSHNER, H. J.: A new method of locating the maximum point of an arbitrary multi-peak curve in the presence of noise. In: *Journal of Basic Engineering* 86. no. 1 (1964), pp. 97–106. DOI: 10.1115/1.3653121.
- [66] LARMINIE, J. ; DICKS, A.: *Fuel cell systems explained*. 1st edition. Wiley, 2003. DOI: 10.1002/9781118878330.

- [67] LIM, B. H. ; MAJLAN, E. H. ; DAUD, W. R. W. ; HUSAINI, T. ; ROSLI, M. I.: Effects of flow field design on water management and reactant distribution in PEMFC: A review. In: *Ionics* 22. no. 3 (2016), pp. 301–316. DOI: 10.1007/s11581-016-1644-y.
- [68] LU, Z. ; DAINO, M. M. ; RATH, C. ; KANDLIKAR, S. G.: Water management studies in PEM fuel cells, part III: Dynamic breakthrough and intermittent drainage characteristics from GDLs with and without MPLs. In: *International Journal of Hydrogen Energy* 35. no. 9 (2010), pp. 4222–4233. DOI: 10.1016/j.ijhydene.2010.01.012.
- [69] LUNA, J. ; JEMEI, S. ; YOUSFI-STEINER, N. ; HUSAR, A. ; SERRA, M. ; HISSEL, D.: Nonlinear predictive control for durability enhancement and efficiency improvement in a fuel cell power system. In: *Journal of Power Sources* 328 (2016), pp. 250–261. DOI: 10.1016/j.jpowsour.2016.08.019.
- [70] LUNA, J. ; USAI, E. ; HUSAR, A. ; SERRA, M.: Enhancing the efficiency and lifetime of a proton exchange membrane fuel cell using nonlinear model-predictive control with nonlinear observation. In: *IEEE Transactions on Industrial Electronics* 64. no. 8 (2017), pp. 6649–6659. DOI: 10.1109/TIE.2017.2682787.
- [71] MATÉRN, B.: *Spatial variation*. Springer New York, 1986. DOI: 10.1007/978-1-4615-7892-5.
- [72] MOČKUS, J.: On bayesian methods for seeking the extremum. In: *Optimization Techniques*. Springer Berlin Heidelberg, 1975, pp. 400–404. DOI: 10.1007/3-540-07165-2{\textunderscore}55.
- [73] NEISEN, V.: *Hierarchische Modellprädiktive Regelung für Brennstoffzellen-Batteriesysteme auf multiplen Zeitskalen*. PhD thesis. RWTH Aachen University, 2024.
- [74] NEISEN, V. ; MANNHARDT, J. ; ABEL, D.: Dynamic tracking of power demand for integrated fuel cell systems using nonlinear model predictive control. In: *IFAC-PapersOnLine* 53. no. 2 (2020), pp. 13216–13223. DOI: 10.1016/j.ifacol.2020.12.148.
- [75] NOCEDAL, J. ; WRIGHT, S. J., eds.: *Numerical optimization*. 2nd edition. Springer Series in Operations Research and Financial Engineering. Springer, 2006. DOI: 10.1007/b98874.
- [76] NUSS, E.: *Modeling and nonlinear model predictive control for gasoline-controlled auto-ignition*. PhD thesis. RWTH Aachen University, 2024.
- [77] OKONKWO, P. C. ; OTOR, C.: A review of gas diffusion layer properties and water management in proton exchange membrane fuel cell system. In: *International Journal of Energy Research* 45. no. 3 (2021), pp. 3780–3800. DOI: 10.1002/er.6227.

- [78] OLSON, M. ; SANTORELLA, E. ; TIAO, L. C. ; CAKMAK, S. ; GARRARD, M. ; DAULTON, S. ; LIN, Z. J. ; AMENT, S. ; BECKERMAN, B. ; ONOFREY, E. ; IGUSTI, P. ; LARA, C. ; LETHAM, B. ; CARDOSO, C. ; SHEN, S. S. ; LIN, A. C. ; GRANGE, M. ; KASHTELYAN, E. ; ERIKSSON, D. ; BALANDAT, M. ; BAKSHY, E.: Ax: A Platform for Adaptive Experimentation. In: *AutoML 2025 ABCD Track* (2025).
- [79] OMRANI, R. ; SHABANI, B.: Gas diffusion layer modifications and treatments for improving the performance of proton exchange membrane fuel cells and electrolyzers: A review. In: *International Journal of Hydrogen Energy* 42. no. 47 (2017), pp. 28515–28536. DOI: 10.1016/j.ijhydene.2017.09.132.
- [80] PADGET, E. ; KLEEN, G.: *Automotive fuel cell targets and status*. 2020.
