Optimal Robust Distributed Control Towards Attosecond Synchronization Systems

Vom Promotionsausschuss der
Technischen Universität Hamburg

zur Erlangung des akademischen Grades

Doktor-Ingenieur (Dr.-Ing.)

genehmigte Dissertation (Monografie)

von

Maximilian Schütte

aus

Donaueschingen

2026

1. Gutachter:

Prof. Dr. Herbert Werner

Institut für Regelungstechnik
Technische Universität Hamburg

2. Gutachterin:

Prof. Dr.-Ing. Annika Eichler

Maschine Strahlenkontrollen
Deutsches Elektronen-Synchrotron DESY

Institut für Regelungstechnik
Technische Universität Hamburg

3. Gutachter:

Prof. Dr.-Ing. Gerwald Lichtenberg

Fakultät Life Sciences
Hochschule für Angewandte Wissenschaften Hamburg

Vorsitzender des Prüfungsauschusses:

Prof. Dr.-Ing. Görschwin Fey

Institut für Eingebettete Systeme
Technische Universität Hamburg

Tag der Disputation: 10. November 2025

10.15480/882.17063
0000-0002-4684-0633

ISBN 978-3-8439-5771-7 (softcover)

First Edition

Abstract

This thesis is concerned with the modelling of, and controller design for precision synchronization networks, posing a distributed control problem. One application for such systems from the frontiers of experimental physics are free-electron lasers, tools to study matter down to the submolecular level. When aiming to spatially and temporally resolve matter and processes on atomic scales, the elementary diffraction limit and the attosecond dynamics of electrons formulate the need for pulses of light at very short wavelengths in the hard X-ray regime and shorter than a few femtoseconds. Such pulses today can only be created in large facilities such as the European XFEL, generating coherent synchrotron radiation from accelerated packets of electrons. The short timescales and experimental techniques such as pump-probe schemes pose stringent synchronization requirements on the facility. Given the large spatial extent of the facility of a few kilometres, laser-based synchronization networks have become indispensable to achieve femtosecond timing stability. Such systems consist of time-of-flight stabilized optical fibre links and several phase-locked loops interconnecting conventional RF oscillators and mode-locked lasers, functioning as oscillators in the optical domain.

By taking into account the interconnection structure of precision synchronization networks, this work aims to improve the overall synchronization performance to attosecond levels RMS, while simultaneously ensuring robust operation. The latter is crucial because of the sensitive optical components and high cost of facility downtime. From a control theoretic perspective, the performance goal is expressed using the well-known $\mathbf{H}_2$-norm. In combination with a distributed controller architecture, this problem is however also known to be NP-hard. Nevertheless, approaches towards acceptable suboptimal solutions exist in the form of convex approximations, and nonsmooth optimization. For robustness, the method of bounding model uncertainties using multipliers has received substantial attention over the last decades and was generalized in the integral quadratic constraint (IQC) framework. One particularly relevant feature is the possibility to account for uncertain time delays, which are non-negligible in large synchronization systems.

Despite mature individual solutions for the performance, robustness and distributed system aspects of control system design, the joint consideration remains challenging. After developing the required models for phase-locked loops and stabilized timing links, identifying them as singular variants of the standard $\mathbf{H}_2$-problem, this work presents novel solutions to the robust distributed output feedback problem by means of linear matrix inequalities and, alternatively, by exact nonsmooth minimum singular value problems in the frequency domain. Particular contributions are the extension of the sparsity invariance framework for least conservative convex restrictions to the output feedback case via a specific state-space form and the incorporation of static multipliers via the full-block S-procedure for robust stability and performance. In addition, the full versatility of the IQC framework, including dynamic multipliers, is made available to the problem via the nonsmooth optimization approach. Here a suitable performance multiplier bounding the $\mathbf{H}_2$-norm is designed along with implementation of an efficient analytic solution to the attached frequency domain integral.

Using these theoretic results and a model of the laser-based synchronization system at European XFEL identified from extensive measurements, a detailed distributed controller design study is conducted. It is shown how these methods perform with challenging large-scale models, both in terms of system order, number and size of internal delays, and number of disturbance inputs. Immediate recommendations are derived for future development of the laser-based synchronization system, highlighting the importance of the distributed system aspect. Although it cannot be seen as a standalone remedy, optimal controller synthesis is an important step on the path towards attosecond synchronization.

It is only with the heart
that one can see rightly;
what is essential is invisible to the eye.

– The fox to the little prince

Acknowledgements

This dissertation came to fruition over the course of five years of work alongside my role in the beam machine controls group at DESY. During this time, I had the pleasure of working on a variety of fascinating technical topics, many of which unfortunately did not find their way into this work, as is the nature of such an endeavour. While challenging at times, it was this variety that still keeps me curious and helped me push through inevitable moments of doubt and frustration. An even bigger contribution to the success of this project is owed however to the many people who supported me along the way in various and many ways.

First, I'd like to thank Prof. Herbert Werner for supervising this thesis and providing valuable feedback on both research papers and this manuscript, and his research group for stimulating exchange on the topic of control theory. I'd also like to thank Prof. Gerwald Lichtenberg for his enthusiasm for our work at DESY, his many interesting ideas for research and collaborations, and for giving me opportunity to teach at university, something I enjoyed tremendously. Of all my collaborators, my most heartfelt gratitude is owed to Prof. Annika Eichler, who, after supervising my master's thesis at ETH Zürich, offered me the interview for a position as doctoral researcher at DESY and thus made all of this possible. The depth of her knowledge and her persistence in pursuing her goals while still always taking time for her students, whom she treats almost like her family, not only benefitted me but humbled me deeply. She has been my defining academic mentor and role model for the past decade and I am forever grateful for her support and guidance.

Furthermore, I'd like to express my appreciation for the many amazing colleagues at DESY and more specifically the optical synchronization team. Special thanks goes to Holger Schlarb, Jost Müller and Sebastian Schulz, who entrusted me with a great deal of freedom in my work and always were supportive of the ideas I put forward and requests I had and also let me travel to several great conferences and workshops. Among many more, I'd like to also especially thank Matthias Felber, Thorsten Lamb, Falco Zummack, Mikhail Titberidze, Tomasz Kozak, Anne-Laure Calendron, Nick Kschuev, Sven Pfeiffer, Frank Ludwig, Christian Schmidt, Marie Kristin Czwalinna, Vladimir Rybnikov, and Tim Wilksen for their support and assistance, their many teachings and explanations, fruitful discussions, or just cheerful banter while working in the lab.

I'm indebted to my students who repeatedly put my expertise to test and provided helpful input to this thesis, especially Jannis. I'm grateful for the technical and emotional support from my fellow doctoral researchers – in particular Arne, Jan, Adwait, Bindu, and Christian, with whom I had enlightening discussions on LMIs and numerical stability. Beyond that, they are simply a wonderful bunch who I'm glad to consider my friends. I hope our paths continue to cross in the years ahead.

Above all, I'm thankful for my close friends in Hamburg and across the world – those who smile when reading this – and for my family, especially my loving parents and brother. I never got anything but support and love from all of them, and whatever was weighing on my mind, I could be sure to find open ears and caring eyes. Their belief and trust in me as much as mine in them is the bedrock on which everything is built, including this thesis. None of this would have felt worth doing without them.
Thank you for everything. 💛

Contents

Symbols

$\mathbf{N}$	Set of natural numbers
$\mathbf{Z}$	Set of integers
$\mathbf{R}$	Field of real numbers
$\mathbf{R}_+$	Field of nonnegative real numbers
$\mathbf{C}$	Field of complex numbers
$\mathbf{C}_+$	Open right-haft plane
$\bar{\mathbf{C}}_+$	Closed right-haft plane
$\mathbf{F}$	Either field of reals or complex numbers
$\mathbf{S}^n$	Real symmetric matrices of size $n \times n$
v	Signal, a function of time
$\hat{v}$	Fourier transformed signal
I_n	Identity matrix, size $n \times n$
0_n	Zero matrix, size $n \times n$
diag$(A, B, \dots)$	Block diagonal matrix with blocks $A, B, \dots$ on the main diagonal
daug$(A, B, \dots)$	Diagonally augmented matrix
trace(A)	Trace of A
A^{H}	Conjugate transpose of A
$\bar{\sigma}(A)$	Largest singular value of A
im(A)	Image of A
$\otimes$	Kronecker product
BV	Banach space of real-valued functions of bounded variation
BV$_\downarrow$	Space of functions in BV bounded by a monotonically decreasing function
$\mathcal{C}_{\mathsf{H}}$	Banach space of continuous, Hermitian matrix-valued functions
$\mathbf{L}_\infty$	Banach space of complex, matrix-valued functions essentially bounded on $i\mathbf{R}$
$\mathbf{L}_2$	Hilbert space of complex, matrix-valued functions with bounded integral on $i\mathbf{R}$
$\mathbf{H}_\infty$	Subspace of $\mathbf{L}_\infty$, analytic in the open RHP
$\mathbf{H}_2$	Subspace of $\mathbf{L}_2$, analytic in the open RHP
$\mathbf{L}_{2_e}$	Extended space
prefix $\mathbf{R}$	Real rational variant, e.g. $\mathbf{RL}_\infty$
sect(X)	Bilinear sector transform of X

G	Physical system (plant), given by a (bi)proper transfer function
T	(Real rational) transfer function (not necessarily proper)
P	Generalized plant
K	Feedback controller
Π	(Dynamic) multiplier
H	(Nonlinear, time-varying) Operator
$\mathbf{\Delta}_{\mathrm{NLTV}}$	Set of causal, bounded, non-linear, time-varying operators on space $\mathbf{L}_{2_e}$
$\mathcal{F}_l$	Lower linear fractional transformation
$\mathcal{F}_u$	Upper linear fractional transformation
$R_{\mathbf{x}}$	Autocorrelation function of $\mathbf{x}$
S_x	Two-sided power spectral density of $\mathbf{x}$
S'_x	One-sided power spectral density of $\mathbf{x}$
c_0	Speed of light

Acronyms

ADC	analog-to-digital converter
AM	amplitude modulation
AMC	advanced mezzanine card
ANC	active noise cancellation
ARE	algebraic Riccati equation
ASE	amplified spontaneous emission
ASPECT	attosecond pulses with eSASE and chirp-taper schemes
AWGN	additive white Gaussian noise
BAM	bunch arrival time monitor
BBO	beta barium borate
BIBO	bounded-input, bounded-output
BO	Bayesian optimization
BOXC	balanced optical cross-correlator
DAC	digital-to-analog converter
DC	direct current
DESY	Deutsches Elektronen-Synchrotron
DFT	discrete Fourier transform
DMA	direct memory access
DOOCS	Distributed Object-Oriented Control System
DRO	dielectric resonator oscillator
DSP	digital signal processor
DTFT	discrete time Fourier transform
DUT	device under test
EDFA	erbium-doped fibre amplifier
EMI	electromagnetic interference
FDI	frequency domain inequality
FEL	free-electron laser
FLASH	Free-Electron Laser in Hamburg
FPGA	field-programmable gate array
GMT	Greenwich Mean Time

GNSS	global navigation satellite system
HPA	high-power amplifier
IC	integrated circuit
IF	intermediate frequency
IIR	infinite impulse response
ILC	International Linear Collider
IQ	in-phase, quadrature
IQC	integral quadratic constraint
ISE	integral square error
KYP	Kalman-Yakubovich-Popov
LAM	laser arrival time monitor
LbSync	laser-based synchronization
LFR	linear fractional representation
LFT	linear fractional transformation
LHP	left-half plane
LINAC	linear accelerator
LMI	linear matrix inequality
LO	local oscillator
LQG	linear-quadratic-Gaussian
LQR	linear-quadratic regulator
LTI	linear time-invariant
LTV	linear time varying
MicroTCA	Micro Telecommunications Computing Architecture
MIMO	multiple input, multiple output
MLL	mode-locked laser
MLO	master laser oscillator
MO	master oscillator
MUTC-PD	modified uni-traveling-carrier photodiode
MZM	Mach-Zehnder modulator
OFD	optical frequency division
OXCO	oven-controlled crystal oscillator
PAM	photon arrival time monitor
PCIe	Peripheral Component Interconnect Express
PE	polyethylene
PFA	photodiode, filter, amplifier
PI	proportional-integral
PID	proportional-integral-derivative
PLL	phase-locked loop
PM	phase modulation

PM	polarization maintaining
PPKTP	periodically poled potassium titanyl phosphate
PPL	pump-probe laser
PSD	power spectral density
PTFE	polytetrafluorethylen
QI	quadratic invariance
REFM-OPT	optical reference module
RF	radio frequency
RHP	right-half plane
RIN	relative intensity noise
RMS	root mean square
RSS	root sum square
RTM	Rear Transition Modules
SASE	self-amplified spontaneous emission
SDP	semi-definite program
SFG	sum-frequency generation
SHG	second-harmonic generation
SI	sparsity invariance
SISO	single input, single output
SLO	subsidiary laser oscillator
SNR	signal-to-noise ratio
SPI	Serial Peripheral Interface
SSB	single-sideband
SXP	soft X-ray port
TDR	technical design report
TDS	time-delay system
TESLA	Tera-Electronvolt Energy Superconducting Linear Accelerator
TTF	TESLA test facility
VCO	voltage-controlled oscillator
VHDL	VHSIC Hardware Description Language

Chapter 1

Introduction

1.1 Time and Synchronization

The need to synchronize agents, machines, processes and events is ubiquitous in the modern world. While rudimentary timepieces have been around since the dawn of civilization, the accuracy of clocks first became a limiting factor and serious technical challenge for cartographers and marine navigators with what is known as the longitude problem during the 17^{th} and 18^{th} centuries [Wil92]. Determining one's own longitudinal position anywhere on earth by measuring the time offset of an astronomical event such as high noon with a clock showing the time of a reference location (usually that of the Royal Observatory, Greenwich), is prone to an error of $1'$ per 4 s of timing deviation, which on the equator translates to just over one nautical mile. Before Harrison invented the famous H1 marine chronometer, such errors could easily accumulate to a few hundred kilometres over the course of a long voyage, leading to many of histories most perilous disasters at sea. The critical need for timekeeping on ships lead to the implementation of Greenwich Mean Time (GMT) as a first time standard, and the use of acoustic and visual signals in ports to transmit this reference time to ships within range for local correction of the ship's less accurate marine chronometers [BD81]. The limited range of these synchronization signals was soon drastically extended by the electric telegraph and today, highly accurate time signals can be received worldwide by means of radio communication, e.g. via global navigation satellite systems (GNSSs).

Unsurprisingly however, modern technologies and especially cutting-edge experimental physics have pushed the requirements for timing synchronization far beyond what can be achieved with commonly available (these days usually fully digital) solutions. Albert Einstein's *annus mirabilis* paper on special relativity, "Zur Elektrodynamik bewegter Körper" [Ein05], notably starts with a revised definition of synchronism and ultimately shows that it has no *absolute* meaning, but depends on the observer's frame of reference. While Einstein's famed *Gedankenexperimente* had the luxury of just conjuring up perfectly stable and synchronized clocks in different places, the numerous validation experiments for special and general relativity had to overcome significant technological challenges in order to measure the oftentimes only minuscule effects. For example, in [Cho+10], the authors demonstrate the effect of gravitational time dilation resulting from a height difference of less than half a meter on the earth's surface. This was achieved by measuring the expected relative frequency change of just $\sim 4 \cdot 10^{-17}$ between two ultra-stable optical clocks, where the second clock's analog signal is transferred to the first clock for comparison using a phase-stabilized optical fibre link.

(a) Illustration of an array of radio telescopes with star-topology RF network.

(b) Aerial photograph of the European XFEL in Hamburg, Germany. ©European XFEL GmbH.

Figure 1.1: Example applications for distributed precision synchronization networks.

As discussed in [For+07], using stabilized fibre links to distribute timing signals over long distances is a key-enabler also for other science applications. In radio astronomy, synchronously collected data from multiple distributed radio telescopes can be used to increase the effective aperture of the full array of telescopes [CS06]. Due to their kilometre-long interferometer arms, gravitational wave observatories are also becoming increasingly relevant. Another application, one that will be discussed extensively in this work, is the synchronization of subsystems in particle accelerators, which, in the example of the European XFEL, can extend to multiple kilometres and require a timing error of <10 fs RMS over 100 ms [Czu+13]. Future experiments studying electron dynamics in matter will require even higher levels of sub-femtosecond timing stability, see Section 5.1.1. Subsystems in need of synchronization are in particular the microwave oscillators generating the accelerating RF fields and laser oscillators in pump-probe style experiments as well as certain beam diagnostic devices. The common feature is that these subsystems generate, measure, or interact with ultra-short packets of particles moving at (close to) speed-of-light and hence have associated ultra-fast dynamics in the (sub)femtosecond regime.

As shown in Figure 1.1, applications commonly implement a timing distribution network of star- or tree-like topology, where the root is a reference timing source, the edges are stabilized transmission links and the intermediate nodes and leaves are local oscillators that are to be synchronized to the upstream reference. We therefore come to the realization that facility-wide ultra-high-precision timing distribution is a coupled control problem of many instances of two basic types of feedback loops, namely phase-locked loop (PLL)-like oscillator synchronization and transmission path stabilization. The figure of merit of such a network is the (possibly weighted) average synchronization error of any node with respect to a dedicated reference.

1.2 Control Challenges in Distributed Synchronization Systems

To achieve the required synchronization performance at all stations, several design aspects need to be considered simultaneously. Existing literature on femtosecond synchronization is overall scarce due to the rather specific field of application, and many publications deal with questions pertaining the implementation of individual system components, such as for example optical frequency division and interferometry. While these methods provide the technological foundation for achieving femtosecond precision synchronization, their potential can only be exhausted if the error information obtained from these devices is fed back to the correcting actuators by optimally tuned controllers. This is of particular importance because the current state-of-the-art generation of synchronization system components is already very expensive, challenging to manufacture, and scraping on physical limits in terms of noise and resolution, cf. Section 5.2.

The theory of PLLs is well known and essentially complete [Gar05; Ste97; Ban17; Abr02]. However, most of the works view the PLL from the perspective of an electrical engineering practitioner. An explicit link to control theoretic perspective of the underlying disturbance rejection problem and linear-quadratic-Gaussian (LQG) and $\mathbf{H}_2$ theory is rarely made. While inherently nonlinear, linearized models bode very well in practice for low-noise systems. Due to low operational power requirements, the optimal control problem is essentially "cheap", leading to a singular problem formulation, for which the standard algebraic Riccati equation (ARE)-based synthesis methods yield no solution [Sto92]. Though less prominently studied compared to their regular counterparts, singular problems have known solutions in the space of distributions [HS83], and perturbation and linear matrix inequality (LMI) methods yield practical but (controlled) suboptimal solutions [SL99].

Stabilized transmission links are similar in that sense, however generally do not feature integral dynamics like PLL voltage-controlled oscillators (VCOs), thus avoiding the problem of unstable pole-zero cancellations between the optimal controller and the plant. On the other hand, for sufficiently long transmission links (approx. 1 km and over), time-delay dynamics become non-negligible. Due to their infinite-dimensional nature, time-delay systems (TDSs) do not apply for many linear time-invariant (LTI) analysis and synthesis methods. Because the delays occur as state delays, most predictor methods such as the standard Smith predictor are eliminated. Finite-order real-rational approximations such as the Padé-approximant can be used, but fail to give any hard stability guarantees for the real system. Robust control techniques may be used to assert this property a-posteriori, or it can be taken into consideration directly in the synthesis step, these days commonly in the form of integral quadratic constraints (IQCs) [MR97; PS15], circumventing approximations entirely.

When considering timing synchronization networks, it is well known that, due to the coupled dynamics of PLLs and attached stabilized transmission links, individual tuning of the constituting interconnected feedback loops does not lead to optimal overall performance and may even deteriorate exponentially with increasing system size [Sch+21]. In order to perform joint synthesis of the network's controller parameters, a distributed system model, separating the network's topology from the individual subsystem dynamics, is needed. Then, during synthesis, the informational constraints imposed by the network topology, i.e. which state or error information is available to which controller and which controller is able to actuate which subsystem, need to be respected and fused with other design criteria such as robustness and output feedback constraints. At this point, convexity of the underlying optimization problem is generally lost [RL06], except for some special cases, see e.g. [EHW14; RL06]. Due to the increased interconnectivity of devices and agents in the modern world, distributed systems have attracted significant research interest over the last decade. For LTI systems, research into convex restrictions of the non-convex problem has recently originated the sparsity invariance (SI) framework [Fur+20] for static state-feedback or – via a Youla parameterization – dynamic output feedback of structured control problems. It was shown that this approach covers the least conservative convex restriction of the original problem, however deriving the convexification parameters yielding this particular restriction is still partially unsolved. Other approaches involve (nonsmooth) optimization over the non-convex cost function [AN06b; CS20] or IQC methods [Eic16].

Failure of any single node in the synchronization network may render the entire facility inoperable, causing significant financial cost. Since the controllers are designed toward a $\mathbf{H}_2$ specification, for which it is known that no stability margins exist [Doy78], a robust design, explicitly taking into account degradation and uncertainty in the model parameters is generally required in practice – especially for highly sensitive optical devices. Because generalization of the $\mathbf{H}_2$-norm for possibly nonlinear uncertainties is challenging [PF00], the mixed $\mathbf{H}_2/\mathbf{H}_\infty$ control problem is frequently studied instead in the literature [Sch96; ANR08]. Here, robust stability is guaranteed by a $\mathbf{H}_\infty$-criterion, while simultaneously optimizing $\mathbf{H}_2$-performance for the nominal system. Because this approach however gives no worst-case performance guarantees within the stability margins, it is less appealing in practice than a proper robust performance design, which is possibly conservative due to the generalization to the nonlinear regime [Cha+20; PF00].

1.3 Thesis Outline and Contribution

This work aims to provide a comprehensive treatment of distributed timing synchronization systems and the related control theoretic concepts, in particular optimal robust structured output feedback control. While not claiming full generality, the discussion and the results are abstracted to the largest reasonable extent from the application up to Chapter 5, where the results are applied to a specific application. To support the theoretic discussions, artificial numerical examples are added at the end of relevant sections.

The contents of this thesis are structured as follows. In Chapter 2, starting with a short summary of the notation used in this thesis, we provide preliminary materials that revisit the mathematical foundations of vector spaces, signals and operators as well as random variables and processes. These provide the basis for the subsequent review of standard modern control theory: LTI system models, feedback interconnections, and stability and performance analysis in the nominal setting. For robust analysis, we provide an introduction to the powerful IQC framework, which generalizes multiple previous approaches to robust control. Being of particular relevance to large-scale applications, a dedicated introduction to TDS and IQC multipliers for time delays is given thereafter. Finally, to familiarize the reader with the domain-specific nomenclature and the fundamentals that motivate the system models derived in the later chapters, an introduction to time and frequency metrology is given.

Chapter 3 treats the aforementioned two basic building blocks of synchronization systems and formulates the associated control problems of $\mathbf{H}_2$-optimal phase synchronization (Problem 3.5) and phase distribution (Problem 3.16). For these singular synthesis problems, state-of-the-art solution methods based on LMIs are provided and compared to solutions obtained from standard AREs for system models regularized by coefficient perturbation. For the stabilized timing link with non-negligible delay terms, we compare solutions obtained from standard methods for a Padé-approximation of the system with robust solutions obtained via the IQC framework.

Having established the models for the two subsystem types, along with the $\mathbf{H}_2$-synthesis LMIs, we turn to the networked setting in Chapter 4, starting with a motivating example and a sufficiently general linear fractional representation (LFR) model for the interconnected systems considered in this work. The control problem formulation is updated for the distributed control setting. We then review the SI framework for structured controller synthesis as a way to render the synthesis problem trackable and describe a novel extension to static and fixed-order output feedback via a specific state-space realization. Subsequently, for the common class of linear parameter-varying uncertainties, compatibility of this LMI approach with full-block multipliers is shown. For uncertainties requiring dynamic multipliers, such as time delays, we then explore a novel frequency domain condition [CS20] for the IQC framework suited to structured control problems. We discuss how to formulate the $\mathbf{H}_2$-performance objective in this setting using a multiplier derived from [Pag99] and give details for efficient implementation. Finally, the structured synthesis approaches are compared in terms of qualitative factors such as numerical sensitivity, speed, applicability and conservatism.

Chapter 5 introduces the motivating application of this work, the laser-based synchronization (LbSync) systems at the free-electron lasers (FELs) Free-Electron Laser in Hamburg (FLASH) and European XFEL. We review the FEL principle of operation and from that derive the requirements and challenges for the synchronization system. The DESY developed pulsed optical fibre timing synchronization is discussed by short expositions on core components. Then, using a model of the specific system at European XFEL, we perform numerical studies on achievable performance and performance-limiting factors, and verify the numerical results by comparing model based simulation data to measurements from the European XFEL system.