- [81] PATHAK, M. ; SLADE, R. ; SHUKLA, P. R. ; SKEA, J. ; PICHES-MADRUGA, R., ÜRGEVORSATZ, D.: Mitigation of climate change: Technical summary. In: *Climate change 2022*. Cambridge University Press, 2023, pp. 51–148. DOI: 10.1017/9781009157926.002.
- [82] POWELL, M. J. D.: A fast algorithm for nonlinearly constrained optimization calculations. In: *Numerical Analysis*. Lecture Notes in Mathematics Ser. Springer Berlin Heidelberg, 1978, pp. 144–157. DOI: 10.1007/bfb0067703.
- [83] PUKRUSHPAN, J. T. ; PENG, H. ; STEFANOPOULOU, A. G.: Control-oriented modeling and analysis for automotive fuel cell systems. In: *Journal of Dynamic Systems, Measurement, and Control* 126. no. 1 (2004), pp. 14–25. DOI: 10.1115/1.1648308.
- [84] PUKRUSHPAN, J. T. ; STEFANOPOULOU, A. G. ; PENG, H.: Control of fuel cell breathing. In: *IEEE Control Systems Magazine* 24. no. 2 (2004), pp. 30–46. DOI: 10.1109/MCS.2004.1275430.
- [85] PUKRUSHPAN, J. T. ; STEFANOPOULOU, A. G. ; PENG, H.: *Control of fuel cell power systems: Principles, modeling, analysis and feedback design*. Springer Science & Business Media, 2004.
- [86] PUKRUSHPAN, J. T. ; STEFANOPOULOU, A. G. ; PENG, H.: *Fuel cell stack simulation package*. 2002.
- [87] QIU, D. ; PENG, L. ; LAI, X. ; NI, M. ; LEHNERT, W.: Mechanical failure and mitigation strategies for the membrane in a proton exchange membrane fuel cell. In: *Renewable and Sustainable Energy Reviews* 113 (2019), p. 109289. DOI: 10.1016/j.rser.2019.109289.
- [88] QUIRYNEN, R.: *Numerical simulation methods for embedded optimization*. PhD thesis. University of Freiburg, 2017.
- [89] RAJASINGHAM, T. A.: *Nonlinear model predictive control of combustion engines: From fundamentals to applications*. Springer Nature, 2021.

- [90] RAO, C. V. ; WRIGHT, S. J. ; RAWLINGS, J. B.: Application of interior-point methods to model predictive control. In: *Journal of Optimization Theory and Applications* 99. no. 3 (1998), pp. 723–757. DOI: 10.1023/A:1021711402723.
- [91] RASMUSSEN, C. E. ; WILLIAMS, C. K. I.: *Gaussian processes for machine learning*. The MIT Press, 2005. DOI: 10.7551/mitpress/3206.001.0001.
- [92] RAWLINGS, J. B. ; MAYNE, D. Q. ; DIEHL, M. M.: *Model predictive control: Theory, computation, and design*. 4th printing. Nob Hill Publishing, 2022.
- [93] REN, P. ; PEI, P. ; LI, Y. ; WU, Z. ; CHEN, D. ; HUANG, S.: Degradation mechanisms of proton exchange membrane fuel cell under typical automotive operating conditions. In: *Progress in Energy and Combustion Science* 80 (2020), p. 100859. DOI: 10.1016/j.pecs.2020.100859.
- [94] RITZBERGER, D. ; HAMETNER, C. ; JAKUBEK, S.: A real-time dynamic fuel cell system simulation for model-based diagnostics and control: Validation on real driving data. In: *Energies* 13. no. 12 (2020), p. 3148. DOI: 10.3390/en13123148.
- [95] RITZBERGER, D. ; HÖFLINGER, J. ; DU, Z. P. ; HAMETNER, C. ; JAKUBEK, S.: Data-driven parameterization of polymer electrolyte membrane fuel cell models via simultaneous local linear structured state space identification. In: *International Journal of Hydrogen Energy* 46. no. 21 (2021), pp. 11878–11893. DOI: 10.1016/j.ijhydene.2021.01.037.
- [96] SAUERMOSEER, M. ; KIZILOVA, N. ; POLLET, B. G. ; KJELSTRUP, S.: Flow field patterns for proton exchange membrane fuel cells. In: *Frontiers in Energy Research* 8 (2020). DOI: 10.3389/fenrg.2020.00013.
- [97] SCHMITT, L. ; ABEL, D.: Nonlinear MPC for fuel cell air path control with experimental validation. In: *2023 31st Mediterranean Conference on Control and Automation (MED)*. IEEE, 2023, pp. 89–94. DOI: 10.1109/med59994.2023.10185785.