In the final Chapter 6, we draw a conclusion about the research progress achieved with this thesis and share our view on future work in this field.

The relation of the individual chapters is depicted in Figure 1.2. A reader interested in the theory only may skip Chapter 5 as it contains mostly results specific to the DESY LbSync system implementation. Vice versa, readers coming from a physics and FEL background may skip the theory chapters Chapters 3 and 4. Practitioners wishing to get a better understanding of the control theoretic implications on the subsystem system design are however recommended to include Chapter 3.

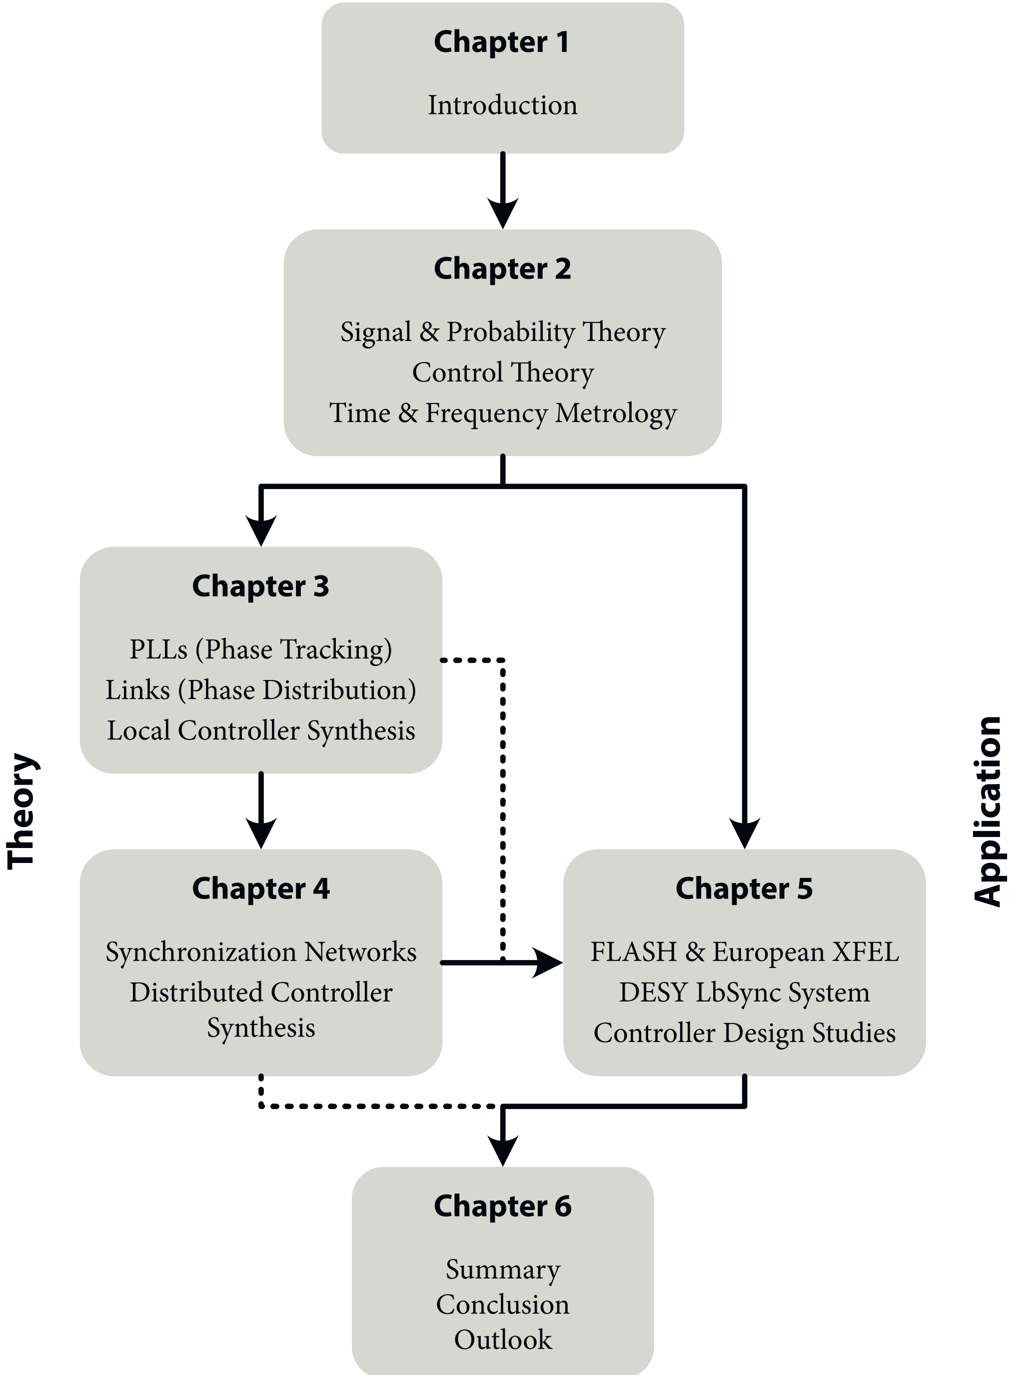

Figure 1.2: Relational diagram of the topics and chapters covered in this work.

In the following, we summarize the key contributions of this work towards the understanding and development of timing synchronization systems and robust structured optimal control problems:

1. A unifying abstraction for timing synchronization systems is developed in the form of a graph description of the network topology (Definition 4.3), an associated LFR model description (Section 4.2.1), and an optimal control problem formulation (Problem 4.5). This separates the details of heterogeneous implementations in application from the common performance analysis and optimization techniques.

2. The PLL and the stabilized transmission link systems are modelled and analysed in detail from a control perspective (Sections 3.1 and 3.3), and generalized to LTI singular $\mathbf{H}_2$-optimal control problems. Lacking implementable optimal solutions, LMI Lyapunov shaping methods are investigated to explicitly design the suboptimal solution in Section 3.2.2. For non-negligible time delays, suitable IQC multipliers are evaluated (Sections 2.3.4 and 3.4.3). The introduction of the IQC framework also allows for robust designs against a wide range of other, possibly nonlinear time-varying uncertainty.

3. Convex conditions for the synthesis of fixed-order controllers with structural constraints are given, which contain the tightest convex restriction of the underlying non-convex problem. This is achieved by extending the SI framework [Fur+20] based on separable Lyapunov functions to static and fixed-order output feedback settings in Proposition 4.32 and Proposition 4.34, respectively. The results, also reported in [Sch+21], are based on a state-output transformation (Definition 4.29), which is shown to be always available (Lemma 4.31).

4. Reliable and safe operation in the face of uncertainty in large system models and degradation of high-fidelity devices is achieved by unifying the SI framework for structured control problems and the S-procedure with full-block multipliers [Sch00] into a set of robust synthesis LMIs for linear time-varying parametric uncertainty (Proposition 4.41). The aforementioned extension to static and fixed-order output feedback bears no consequence of the compatibility of these two concepts. In numerical examples, we show that the combination of the two approaches adds no excess conservatism. These results were reported in [Sch+22].

5. For even broader classes of uncertainty, such as uncertain time delays, an alternate approach based on the IQC framework is developed. By adapting a frequency domain multiplier condition [Pag99] for $\mathbf{H}_2$-performance in a white-noise rejection setting to the novel frequency domain IQC stability condition from [CS20], nonsmooth $\mathbf{H}_\infty$-synthesis methods [AN06b] can then be used to solve the structured robust optimal control problem with local certificates of optimality. The result is stated in Corollary 4.51. Additionally, we derive an analytical result allowing for efficient numerical implementation of the multiplier condition (Lemma 4.53) and study the approach using extensive numerical examples. Both theoretical results have been reported in [SEW23].

6. A collection of numerical optimal robust distributed controller studies for the laser-based synchronization system for European XFEL is provided in Section 5.3, and in Section 5.2.8 the used model is validated by comparison of simulation results to measurements from the real system. From these results, we conclude in Section 5.3.6 the performance headroom and bottlenecks in the current system design and propose a set of changes to improve the system performance towards attosecond stability.

Chapter 2

Preliminaries

2.1 Notation

Let $\mathbf{R}$ and $\mathbf{C}$ denote the fields of real and complex numbers, respectively. Let $\mathbf{F}$ denote either field. $\mathbf{F}_-$ and $\mathbf{F}_+$ are the open left- and right-half lines or planes, respectively. $\bar{\mathbf{F}}_-$ and $\bar{\mathbf{F}}_+$ denote their closed complements. For $n \in \mathbf{N}$, $\mathbf{F}^n$ denotes the set of ordered n-tuples with elements in $\mathbf{F}$, i.e.

$$\Big\{(x_1, x_2, ..., x_n) \mid x_i \in \mathbf{F} \;\; \forall\, i = 1, ..., n\Big\}. \tag{2.1}$$

For a complex scalar $z \in \mathbf{C}$ let $\bar{z}$ denote its complex conjugate. Similarly, for a complex-valued matrix $A \in \mathbf{C}^{m \times n}$, let A^{H} denote the conjugate transpose of A. For any matrix A, A_{ij} denotes the element in row i, column j. For positive (semi-)definite matrices we write $A \succ 0$ ($A \succeq 0$) and for negative (semi-)definite matrices $A \prec 0$ ($A \preceq 0$), respectively. Additionally, let, for example, $A \succ B$ be understood as $A - B \succ 0$ for appropriately sized matrices. Further, for any real or complex matrix A, $\bar{\sigma}(A)$ ($\underline{\sigma}(A)$) signifies the maximum (minimum) singular value of A. Similarly, if A is square, $\bar{\lambda}(A)$ ($\underline{\lambda}(A)$) denotes the largest (smallest) eigenvalue.

Let the set of continuous functions bounded on the interval $[a, b]$ be denoted $\mathcal{C}[a, b]$. For repeatedly used functions and operators, arguments will be omitted when they are clear from the context, to keep the exposition concise. Similarly, we will state dimensions of vectors and matrices only when their size is not clear from the context or to support clarity. The bullet symbol is used as placeholder for arguments of arbitrary size that need not be named explicitly. For example $f(\bullet)$ and $A^{\bullet \times 2}$ are a function of one argument and a matrix with two columns and arbitrary number of rows, respectively. Fourier transformed functions are denoted using the hat symbol, e.g. $\hat{v}$.

2.2 Elements from Signal and Probability Theory

2.2.1 Signals and Operators

In this work, a *signal* is a function of time conveying information about a process. The domain of a *continuous-time* signal $v(t)$ is a real scalar interval $T \subseteq \mathbf{R}$. A *discrete-time* signal $v[k]$ is defined analogously for intervals on $\mathbf{Z}$. Signals, like other functions (and later operators), are defined on respective *vector spaces* that set the scope for which the provided theory (in the form of theorems) holds.

A summary of relevant definitions and results related to vector spaces is given in Appendix A.1 based on the works [Mar18; ZDG96].

We are frequently interested in comparing signals based on their size, i.e. a scalar quantity, which may hold physical meaning such as energy or power. The most useful class of such sizes are so-called *norms*, as they add exploitable structure to the space of signals. The most prominent example is the Euclidean norm, the distance covered by a vector x in the standard Euclidean vector space $\mathbf{R}^n$ given by $\|x\|_2 = \sqrt{x_1^2 + \cdots + x_n^2}$.

An *inner product* adds even more exploitable structure to a vector space. Through it, we arrive at the important class of *Hilbert spaces*, which gives access to important results in Fourier theory and geometrical concepts such as orthogonality. Indeed, take for example the dot product, defining the inner product on $\mathbf{R}^n$. Then two vectors $x, y \in \mathbf{R}^n$ are orthogonal if and only if $\langle x, y\rangle = x \cdot y = 0$. In this case, the inner product also induces the Euclidean distance norm, as $\langle x, x\rangle = \|x\|_2$.

Some signal spaces relevant to this work include

$\mathbf{L}_2^n[0,\infty)$ the Hilbert space of all square integrable signals in $\mathbf{R}^n$ defined on the positive real axis. We will assume for $v \in \mathbf{L}_2^n[0,\infty)$, that $v(t) = 0 \;\; \forall\, t < 0$.

$\mathbf{L}_{2_e}^n[0,\infty)$ the extended superset of $\mathbf{L}_2$ including all $\mathbf{R}^n$-valued signals square integrable over finite intervals.

Operators are functions that map elements of one vector space into another. We can thus describe the (multivariable) process driven by some signal $u \in \mathbf{L}_2^m$ and outputting another signal $v \in \mathbf{L}_2^n$ by a suitable operator $H\colon \mathbf{L}_2^m \to \mathbf{L}_2^n$, $v = H(u)$. Notably, the composition $H_1 H_2$ and linear combination $\alpha H_1 + \beta H_2$, $\alpha, \beta \in \mathbf{R}$ of two (size and space) compatible operators H_1 and H_2 are also operators, respectively. The important class of LTI system operators are discussed in detail in Section 2.3.1. Operators are frequently distinguished by the important properties of linearity (satisfying additivity and homogeneity), causality (respecting physical cause-effect laws), and boundedness (having finite gain induced by the input-output signal norms). Like signals, operators live in complex (matrix) function spaces, some of which will be defined next.

$\mathbf{L}_\infty^{m\times n}(i\mathbf{R})$ the Banach space of all complex matrix-valued functions essentially bounded on the imaginary axis.

$\mathbf{L}_2^{m\times n}(i\mathbf{R})$ the Hilbert space of all complex matrix-valued functions with bounded integral below the imaginary axis,

$$\int_{-\infty}^{\infty} \mathrm{trace}(F(i\omega)^{\mathsf{H}} F(i\omega)) \;\mathrm{d}\omega < \infty \;\; \forall\, F \in \mathbf{L}_2^{m\times n}(i\mathbf{R}).$$

$\mathbf{H}_\infty^{m\times n}(i\mathbf{R})$ the subspace of $\mathbf{L}_\infty(i\mathbf{R})$ consisting of functions analytic in the open right-half plane. For $F \in \mathbf{H}_\infty(i\mathbf{R})$ we have

$$\|F\|_\infty := \sup_{\omega\in\mathbf{R}} \bar{\sigma}[F(i\omega)].$$

$\mathbf{H}_2^{m\times n}(i\mathbf{R})$ the subspace of $\mathbf{L}_2(i\mathbf{R})$ consisting of functions analytic in the open right-half plane. After applying the maximum modulus principle, its inner-product induced norm reduces to

$$\|F\|_2 := \sqrt{\frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(F(i\omega)^{\mathsf{H}} F(i\omega)) \;\mathrm{d}\omega}$$

Unless it helps clarity, we will assume that the domain of a space is clear from the context and omit it from the notation. For the frequency-domain spaces, we mostly consider their real-rational subspaces, $\mathbf{RL}_\infty$, $\mathbf{RL}_2$, $\mathbf{RH}_\infty$, $\mathbf{RH}_2$, respectively, as they admit a transfer function form, which will be introduced in Section 2.3.1.

2.2.2 Stochastic Processes and Spectral Estimation

So far, we have made no assumption on whether the signal is (partially) predictable, i.e. following a deterministic law with known parameters, or is *random*, i.e. governed by sufficiently many unknown quantities and relations. Suppressing the impact of random influences on a system is in many cases the main motivation for automatic feedback control, and in the absence of any knowledge about the laws of the process generating a signal, other, statistical, properties need to be considered to design useful control systems and derive meaningful results. In the context of this work, stochastic processes provide the theoretical foundation for (generally non-white) phase noise in oscillators, cf. Section 2.4.2, and control theoretic properties of systems driven by white noise. The theory provided in this section is primarily based on references [PP02; Lap17]. Essential, mathematically rigorous definitions are provided in Appendix A.2.

An *experiment* involving chance is described by its *probability space* $(\Omega, \mathcal{F}, P)$. A stochastic process $\mathbf{x}(t,\xi)$ defined on such a probability space assigns to each experimental outcome $\xi \in \Omega$ a function, in our case a signal in a given space, e.g. $\mathbf{L}_{2_e}[0,\infty)$. For variable t, ξ, $\mathbf{x}(t,\xi)$ describes the ensemble of possible realizations, whereas, for a fixed outcome ξ, $\mathbf{x}(t,\xi)$ describes a fixed realization, also called sample function. We will generally abbreviate this as $x(t)$. On the other hand, for fixed time t, $\mathbf{x}(t,\xi)$ describes a *random variable*, i.e. a number with given probability distribution. Going forward, the dependence on ξ will be assumed and dropped from the notation.

We are primarily interested in stochastic processes that are (wide-sense) stationary. Strict stationarity indicates that the statistical properties of a process are invariant to a shift in time. Wide-sense stationary processes satisfy the weaker conditions of having constant mean and having an autocorrelation function that only depends on the time shift $\tau = t_2 - t_1$. In order to be practical, wide-sense stationary processes however are additionally required to be of finite variance. Wide-sense stationary processes allow for the definition of power spectra, a tool analogous to the Fourier transform for deterministic signals. Specifically, let $\mathbf{x}(t)$ be a wide-sense stationary process with autocorrelation function $R_{\mathbf{x}}(\tau)$. Then, by the Wiener-Khinchin theorem, the *power spectral density (PSD)* $S_{\mathbf{x}}(f)$ of $\mathbf{x}(t)$ is taken as $\hat{R}_{\mathbf{x}}(f)$, i.e. the Fourier transform

$$S_{\mathbf{x}}(f) = \hat{R}_{\mathbf{x}}(f) = \int_{-\infty}^{\infty} R_{\mathbf{x}}(\tau) e^{-i2\pi f\tau} \, \mathrm{d}\tau. \tag{2.2}$$

The PSD is of particular relevance for engineering practice, as it describes the allocation of noise power generated by the stochastic process to specific frequency intervals – usually normalized to a width of 1 Hz – which in turn shapes the frequency response of the optimal feedback controller for systems impacted by this noise source. A stochastic process whose PSD is constant across frequency is called *white*. An example is the additive white Gaussian noise model frequently used in engineering practice. We however stress the fact that a Gaussian distribution does not necessitate a white PSD, nor vice versa. In electronics, noise PSDs scaling with $1/f$, called *pink* noise, are frequently encountered. The Wiener process, also called *Brownian* noise or Brownian motion shows an $1/f^2$ dependence and occurs when white noise is integrated.

Remark 2.1. Note that, by virtue of the Fourier transform, the PSD $S_{\mathbf{x}}(f)$ is defined for positive as well as negative frequencies. For real stochastic processes, the spectrum is symmetric and engineers often prefer to consider the so called *one-sided* (or single-sideband (SSB)) PSD $S'_{\mathbf{x}}(f)$ defined as

$$S'_{\mathbf{x}}(f) = \begin{cases} 0 & \text{for } f < 0 \\ S_{\mathbf{x}}(f) & \text{for } f = 0 \\ 2S_{\mathbf{x}}(f) & \text{for } f > 0. \end{cases} \tag{2.3}$$

Doubling the PSD for positive frequencies ensures that integration of the single-sided PSD yields consistent total power in the sense of Plancherel's theorem (commonly and interchangeably referred to as Parseval's theorem).

In anticipation of Section 2.3.1 and Section 3.1.2, we state the following result on filtering a stochastic process with a linear system.

Lemma 2.2. *The response of an LTI system with frequency response $H(f)$ and corresponding impulse response $h(t)$ to a stochastic process $\mathbf{x}(t)$ with PSD $S_{\mathbf{x}}(f)$ is given by the stochastic process*

$$\mathbf{y}(t) = \int_{-\infty}^{\infty} \mathbf{x}(t-\tau)h(\tau)\ \mathrm{d}\tau, \tag{2.4}$$

with PSD $S_{\mathbf{y}}(f)$ given by

$$S_{\mathbf{y}}(f) = S_{\mathbf{x}}(f)|H(f)|^2. \tag{2.5}$$

While theoretically well-defined, an exact PSD can rarely be computed in practice because the "true" autocorrelation function is not accessible. Indeed, to obtain the full autocorrelation function, one would need to observe a stochastic process for an infinite amount of time. We must thus resort to estimation theory to obtain an approximate PSD from observations (sample functions) of the random process. Processes for which this is possible are called *ergodic*, meaning that time averages of observations can be used to reliably substitute the expectation operator in the computation of various statistical properties (such as the autocorrelation function) in order to obtain a valid estimate of that parameter. The perhaps simplest estimator for the PSD is obtained from direct application of the convolution theorem and computing the discrete Fourier transform (DFT) over the finite sequence $x[k]$, equidistantly sampled from a realization of $\mathbf{x}(t)$, yielding the so-called *periodogram*. Note that the periodogram is not a consistent estimator, and that the standard deviation of the estimated PSD at any given frequency bin is equal to the expected value. A better estimate is obtained by averaging multiple periodograms (Bartlett's method) or averaging periodograms obtained from overlapping and windowed segments of an observation (Welch's method). We direct the reader to references [HRS02; PTV07] for more information and stress the importance of the Nyquist-Shannon sampling theorem and the choice of windowing function in Welch's method with respect to spectral leakage.

2.3 Elements from System and Control Theory

2.3.1 Linear Time-Invariant Systems

Continuous LTI systems are the fundamental modelling concept used in this work. We develop some results central to the following work along references [ZDG96; Jön01; SP09; DFT92; ÅM21; DP00; Hes18; AM06; FPE19].

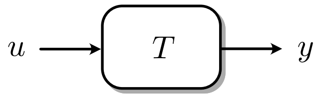

Figure 2.1: Block diagram representation for system Σ with transfer matrix T.

Governing the relation of the input, output and state of a system, any such LTI system Σ can be described using dynamic equations of the form

$$\Sigma\colon \quad \begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) + Du(t), \end{aligned} \tag{2.6}$$

with $x \in \mathbf{R}^{n_x}$ being the state, $u \in \mathbf{R}^{n_u}$ the input and $y \in \mathbf{R}^{n_y}$ the output vector. Additionally, let the initial condition be $x(t_0) = x_0$. We will sometimes refer to the real system matrices A, B, C, D as the state (or just system), input, output and feedthrough matrix, respectively. Together with (2.6), the system matrices constitute the *state-space* representation of the LTI system Σ. Using the Laplace transform, we may find an alternate representation for Σ with zero initial condition in the frequency domain, namely the *transfer (function) matrix* $T(s)$, given by

$$T(s) = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right], \tag{2.7}$$

where the right-hand side, in the style of [ZDG96], denotes

$$\left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right] := C\left(sI - A\right)^{-1} B + D, \tag{2.8}$$

while

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \text{ or } \left[\begin{array}{c:c} A & B \\ \hdashline C & D \end{array}\right]$$

continues to denote a block matrix.

The inverse Laplace transform $K(t) = \mathcal{L}^{-1}\left\{T(s)\right\}$ is the *impulse response* matrix of the system. As such, T further defines an operator by convolution with its impulse response:

$$y(t) = (Tu)(t) = \int_{-\infty}^{\infty} K(t-\tau)u(\tau)\ \mathrm{d}\tau, \tag{2.9}$$

which we will commonly write in the condensed form

$$y = Tu. \tag{2.10}$$

Many problems in control theory involve passive systems, i.e. broadly speaking, those which do not generate energy. Systems of this kind have positive real transfer functions.

Definition 2.3 (Positive Real Function [ZDG96])**.** A square transfer matrix $G(s) \in \mathbf{RH}_\infty^{n\times n}$ is said to be (strictly) positive real if $G(i\omega) + G(i\omega)^{\mathsf{H}} \succeq 0$ $(G(i\omega) + G(i\omega)^{\mathsf{H}} \succ 0)$ for all $\omega \in \mathbf{R}$.

Physical systems, i.e. those which are causal and have a proper transfer function will commonly be denoted G, while T will be used for general transfer functions when studied in a theoretical context.

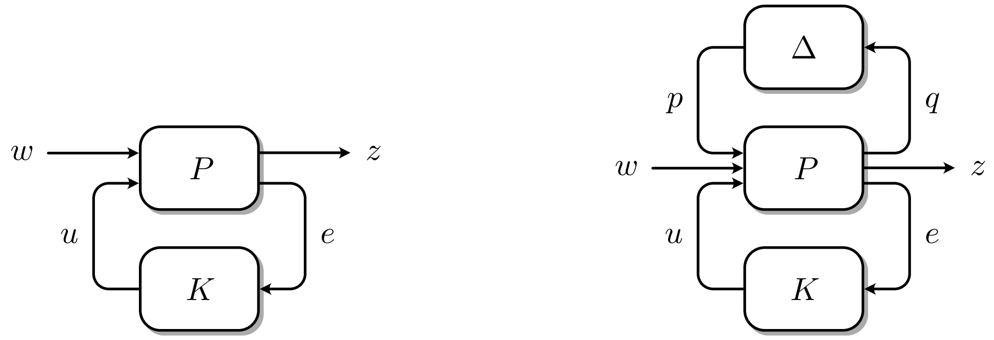

(a) Nominal system (b) Uncertain system

Figure 2.2: Linear fractional representation of a control feedback loop with performance.