- [98] SCHMITT, L. ; ABEL, D.: Numerical integration for nonlinear model predictive control of a fuel cell system. In: *2023 American Control Conference (ACC)*. IEEE, 2023, pp. 2633–2639. DOI: 10.23919/acc55779.2023.10156027.
- [99] SCHMITT, L. ; NICKIG, N. ; BAHR, M. ; GÖSSLING, S. ; ABEL, D.: Review, evaluation and application of condensing algorithms for model predictive control based on a first-order method. In: *2023 European Control Conference (ECC)*. IEEE, 2023, pp. 1–7. DOI: 10.23919/ecc57647.2023.10178157.
- [100] SHAHRIARI, B. ; SWERSKY, K. ; WANG, Z. ; ADAMS, R. P. ; FREITAS, N. de: Taking the human out of the loop: A review of Bayesian optimization. In: *Proceedings of the IEEE* 104. no. 1 (2016), pp. 148–175. DOI: 10.1109/JPROC.2015.2494218.
- [101] SHAHZAD, A. ; KERRIGAN, E. C. ; CONSTANTINIDES, G. A.: *A warm-start interior-point method for predictive control*. IET Digital Library, 2010. DOI: 10.1049/ic.2010.0409.

- [102] SHUKLA, P. R. ; SKEA, J. ; REISINGER, A. ; SLADE, R. ; FRADERA, R. ; PATHAK, M. ; AL KHOURDAJIE, A. ; BELKACEMI, M. ; VAN DIEMEN, R. ; HASIJA, A. ; LISBOA, G. ; LUZ, S. ; MALLEY, J. ; MCCOLLUM, D. ; SOME, S. ; VYAS, P.: *Summary for policymakers*. IPCC 2022: Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, NY, USA: Cambridge University Press, 2022. DOI: 10.1017/9781009157926.001.
- [103] SIMON, D.: *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.
- [104] SINGH, R. S. ; GAUTAM, A. ; RAI, V.: Graphene-based bipolar plates for polymer electrolyte membrane fuel cells. In: *Frontiers of Materials Science* 13. no. 3 (2019), pp. 217–241. DOI: 10.1007/s11706-019-0465-0.
- [105] SPRINGER, T. E. ; ZAWODZINSKI, T. A. ; GOTTESFELD, S.: Polymer electrolyte fuel cell model. In: *Journal of The Electrochemical Society* 138. no. 8 (1991), p. 2334. DOI: 10.1149/1.2085971.
- [106] SRINIVAS, N. ; KRAUSE, A. ; KAKADE, S. M. ; SEEGER, M.: Gaussian process optimization in the bandit setting: No regret and experimental design. In: *arXiv preprint arXiv:0912.3995* (2009). DOI: 10.1109/TIT.2011.2182033.
- [107] STEIN, M. L.: *Interpolation of spatial data: Some theory for kriging*. Vol. 42. Springer Science & Business Media, 1999. DOI: 10.2307/1270967.
- [108] SUH, K. W. ; STEFANOPOULOU, A. G.: Inherent performance limitations of power-autonomous fuel cell system. In: *2006 American Control Conference (ACC)*. IEEE, 2006. DOI: 10.1109/acc.2006.1655404.
- [109] SUH, K. W. ; STEFANOPOULOU, A. G.: Performance limitations of air flow control in power-autonomous fuel cell systems. In: *IEEE Transactions on Control Systems Technology* 15. no. 3 (2007), pp. 465–473. DOI: 10.1109/tcst.2007.894640.
- [110] TANIGUCHI, A. ; AKITA, T. ; YASUDA, K. ; MIYAZAKI, Y.: Analysis of degradation in PEMFC caused by cell reversal during air starvation. In: *International Journal of Hydrogen Energy* 33. no. 9 (2008), pp. 2323–2329. DOI: 10.1016/j.ijhydene.2008.02.049.
- [111] THOMPSON, W. R.: On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. In: *Biometrika* 25. no. 3/4 (1933), p. 285. DOI: 10.2307/2332286.
- [112] TREMBLAY, O. ; DESSAINT, L.-A. ; DEKKICHE, A.-I.: A generic battery model for the dynamic simulation of hybrid electric vehicles. In: *2007 IEEE Vehicle Power and Propulsion Conference*. IEEE, 2007. DOI: 10.1109/vppc.2007.4544139.

-
- [113] VAHIDI, A. ; KOLMANOVSKY, I. ; STEFANOPOULOU, A.: Constraint management in fuel cells: a fast reference governor approach. In: *2005 American Control Conference (ACC)*. Vol. 6. IEEE, 2005, pp. 3865–3870. DOI: 10.1109/acc.2005.1470577.
- [114] VAHIDI, A. ; STEFANOPOULOU, A. ; HUEI PENG: Model predictive control for starvation prevention in a hybrid fuel cell system. In: *2004 American Control Conference (ACC)*. IEEE, 2004. DOI: 10.23919/acc.2004.1383709.
- [115] VAHIDI, A. ; STEFANOPOULOU, A. ; PENG, H.: Current management in a hybrid fuel cell power system: A model-predictive control approach. In: *IEEE Transactions on Control Systems Technology* 14. no. 6 (2006), pp. 1047–1057. DOI: 10.1109/tcst.2006.880199.
- [116] VAHIDI, A. ; GREENWELL, W.: A decentralized model predictive control approach to power management of a fuel cell-ultracapacitor hybrid. In: *2007 American Control Conference (ACC)*. IEEE, 2007. DOI: 10.1109/acc.2007.4282904.
- [117] VAHIDI, A. ; KOLMANOVSKY, I. ; STEFANOPOULOU, A.: Constraint handling in a fuel cell system: A fast reference governor approach. In: *IEEE Transactions on Control Systems Technology* 15. no. 1 (2007), pp. 86–98. DOI: 10.1109/TCST.2006.883242.
- [118] VERSCHUEREN, R.: *Convex approximation methods for nonlinear model predictive control*. PhD thesis. University of Freiburg, 2018.
- [119] VERSCHUEREN, R. ; FRISON, G. ; KOUZOUPIS, D. ; FREY, J. ; VAN DUIJKEREN, N. ; ZANELLI, A. ; NOVOSELNİK, B. ; ALBIN, T. ; QUIRYNEN, R. ; DIEHL, M.: acados—a modular open-source framework for fast embedded optimal control. In: *Mathematical Programming Computation* 14. no. 1 (2022), pp. 147–183. DOI: 10.1007/s12532-021-00208-8.
- [120] VRLIĆ, M. ; RITZBERGER, D. ; JAKUBEK, S.: Efficient and life preserving power tracking control of a proton exchange membrane fuel cell using model predictive control. In: *2020 SICE International Symposium on Control Systems (SICE ISCS)*. IEEE, 2020. DOI: 10.23919/siceiscs48470.2020.9083653.
- [121] VRLIĆ, M. ; RITZBERGER, D. ; JAKUBEK, S.: Safe and efficient polymer electrolyte membrane fuel cell control using successive linearization based model predictive control validated on real vehicle data. In: *Energies* 13. no. 20 (2020), p. 5353. DOI: 10.3390/en13205353.
- [122] VRLIĆ, M. ; RITZBERGER, D. ; JAKUBEK, S.: Model-predictive-control-based reference governor for fuel cells in automotive application compared with performance from a real vehicle. In: *Energies* 14. no. 8 (2021), p. 2206. DOI: 10.3390/en14082206.
- [123] WÄCHTER, A. ; BIEGLER, L. T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. In: *Mathematical Programming* 106. no. 1 (2006), pp. 25–57. DOI: 10.1007/s10107-004-0559-y.

- [124] WANG, W. ; SEBAG, M.: Hypervolume indicator and dominance reward based multi-objective Monte-Carlo Tree Search. In: *Machine Learning* 92. no. 2 (2013), pp. 403–429. DOI: 10.1007/s10994-013-5369-0.
- [125] WELCH, G. ; BISHOP, G.: An introduction to the Kalman filter. In: 8 (2006).
- [126] WU, J. ; YUAN, X. Z. ; MARTIN, J. J. ; WANG, H. ; ZHANG, J. ; SHEN, J. ; WU, S. ; MERIDA, W.: A review of PEM fuel cell durability: Degradation mechanisms and mitigation strategies. In: *Journal of Power Sources* 184. no. 1 (2008), pp. 104–119. DOI: 10.1016/j.jpowsour.2008.06.006.
- [127] XU, Z. ; QIU, D. ; YI, P. ; PENG, L. ; LAI, X.: Towards mass applications: A review on the challenges and developments in metallic bipolar plates for PEMFC. In: *Progress in Natural Science: Materials International* 30. no. 6 (2020), pp. 815–824. DOI: 10.1016/j.pnsc.2020.10.015.
- [128] YANG, K. ; EMMERICH, M. ; DEUTZ, A. ; BÄCK, T.: Multi-objective Bayesian global optimization using expected hypervolume improvement gradient. In: *Swarm and Evolutionary Computation* 44 (2019), pp. 945–956. DOI: 10.1016/j.swevo.2018.10.007.
- [129] ZHANG, W. ; LI, J. ; XU, L. ; OUYANG, M.: Optimization for a fuel cell/battery/capacity tram with equivalent consumption minimization strategy. In: *Energy Conversion and Management* 134 (2017), pp. 59–69. DOI: 10.1016/j.enconman.2016.11.007.