System Interconnection

Starting from a (generalized) *plant* P, i.e. a given and fixed system, a central concern in control theory is the design of a controller, i.e. a variable system K, that through interconnection with P achieves a desired overall system behaviour, e.g. expressed in terms of signal specifications on some outputs of P. By appropriately modelling the internal structures of plant and controller, any such interconnection can be represented as depicted in Figure 2.2a, and given by the system of equations

$$\begin{bmatrix} z \\ e \end{bmatrix} = P \begin{bmatrix} w \\ u \end{bmatrix} \tag{2.11a}$$

$$u = Ke, \tag{2.11b}$$

where the inputs and outputs of P are partitioned into exogenous inputs w, controlled inputs u, regulated outputs z and measured outputs e. In this configuration, we say that the controller K closes the *control channel* $e \to u$. The remaining inputs w are outside the immediate control of the system designer and comprise disturbance sources acting on the system as well as e.g. reference signals to be tracked. In contrast to the measured signals, which can and need to be actually observed for the operation of the feedback loop, the regulated output can be virtual metrics exclusive to the mathematical system model. In this work, as opposed to G, which is primarily used for stability analysis, P is used in performance analysis and controller design formulations and commonly includes artificial dynamics (e.g. weightings) that are physically motivated or heuristically selected to achieve the desired closed-loop behaviour.

Models developed for real world systems are always subject to (at least small) structural or parameter uncertainty originating from limited system knowledge or imprecise system identification. In cases where the uncertainty cannot be neglected, robust controllers guaranteeing some goal, e.g. system stability, for a range of possible realizations of the uncertainty, are desired. In the same fashion as before, where the system was separated into fixed plant and variable controller, it is common in robust control theory to separate an uncertain generalized plant model into a nominal and an uncertain part. The latter is expressed in the form of a not necessarily LTI but causal operator Δ belonging to the set of all possible realizations $\mathbf{\Delta}$. Virtually all forms of uncertainty can be covered in this fashion.

This is depicted in Figure 2.2b and described by the system of equations extended from (2.11) as

$$p = \Delta q \tag{2.12a}$$

$$\begin{bmatrix} q \\ z \\ e \end{bmatrix} = P \begin{bmatrix} p \\ w \\ u \end{bmatrix} \tag{2.12b}$$

$$u = Ke, \tag{2.12c}$$

where the uncertainty Δ closes the uncertainty channel $q \to p$.

Interconnections of this form are known as LFRs of the closed system from exogenous inputs to outputs $w \to z$. The closed form is computed from the individual blocks using linear fractional transformations (LFTs).

Definition 2.4 (Linear Fractional Transformation [ZDG96])**.** Given a complex valued matrix $M \in \mathbf{C}^{(q_1+q_2)\times(p_1+p_2)}$ partitioned as

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}, \tag{2.13}$$

the *lower* LFT with a complementary complex matrix $\Delta \in \mathbf{C}^{q_2 \times p_2}$ is defined as

$$\mathcal{F}_l(M, \Delta) = M_{11} + M_{12}\Delta(I - M_{22}\Delta)^{-1} M_{21}. \tag{2.14}$$

Analogously, the *upper* LFT with a complementary complex matrix $\Delta \in \mathbf{C}^{q_1 \times p_1}$ is defined as

$$\mathcal{F}_u(M, \Delta) = M_{22} + M_{21}\Delta(I - M_{11}\Delta)^{-1} M_{12}. \tag{2.15}$$

The transformations are well-defined as long as $(I - M_{22}\Delta)$ and $(I - M_{11}\Delta)$ are invertible for all possible realizations $\Delta \in \mathbf{\Delta}$, respectively. In this case we say that the interconnection of M and Δ is well-posed.

Applying this to the LFR (2.12), one obtains the partially closed systems

$$M = \mathcal{F}_l(P, K), \tag{2.16}$$

the controlled nominal system, and

$$P(\Delta) = \mathcal{F}_u(P, \Delta), \tag{2.17}$$

the uncertain generalized plant. Closing either of the partially closed systems with the remaining block obtains the controlled uncertain system

$$M(\Delta) = \mathcal{F}_u(M, \Delta) = \mathcal{F}_l(P(\Delta), K). \tag{2.18}$$

2.3.2 Nominal Stability and Performance Analysis

Stability

Having established a formalism to model and describe dynamics of a given LTI system, we turn to analysing its characteristic properties. The central objective in control theory is to ensure that the regulated system is stable, meaning that, broadly speaking, whatever happens, the system state (and hence output) remains within a well-defined, bounded region around some desired point of operation or *equilibrium* (w.l.o.g. assumed to be the origin). Stability analysis for linear and nonlinear systems is a mature field and we refer the reader to references [DV75; Baş00; ZDG96] for an in depth treatment of various notions of stability. For the purpose of this work, stability in the sense of Lyapunov and bounded-input, bounded-output (BIBO) stability are central and shall be defined next.

Definition 2.5 (Lyapunov Stability [Hes18])**.** The LTI system (2.6) is said to be (in the sense of Lyapunov)

1. *(marginally) stable* or *internally stable* if, for every initial condition $x(t_0) = x_0 \in \mathbf{R}^{n_x}$, $t_0 \geq 0$ the homogeneous state response
$$x(t) = \mathrm{e}^{A(t-t_0)}x_0, \quad \forall\, t \geq 0$$
is uniformly bounded;
2. *globally asymptotically stable* if, in addition, for every initial condition $x(t_0) = x_0 \in \mathbf{R}^{n_x}$, $t_0 \geq 0$ the homogeneous state response satisfies $x(t) \to 0$ as $t \to \infty$;
3. *globally exponentially stable* if, in addition, there exist constants $c, \lambda > 0$ such that, for every initial condition $x(t_0) = x_0 \in \mathbf{R}^{n_x}$, $t_0 \geq 0$ the homogeneous state response satisfies
$$\|x(t)\| \leq c\,\mathrm{e}^{-\lambda(t-t_0)}\|x(t_0)\|, \quad \forall\, t \geq 0; \quad \text{or}$$
4. *unstable* if it is not marginally stable.

As such, Lyapunov stability is concerned with the unforced system response without inputs. On the other hand, the following definition of BIBO stability considers the case of applying inputs to a system with zero initial state:

Definition 2.6 (Bounded-input, bounded-output Stability [Hes18])**.** The system LTI (2.6) is said to be *(uniformly) bounded-input, bounded-output stable* if there exists a finite constant g such that, for every input $u(\cdot) \in \mathbf{L}_p$, its forced response from zero initial state

$$y(t) = \int_0^t C\mathrm{e}^{A(t-\tau)}Bu(\tau)\ \mathrm{d}\tau + Du(t)$$

satisfies

$$\sup_{t\in[0,\infty]} \|y(t)\| \leq g \sup_{t\in[0,\infty]} \|u(t)\|.$$

Remark 2.7. The choice of the norm in Definition 2.6 is irrelevant and only results in a different value of g [Hes18].

Remark 2.8. Definition 2.6 also applies to the more general class of linear time varying (LTV) systems [Hes18].

To arrive at the more general result of BIBO stability for arbitrary initial states, we state

Theorem 2.9 ([SP09])**.** *In the absence of any hidden unstable modes, an LTI system (2.6) is globally exponentially stable if and only if it is BIBO stable.*

In order to verify the stability property, a plethora of time and frequency stability criteria can be used, see e.g. [ÅM21; Hes18; SP09; SW17]. In general, Lyapunov stability is proven by finding a suitable Lyapunov function, i.e. an energy storage function monotonically decreasing along all system trajectories [ÅM21]. For LTI systems, such functions are fully parameterized by an according real symmetric Lyapunov matrix $\mathcal{X} = \mathcal{X}^\top \in \mathbf{S}^n$ yielding the Lyapunov function $v(x) := x^\top \mathcal{X} x$. As stated in the next theorem, this is equivalent to the common test that the system's poles lie within a suitable stability region, i.e., in continuous-time, the open left-half plane (LHP) [SW17].

Theorem 2.10 ([Hes18])**.** *The following statements are equivalent.*

- *The LTI system (2.6) is asymptotically stable.*
- *The LTI system (2.6) is exponentially stable.*
- *All eigenvalues of A have strictly negative real parts, in other words, A is* Hurwitz.
- *There exists a symmetric positive-definite matrix $\mathcal{X} \succ 0$ for which the following Lyapunov matrix inequality holds:*
$$A^\top \mathcal{X} + \mathcal{X} A \prec 0. \tag{2.19}$$

Remark 2.11. Recall that, by (2.8), the eigenvalues of A are identical to the roots of the denominator polynomial of the corresponding transfer function if the state-space realization is minimal. Both are also referred to as *poles* of the system (or transfer function).

The theory in Chapters 3 and 4 will make extensive use of the Lyapunov matrices and LMI based criteria.

Performance

Beyond stability, the goal of system or controller design is to achieve optimal performance in closed-loop operation. Again, depending on the application and its constraints, various performance criteria in the time and frequency domain may be considered. For the application considered in this work, the $\mathbf{H}_2$- and $\mathbf{H}_\infty$-performance metrics frequently used in modern control theory are crucial, see Section 3.2.1, where $\mathbf{H}_2$-performance arises as a natural choice for synchronization systems, cf. Sections 3.2.1 and 4.1. The $\mathbf{H}_2$ and $\mathbf{H}_\infty$-performance metrics are defined in terms of the norms of their respective Hardy spaces [ZDG96], see Appendix A.1.

Definition 2.12 (Hardy Space $\mathbf{H}_2$ [ZDG96])**.** $\mathbf{H}_2$ is a (closed) subspace of $\mathbf{L}_2(i\mathbf{R})$ with matrix functions $T(s)$ analytic in the open right-half plane (RHP) and square integrable on the imaginary axis including infinity. The corresponding inner-product induced norm is defined as

$$\|T\|_2^2 := \frac{1}{2\pi} \int_{-\infty}^{\infty} \text{trace}(T(i\omega)^{\mathsf{H}} T(i\omega))\ \mathrm{d}\omega. \tag{2.20}$$

Consistently, the subspace $\mathbf{RH}_2$ consists of all strictly proper and real rational stable transfer matrices.

Definition 2.13 (Hardy Space $\mathbf{H}_\infty$ [ZDG96])**.** $\mathbf{H}_\infty$ is a (closed) subspace of $\mathbf{L}_\infty(i\mathbf{R})$ with matrix functions $T(s)$ analytic in the open RHP and bounded on the imaginary axis. The corresponding norm is defined as

$$\|T\|_\infty := \sup_{\omega \in \mathbf{R}} \bar{\sigma}(T(i\omega)). \tag{2.21}$$

Consistently, the subspace $\mathbf{RH}_\infty$ consists of all proper and real rational stable transfer matrices.

The $\mathbf{H}_2$ and $\mathbf{H}_\infty$-norms have important interpretations in terms of the input-output characteristics of the operator defined by a transfer matrix element from either space. Specifically, by the Plancherel theorem, the $\mathbf{H}_2$-norm can be defined in terms of the system's impulse response energy. For $T \in \mathbf{H}_2^{m\times n}$ we have [Sto93; ZDG96]

$$\|T\|_2^2 = \sum_{i=1}^{m} \|Ke_i\|_2^2, \tag{2.22}$$

where $K = \mathcal{L}^{-1}\{T\}$ is the impulse response matrix and e_i is the i^{th} standard basis vector of $\mathbf{R}^m$.

The $\mathbf{H}_2$-norm further admits a stochastic interpretation, namely the expected energy of the response to an input impulse of random direction or the output variance of the system when excited with unitary white noise on all input channels [Sto93],

$$\|T\|_2^2 = \lim_{t\to\infty} \mathrm{E}\left\{\frac{1}{t}\int_0^t z^\top(\tau)z(\tau)\ \mathrm{d}\tau\right\}. \tag{2.23}$$

The $\mathbf{H}_2$-norm therefore expresses a metric for the system's ability to suppress or attenuate disturbances across all frequencies. It is from this interpretation that the $\mathbf{H}_2$-theory yields its equivalence to the classical LQG problem.

On the other hand, the $\mathbf{H}_\infty$-norm expresses the peak amplification of any particular frequency across the spectrum. This fact is frequently exploited in robust control, ensuring stability margins by means of the small gain theorem [DV75], however is also sometimes criticized from an engineering perspective as too focussed on the worst-case scenario [Cha+20]. Indeed, optimizing for $\mathbf{H}_\infty$-performance does not easily allow trading-off slightly higher sensitivity at one frequency for significantly reduced sensitivity over another frequency range[1].

It can be shown that the $\mathbf{H}_2$-norm of a system can be computed using either its controllability or observability Gramian, which is commonly done in practice, while the $\mathbf{H}_\infty$-norm conventionally requires an iterative bisection procedure using an associated Hamiltonian matrix [ZDG96]. In view of the theory developed in this work, we again provide, albeit closely related, more relevant LMI conditions for verifying a given $\mathbf{H}_2/\mathbf{H}_\infty$-performance level $\gamma > 0$ that can be optimized using semi-definite programming to achieve tightness.

Theorem 2.14 ([SW17]). *Consider the LTI system (2.6) with A Hurwitz and $D = 0$ such that the corresponding transfer function is $T(s) \in \mathbf{RH}_2^{\bullet\times\bullet}$. Let $\gamma > 0$. The following statements are equivalent:*

1. $\|T\|_2 < \gamma$.
2. *There exists a solution $X = X^\top \succ 0$ to the LMIs*

$$AX + XA^\top + BB^\top \prec 0, \quad \text{trace}(CXC^\top) < \gamma^2.$$

3. *There exists a solution $Y = Y^\top \succ 0$ to the LMIs*

$$A^\top Y + YA + C^\top C \prec 0, \quad \text{trace}(B^\top YB) < \gamma^2.$$

[1]To alleviate this, carefully designed weighting functions may be included in the $\mathbf{H}_\infty$-problem

(a) Stability analysis: System is disturbed by exogenous noise.

(b) Stability and performance analysis: An additional performance output is considered.

Figure 2.3: LFRs of an LTI system affected by uncertain dynamics.

Theorem 2.15 ([SW17]). *Consider the LTI system (2.6) with A Hurwitz, i.e. $T \in \mathbf{H}_\infty^{\bullet\times\bullet}$. Let $\gamma > 0$. The following statements are equivalent:*

1. $\|T\|_\infty < \gamma$.
2. *There exists a solution $X = X^\top \succ 0$ to the LMI*

$$\begin{bmatrix} I & 0 \\ A & B \\ \hline 0 & I \\ C & D \end{bmatrix}^\top \left[\begin{array}{cc|cc} 0 & X & 0 & 0 \\ X & 0 & 0 & 0 \\ \hline 0 & 0 & -\gamma I & 0 \\ 0 & 0 & 0 & \frac{1}{\gamma} I \end{array}\right] \begin{bmatrix} I & 0 \\ A & B \\ \hline 0 & I \\ C & D \end{bmatrix} \prec 0. \tag{2.24}$$

2.3.3 Robust Stability and Performance Analysis

Integral Quadratic Constraints

Consider a nominally stable LTI plant $G \in \mathbf{RH}_\infty^{n_q \times n_p}$ affected by uncertain dynamics, represented through interconnection with an uncertain but bounded and causal operator $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ as depicted in Figure 2.3a and given by

$$p = \Delta(q) \tag{2.25a}$$
$$q = Gp + q_0, \tag{2.25b}$$

where q_0 is an external disturbance. The immediate methodological challenge in robust control is dealing with the unknown realization of the uncertainty block Δ affecting the plant. Many early approaches to robust control have since been unified in a joint framework, based on the concept of IQCs. The central idea is that even though the uncertain operator Δ cannot be used explicitly in formulae, the set of possible realizations $\boldsymbol{\Delta}$ can be described by bounding the effect any such $\Delta \in \boldsymbol{\Delta}$ can have on an applied input signal. One useful representation of such a bound is an IQC, relating input and output of the operator Δ with an explicit multiplier $\Pi\colon i\mathbf{R} \to \mathbf{C}^{(n_q+n_p)\times(n_q+n_p)}$ which encodes information about the uncertainty.

This work relies on IQC theory as laid out by Megretski and Rantzer [MR97], based on preliminary work by Yakubovich. They define IQCs as follows using a frequency domain criterion, but an equivalent and just as useful time domain equivalent exists as well.

Definition 2.16 (Integral Quadratic Constraint [MR97; VSK16])**.** Two signals $q \in \mathbf{L}_2^{n_q}$ and $p \in \mathbf{L}_2^{n_p}$ are said to satisfy the IQC defined by a bounded and Hermitian-valued multiplier $\Pi\colon i\mathbf{R} \to \mathbf{C}^{(n_q+n_p)\times(n_q+n_p)}$ if

$$\mathcal{I}(\Pi, q, p) := \int_{-\infty}^{\infty} \begin{bmatrix} \hat{q}(i\omega) \\ \hat{p}(i\omega) \end{bmatrix}^{\mathsf{H}} \Pi(i\omega) \begin{bmatrix} \hat{q}(i\omega) \\ \hat{p}(i\omega) \end{bmatrix} \, \mathrm{d}\omega \geq 0. \tag{2.26}$$

By extension, we say that a causal and bounded operator $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ satisfies the IQC defined by Π if

$$\mathcal{I}(\Pi, q, \Delta(q)) \geq 0 \;\; \forall\, q \in \mathbf{L}_2^{n_q}. \tag{2.27}$$

The following summary of the IQC robust stability framework is based on references [VSK16; MR97].

Stability by Integral Quadratic Constraints

An already renowned stability theorem based on IQCs has been derived by Megretski and Rantzer based on an elegant proof by induction:

> "First, the system is considered as depending on a parameter $\alpha \in [0, 1]$, such that stability is obvious for $\alpha = 0$, while $\alpha = 1$ gives the system to be studied. Then, the IQCs are used to show that as α increases from zero to one, there can be no transition from stability to instability." (Megretski & Rantzer [MR97])

Theorem 2.17 (IQC Stability [MR97])**.** *Let $G \in \mathbf{RH}_\infty^{n_q\times n_p}$ be a stable transfer function and $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ a causal and bounded operator. The feedback interconnection (2.25) of G and Δ is stable if*

1. *for all $\alpha \in [0, 1]$ the interconnection is well-posed for $\alpha\Delta$ in place of Δ,*
2. *for all $\alpha \in [0, 1]$ the IQC defined by some $\Pi = \Pi^{\mathsf{H}} \in \mathbf{RL}_\infty^{(n_q+n_p)\times(n_q+n_p)}$ is satisfied for $\alpha\Delta$ in place of Δ, and*
3. *the frequency domain inequality (FDI)*

$$\begin{bmatrix} G(i\omega) \\ I \end{bmatrix}^{\mathsf{H}} \Pi(i\omega) \begin{bmatrix} G(i\omega) \\ I \end{bmatrix} \prec 0 \;\; \forall\, \omega \in \mathbf{R} \cup \{\infty\} \tag{2.28}$$

is satisfied.

To apply this theorem for *robust* stability analysis, one then proceeds in the following fashion. First, preliminary knowledge and assumptions about the uncertainty Δ are used to define a set of all possible realizations $\boldsymbol{\Delta}$. Here, a tactical assumption has to be made in view of the employed proof by induction:

Assumption 2.18. $\Delta \in \boldsymbol{\Delta}$ implies that $\alpha\Delta \in \boldsymbol{\Delta} \;\; \forall\, \alpha \in [0, 1]$.

Remark 2.19. The linear scaling $\alpha\Delta$ may be replaced by a more generic parameterization $\alpha \circ \Delta$ satisfying certain properties detailed in [JG01]. This will be relevant when discussing time-delay uncertainties, but we stick to the linear parameterization unless otherwise necessary.

Exploiting the structure of $\boldsymbol{\Delta}$, a family of multipliers $\boldsymbol{\Pi} \subset \mathbf{RL}_\infty^{(n_q+n_p)\times(n_q+n_p)}$ is then designed such that the IQC defined by any $\Pi \in \boldsymbol{\Pi}$ is satisfied for any $\Delta \in \boldsymbol{\Delta}$. From this, a result on robust stability analysis follows:

Corollary 2.20 (Robust IQC Stability [VSK16]). *Let $G \in \mathbf{RH}_\infty^{n_q \times n_p}$ be a stable transfer function and $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ a causal and bounded operator. Assuming*

1. *for all $\Delta \in \mathbf{\Delta}$ the interconnection (2.25) of G and Δ is well-posed, and*

2. *for all $\Delta \in \mathbf{\Delta}$ and for all $\Pi \in \mathbf{\Pi}$ the IQC (2.27) holds true,*

then the interconnection G – Δ is robustly stable if there exists one $\Pi \in \mathbf{\Pi}$ such that the FDI (2.28) is satisfied.

For many common types of uncertainty (e.g. uncertain time-invariant parameters, LTI dynamics, arbitrarily fast or rate-bounded LTV parameters, time delays, ...), suitable families of multipliers can be found in the literature, see e.g. [MR97; JG01; JS02; KL04; KR07; VSK16]. Naturally, one would like the multiplier family to be sufficiently large in order to increase the chance of finding a suitable member satisfying the IQC FDI (2.28). Whether the robust stability test succeeds or not is then a question of how large the robust stability margins of the system are relative to how tightly the uncertainty can be bounded by the multiplier satisfying the IQC FDI.

In order to find tighter multipliers, the following lemma is frequently exploited.

Lemma 2.21 ([VSK16]). *Assume $\alpha\Delta$, $\alpha \in [0, 1]$ satisfies the IQC defined by every multiplier $\Pi_1, \ldots, \Pi_k$. Then $\alpha\Delta$ also satisfies the IQC defined by any conic combination $\sum_{i=1}^k \lambda_i \Pi_i, \lambda_i \geq 0$.*

Further, the robustness analysis can often be rendered less conservative by considering frequency weighted multipliers, which is explored further in Section 2.3.4. This is analogous to the earlier concept of D/G-scales [YP93; PD93].

For systems affected by multiple uncertainties $\Delta_1, \ldots, \Delta_k$ with $\Delta_i \in \mathbf{\Delta}_i$, $i = 1, \ldots, k$, one easily finds that for $q = [q_1, \ldots, q_k]^\top$ and $\mathcal{I}(\Pi_i, q_i, \Delta_i(q_i)) \geq 0$ satisfied for all $q_i \in \mathbf{L}_2^{n_{q_i}}$, the diagonally augmented multiplier

$$\Pi = \operatorname{daug}(\Pi_1, \ldots, \Pi_k) := \begin{bmatrix} \operatorname{diag}(\Pi_{1,11}, \ldots, \Pi_{k,11}) & \operatorname{diag}(\Pi_{1,12}, \ldots, \Pi_{k,12}) \\ \operatorname{diag}(\Pi_{1,21}, \ldots, \Pi_{k,21}) & \operatorname{diag}(\Pi_{1,22}, \ldots, \Pi_{k,22}) \end{bmatrix} \tag{2.29}$$

satisfies the composite IQC

$$\mathcal{I}(\Pi, q, \Delta(q)) \geq 0 \ \ \forall\, q \in \mathbf{L}_2^{n_{q_i}},$$

where Π_i is partitioned according to $[q_i, \Delta_i(q_i)]^T$ and $\Delta(q) = [\Delta_1(q_1), \ldots, \Delta_k(q_k)]^\top$ [VSK16]. Augmented multipliers for repeated uncertainties (i.e. the same uncertain parameter occurring multiple times in the LFR) are easily constructed using the Kronecker product and the identity matrix of appropriate size.

The search for a suitable multiplier $\Pi \in \mathbf{\Pi}$ satisfying FDI (2.28) on the entire imaginary axis is rendered computationally feasible by first factorizing the transfer functions belonging to the multiplier family $\mathbf{\Pi}$ into a fixed dynamic basis $\Psi \in \mathbf{RH}_\infty^{\bullet \times (n_q + n_p)}$ and static projection $\Theta \in \mathbf{\Theta} \subset \mathbf{S}$ such that $\mathbf{\Pi} = \{\Psi^\mathsf{H} \Theta \Psi \mid \Theta \in \mathbf{\Theta}\}$. The fixed basis function Ψ is then absorbed in the system to be analysed while the search over the infinite dimensional FDI test is reformulated into a tractable optimization problem on Θ. The latter is conventionally achieved by applying the Kalman-Yakubovich-Popov (KYP) lemma in order to obtain an equivalent LMI condition, which poses a tractable semi-definite program (SDP) problem, for which efficient solvers exist [MR97; VSK16].

Lemma 2.22 (Kalman-Yakubovich-Popov Lemma [VSK16]). *Let $Y \in \mathbf{S}^{\bullet}$ and let $G \in \mathbf{RL}_{\infty}^{\bullet \times \bullet}$ admit realization (A, B, C, D) with $A \in \mathbf{R}^{n_x \times n_x}$ and $\operatorname{eig}(A) \cap \mathbf{C}_0 = \emptyset$. Then the following statements are equivalent:*

1. $G(i\omega)^{\mathsf{H}} Y G(i\omega) \prec 0 \;\; \forall\, \omega \in \mathbf{R} \cup \{\infty\}$.
2. *There exists a matrix $X \in \mathbf{S}^{n_x}$ such that*

$$\begin{bmatrix} I & 0 \\ A & B \\ C & D \end{bmatrix}^{\top} \begin{bmatrix} 0 & X & 0 \\ X & 0 & 0 \\ 0 & 0 & Y \end{bmatrix} \begin{bmatrix} I & 0 \\ A & B \\ C & D \end{bmatrix} \prec 0. \tag{2.30}$$

This equivalence persists to hold for

- *non-strict inequalities, if, in addition, the pair (A, B) is controllable.*
- *equalities, if, in addition, A is Hurwitz and the pair (A, B) is controllable.*

More recently, a frequency domain approach has been developed, wherein the test is reformulated as an equivalent maximum singular value problem [CS20], which is feasibly solved using for example the nonsmooth optimization methods developed by Apkarian and Noll [AN06b]. The detailed discussion and comparison of these methods is deferred to Section 4.5.

Performance by Integral Quadratic Constraints

Interestingly, IQCs can not only be used to guarantee robust stability, but also to verify signal based interpretations of various performance criteria, and when combined with IQC based robust stability analysis, one immediately yields worst-case performance bounds over all possible realizations of the uncertain system. Consider therefore an uncertain system with exogenous inputs and performance outputs as depicted in Figure 2.3b and given by

$$p = \Delta(q) \tag{2.31a}$$

$$\begin{bmatrix} q \\ z \end{bmatrix} = \underbrace{\begin{bmatrix} G_{qp} & G_{qw} \\ G_{zp} & G_{zw} \end{bmatrix}}_{G} \begin{bmatrix} p \\ w \end{bmatrix}, \tag{2.31b}$$

The performance IQC on the channel $w \to z$ is stated as

$$\mathcal{I}_p(\Pi_p, z, w) := \int_{-\infty}^{\infty} \begin{bmatrix} \hat{z}(i\omega) \\ \hat{w}(i\omega) \end{bmatrix}^{\mathsf{H}} \Pi_p(i\omega) \begin{bmatrix} \hat{z}(i\omega) \\ \hat{w}(i\omega) \end{bmatrix} \, \mathrm{d}\omega \leq -\varepsilon \|w\|^2, \tag{2.32}$$

where

$$\Pi_p \in \left\{ \begin{bmatrix} \Pi_{p11} & \Pi_{p12} \\ \Pi_{p12}^{\mathsf{H}} & \Pi_{p22} \end{bmatrix} \in \mathbf{RL}_{\infty}^{(n_z+n_w)\times(n_z+n_w)} \mid \Pi_{p11} \succeq 0 \right\}. \tag{2.33}$$

Again one will usually construct a family $\mathbf{\Pi}_p$ such that signals $w \in \mathbf{L}_2^{n_w}$ and $z \in \mathbf{L}_2^{n_z}$ satisfying the performance IQC (2.32) defined by any $\Pi_p \in \mathbf{\Pi}_p$ is equivalent to some performance criterion being fulfilled.

Corollary 2.23 (Robust IQC Performance [VSK16]). *Let $G \in \mathbf{RH}_\infty^{(n_q+n_z)\times(n_p+n_w)}$ be a stable transfer function and $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ a causal and bounded operator. Assuming*

1. *for all $\Delta \in \mathbf{\Delta}$ the interconnection (2.31) is well-posed, and*
2. *for all $\Delta \in \mathbf{\Delta}$ and for all $\Pi \in \mathbf{\Pi}$ the IQC (2.27) holds true,*

then the interconnection G – Δ is robustly stable and robust performance with respect to $\mathbf{\Pi}_p$ on the channel $w \to z$ is guaranteed if there exist $\Pi \in \mathbf{\Pi}$ and $\Pi_p \in \mathbf{\Pi}_p$ satisfying the FDI

$$\begin{bmatrix} G_{qp}(i\omega) & G_{qw}(i\omega) \\ I & 0 \\ \hline G_{zp}(i\omega) & G_{zw}(i\omega) \\ 0 & I \end{bmatrix}^{\mathsf{H}} \begin{bmatrix} \Pi(i\omega) & 0 \\ 0 & \Pi_p(i\omega) \end{bmatrix} \begin{bmatrix} G_{qp}(i\omega) & G_{qw}(i\omega) \\ I & 0 \\ \hline G_{zp}(i\omega) & G_{zw}(i\omega) \\ 0 & I \end{bmatrix} \prec 0 \;\; \forall\, \omega \in \mathbf{R} \cup \{\infty\}. \tag{2.34}$$

Suitable multipliers for various performance indices (such as quadratic performance, passivity, generalized $\mathbf{H}_2$ performance) can be found in the literature, see e.g. [VSK16]. We will state here two performance indices relevant for this work.

Lemma 2.24 ([VSK16]). *Given the assumptions of Corollary 2.23, the interconnection (2.31) is robustly stable and an upper bound $\gamma > 0$ on the induced*
$\mathbf{L}_2$-gain, $\sup_{\Delta\in\mathbf{\Delta}} \|\mathcal{F}_u(G,\Delta)\| < \gamma$, is guaranteed if there exist $\Pi \in \mathbf{\Pi}$ satisfying the FDI (2.34) with

$$\Pi_p = \begin{bmatrix} \gamma^{-1} I_{n_z} & 0 \\ 0 & -\gamma I_{n_z} \end{bmatrix}. \tag{2.35}$$

For robust $\mathbf{H}_2$-performance in the sense of $\sup_{\Delta\in\mathbf{\Delta}} \|\mathcal{F}_u(G,\Delta)\|_2$, several (not necessarily equivalent) signal-based generalizations of the $\mathbf{H}_2$-norm for non-LTI uncertainties exist [Pag99]. A dedicated discussion is given in Section 4.5.2. We state here a result for the impulse-response energy based interpretation, which effectively relies on the KYP lemma.

Lemma 2.25 ([VSK16]). *Given the assumptions of Corollary 2.23, a factorization $\mathbf{\Pi} = \{\Psi^{\mathsf{H}}\Theta\Psi \mid \Theta \in \mathbf{\Theta}\}$, and the resulting state-space realization*

$$\begin{bmatrix} \Psi & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} G_{qp} & G_{qw} \\ I & 0 \\ \hline G_{zp} & G_{zw} \end{bmatrix} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & 0 \\ C_2 & D_{21} & 0 \end{array}\right], \tag{2.36}$$

the interconnection (2.31) is robustly stable and an upper bound $\gamma > 0$ on the $\mathbf{H}_2$-performance is guaranteed by $\operatorname{trace}(B_2^\top X B_2) < \gamma^2$ if there exists $X \in \mathbf{S}^\bullet$ and $\Theta \in \mathbf{\Theta}$ satisfying

$$\left[\begin{array}{c|c} I & 0 \\ A & B_1 \\ \hline C_1 & D_{11} \\ \hline C_2 & D_{21} \end{array}\right]^\top \left[\begin{array}{cc|c|c} 0 & X & 0 & 0 \\ X & 0 & 0 & 0 \\ \hline 0 & 0 & \Theta & 0 \\ \hline 0 & 0 & 0 & I \end{array}\right] \left[\begin{array}{c|c} I & 0 \\ A & B_1 \\ \hline C_1 & D_{11} \\ \hline C_2 & D_{21} \end{array}\right] \prec 0. \tag{2.37}$$

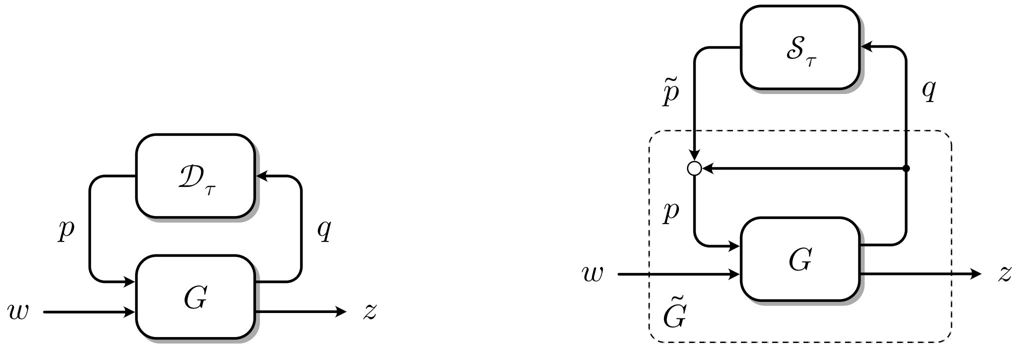

(a) Loop involving time delay operator $\mathcal{D}_\tau$. (b) Shifted loop involving shifted time delay operator $\mathcal{S}_\tau$.

Figure 2.4: Two equivalent LFRs of an LTI system affected by uncertain time delay.

2.3.4 Analysis of Time-Delay Systems

Multipliers for Time Delays

All real-world systems are affected to some extent by time delays. In many cases the size of the delay terms due to e.g. signal propagation or digital processing are small enough to be neglected in the system model, but as the demand on the closed-loop control bandwidths steadily increases, even time-delays in the order of a few microseconds can have detrimental effects on a high-performance system's gain margin and performance. Being such a common source of problems, TDSs have a rich literature dating back to the 50s [Fri14] and renewed interest since the 90s [Ric03] due to the emergence of networked systems, which together with multiagent and distributed systems remain the driving motivation for the continued research in this field [Den+22]. The central complication regarding the analysis of and controller synthesis for systems including time delays stems from the propagation time delays in the system model. Since the delay terms are irrational and infinite-dimensional, most conventional analysis and controller tuning methods are no longer applicable.

The following subsections introduce a simple but useful classical-control approach to handling time delays by rational approximations and a modern-control approach based on multiplier theory, that in conjunction with the IQC framework provides hard stability and performance guarantees where required. Other approaches such as compensation methods are of limited relevance in this work, as concluded in Section 3.3.2.

Real-Rational Approximations

Real-rational approximations of the irrational delay term(s) $e^{-s\tau}$ in a given system model are widely used in practice due to their simplicity and because they allow the use of conventional analysis and synthesis methods that require finite-dimensional systems [Ric03; SP09]. Truncated Taylor series and Padé approximations are most widely studied examples, where the Padé approximant is usually preferred over Taylor series because it remains stable also for higher orders [HP09].

The Padé approximant of order n for $e^{-s\tau}$ is given as the rational function

$$e^{-s\tau} \approx \frac{p_0 + p_1 s\tau + p_2(s\tau)^2 + \cdots + p_m(s\tau)^m}{q_0 + q_1 s\tau + q_2(s\tau)^2 + \cdots + q_n(s\tau)^n} \tag{2.38}$$

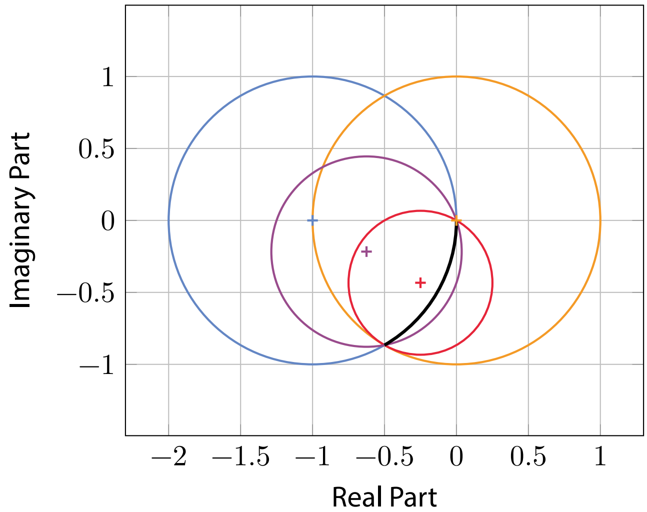

Figure 2.5: Exemplary Nyquist path of $\mathcal{S}_\tau(i\omega)$ (——) and visualization of the corresponding time-delay multipliers Π_1 (——), Π_2 (——) and Π_3 (——), as well as a conic combination of Π_1 and Π_3 (——).

where we assume in general $n \geq m$ to ensure realizability. The coefficients $p_0, \ldots, p_m$ and $q_0, \ldots, q_n$ are computed by equating the approximant to the first $m + n + 1$ terms of the Maclaurin series of $e^{-s\tau}$, yielding [Gv13]

$$\begin{aligned} p_i &= \frac{(m+n-i)!\, m!}{(m+n)!\, i!\, (m-i)!}(-1)^i, && i = 0, \ldots, m \\ q_i &= \frac{(m+n-i)!\, n!}{(m+n)!\, i!\, (n-i)!}, && i = 0, \ldots, n. \end{aligned} \tag{2.39}$$

While the case $m < n$ has known benefits in time domain analysis [Vaj00], we will – as is common – exclusively focus on the case $m = n$, which leads to the simplification

$$\begin{aligned} p_i &= q_i(-1)^i \\ q_i &= \frac{(2n-i)!}{i!\, (n-i)!}, && i = 0, \ldots, n. \end{aligned} \tag{2.40}$$

Consider the LTI time-delay operator $\mathcal{D}_\tau \colon \mathbf{L}_{2_e}[0,\infty] \to \mathbf{L}_{2_e}[0,\infty]$ for $\tau > 0$ given by $\mathcal{D}_\tau(v) = v(t-\tau)$. Recall that for $v \in \mathbf{L}_{2_e}[0,\infty]$, $v(t) = 0$ for $t < 0$. In robust control, the interconnection of the nominal (controlled) plant $G = \mathcal{F}_l(\tilde{P}, K)$ and an (uncertain) time delay $\mathcal{D}_\tau$ is frequently loop shifted as illustrated in Figure 2.4b, yielding the shifted timed-delay operator $\mathcal{S}_\tau = D_\tau - I$, i.e. $\mathcal{S}_\tau(v) = v(t-\tau) - v(t)$. It is clear that both representations are equivalent, i.e. that $\mathcal{F}_u(\tilde{G}, \mathcal{S}_\tau) = \mathcal{F}_u(G, \mathcal{D}_\tau)$. The primary reason to use $\mathcal{S}_\tau$ over $\mathcal{D}_\tau$ is that $\mathcal{S}_\tau$ causes no complications with unstable plants and that $\mathcal{S}_0(v) = 0$, which simplifies the analysis, cf. Assumption 2.18 [VSK16; DP00].

There exists rich literature on IQC-multipliers for (uncertain) time delays, a recent overview is given in [Eic16]. Notable contributions include [MR97; JS02; PS12] for constant time delays, [KL04; KR07] for time-varying delays, and the aggregating discussions in [Eic16; PS15]. As explained in the later references and illustrated in Figure 2.5, IQC multipliers for time delays can be represented as static or dynamic circle or half-plane regions bounding the arc mapped out by the time-delay operator in the Nyquist plane. The tighter the arc can be bounded, the less conservative the analysis (or synthesis) becomes [PS15].

Remark 2.26. One should be careful to note that conic combinations of such multipliers again are circle criteria and not, as one might expect graphically, the intersection area. The bound created by circles (——) and (——) in Figure 2.5 is given by circle (——) and not the tight crescent of intersection. Choosing the right combination of multipliers, in fact, is not straight forward and requires experimentation or numerical optimization [PS15].

Several challenges arise in the design of as-tight-as-possible time delay multipliers. The first is a suitable parameterization of the uncertainty in the sense of Assumption 2.18. Using the standard linear parameterization $\alpha \circ \Delta = \alpha \mathcal{S}_\tau I_{n_q}$ not only raises the issue of finding suitable multipliers such that, for all $\alpha \in [0, 1]$, the IQC defined by that multiplier holds for all $\alpha\Delta$ in place of Δ (cf. Theorem 2.17 and [JG01]), IQCs holding for this parameterization also provide (in this setting) awkward robustness guarantees. Indeed, the loop is in this case guaranteed to be robust against variations in the magnitude of the delayed signal, but not against variations in the delay size τ. A much more popular and for this application far more sensible parameterization is given by $\alpha \circ \Delta = \mathcal{S}_{\alpha\bar{\tau}} I_{n_q}$, such that robustness is guaranteed for all time delays up to a worst-case delay $\bar{\tau}$. This parameterization is valid under the mild assumption that the controlled plant is strictly proper across the uncertainty channel $M_{qp}(\infty) = 0$ [JG01].

Remark 2.27. While making the analysis in a sense more conservative, the parameterization of the uncertainty is a key ingredient for proving robust stability via IQCs per Theorem 2.17 and stability for smaller delays hence comes "free of cost". Certainly, there may be cases where higher performance can be achieved by not requiring stability for smaller time delays, however this edge scenario is not deemed to be of practical relevance in this work.

The second challenge is that the tightest conceivable multipliers are not only frequency dependent but switching, rendering state-space methods such as the common KYP lemma inapplicable for solving the IQC robust stability problem [PS15]. This is a consequence of these multipliers depending on the value of $\mathcal{S}_{\bar{\tau}}(i\omega)$, but needing to cover the full arc to $\mathcal{S}_{\bar{\tau}}(i\omega)$ even as $\omega\alpha\bar{\tau} > \pi$. For state-space methods, we will therefore employ rational approximations of those multipliers, which do not weaken the robustness guarantee but render the analysis slightly conservative.

We now state the three primary time-delay multipliers used in this work, also found and partially derived in [PS15]. Additional multipliers can be found in the other notable works referenced above.

Delay-independent circle criterion This static multiplier covers the entire Nyquist plot of $\mathcal{S}_{\bar{\tau}}$ and is on its own equivalent to verifying infinite delay margin.

$$\Pi_1 := \begin{bmatrix} 0 & -1 \\ -1 & -1 \end{bmatrix}. \tag{2.41}$$

Delay-dependent circle criterion A simple frequency dependent multiplier describes the interior of a circle centred at the origin of radius $|\mathcal{S}_{\bar{\tau}}(i\omega)|$.

$$\Pi_2(i\omega) := \begin{cases} \begin{bmatrix} |\mathcal{S}_{\bar{\tau}}(i\omega)|^2 & 0 \\ 0 & -1 \end{bmatrix} & \text{if } \omega\bar{\tau} < \pi \\ \begin{bmatrix} 4 & 0 \\ 0 & -1 \end{bmatrix} & \text{otherwise.} \end{cases} \tag{2.42}$$

A real rational approximation for this switching multiplier is given by

$$\bar{\Pi}_2(i\omega) := \begin{bmatrix} |\phi_2(i\omega)|^2 & 0 \\ 0 & -1 \end{bmatrix} \tag{2.43}$$

where

$$\phi_2(s) := 2\frac{(s\bar{\tau})^2 + 3.5s\bar{\tau} + 10^{-6}}{(s\bar{\tau})^2 + 4.5s\bar{\tau} + 7.1}. \tag{2.44}$$

Tight delay-dependent circle criterion An even smaller circle can be constructed around the arc described by $\mathcal{S}_{\bar{\tau}}$ by defining the circle going through the origin and $\mathcal{S}_{\bar{\tau}}(i\omega)$ and centred at the half-way point $\frac{1}{2}\mathcal{S}_{\bar{\tau}}(i\omega)$.

$$\Pi_3(i\omega) := \begin{cases} \begin{bmatrix} 0 & \frac{1}{2}\mathcal{S}_{\bar{\tau}}(i\omega) \\ \frac{1}{2}\mathcal{S}_{\bar{\tau}}(i\omega) & -1 \end{bmatrix} & \text{if } \omega\bar{\tau} < \pi \\ \begin{bmatrix} 4 & 0 \\ 0 & -1 \end{bmatrix} & \text{otherwise.} \end{cases} \tag{2.45}$$

A real rational approximation for this switching multiplier is given by

$$\bar{\Pi}_3(i\omega) := \begin{bmatrix} 0 & \phi_3(i\omega)^{\mathsf{H}} \\ \phi_3(i\omega) & -1 \end{bmatrix} \tag{2.46}$$

where

$$\phi_3(s) := \frac{-2.19(s\bar{\tau})^2 + 9.02s\bar{\tau} + 0.089}{(s\bar{\tau})^2 - 5.64s\bar{\tau} - 17}. \tag{2.47}$$

The multipliers that were just introduced are visualized in Figure 2.5.

Lemma 2.28 ([PS15]). *The IQCs defined by multipliers $\Pi_1, \Pi_2, \bar{\Pi}_2, \Pi_3, \bar{\Pi}_3$ from (2.41) to (2.43), (2.45) and (2.46) are satisfied by all $\Delta \in \{\mathcal{S}_\tau \mid 0 \le \tau \le \bar{\tau}\}$.*

As detailed in Section 2.3.3, conic combinations of multipliers (see Lemma 2.21) and frequency weighting can help to find suitable multipliers satisfying the IQC stability FDI (2.28). A useful example of this is given in [PS12]. It combines an $\mathcal{S}_{\bar{\tau}}$ over-bounding weight similar to Π_2 in (2.43) with Π_1 and frequency weighting. The former weight is given as

$$\Psi_{\mathcal{S}_\tau}(s) = 2\frac{(s\bar{\tau})^2 + (\frac{4}{\pi} + \varepsilon)s\bar{\tau} + \frac{4}{\pi}\varepsilon}{(s\bar{\tau})^2 - \pi\cos(\frac{\pi^2}{4})s\bar{\tau} + \frac{\pi^2}{4}}. \tag{2.48}$$

The frequency weighting is achieved by a generic dynamic basis such as

$$\Psi_\nu(s) = \begin{bmatrix} 1 \\ (s-\rho) \\ (s-\rho)^2 \\ \vdots \\ (s-\rho)^{\nu-1} \end{bmatrix}, \tag{2.49}$$

with $\nu \in \mathbf{N}_0$ and $\rho < 0$. For $\nu \to \infty$, the choice of ρ becomes negligible, but for small ν, line search to determine the optimal ρ may be beneficial [VSK16]. This choice offers a flexible parameterization of almost arbitrary frequency weights [PS12; VSK16].

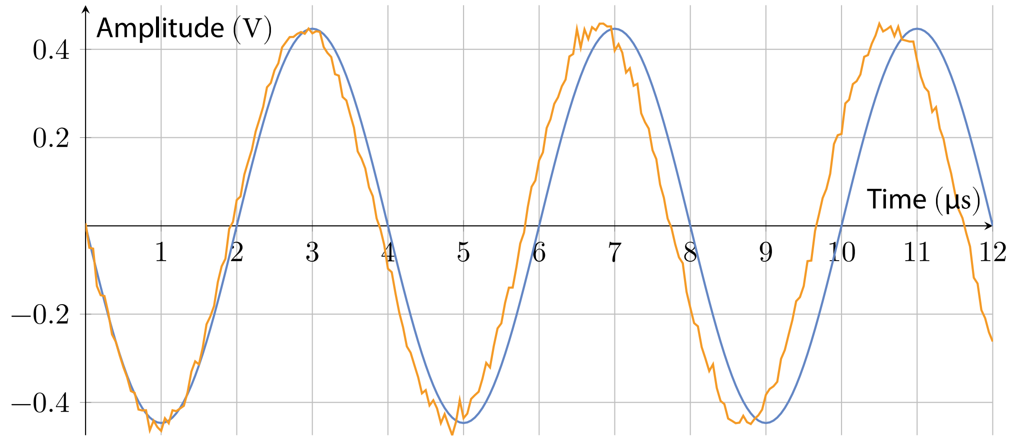

Figure 2.6: Example of a perfect harmonic oscillator signal (——) and a poor, drifting and noise-impaired oscillator signal (——) at 1 GHz and 3 dBm in a 50 Ω system.

Lemma 2.29 ([PS12; VSK16])**.** *The IQC defined by multiplier*

$$\Pi_4 := \begin{bmatrix} \Psi_\nu \Psi_{\mathcal{S}_\tau} & 0 \\ \Psi_\nu & 0 \\ 0 & \Psi_\nu \end{bmatrix}^{\mathsf{H}} \begin{bmatrix} \Theta_{11} & 0 & 0 \\ 0 & 0 & \Theta_{22} \\ 0 & \Theta_{22} & \Theta_{22} - \Theta_{11} \end{bmatrix} \begin{bmatrix} \Psi_\nu \Psi_{\mathcal{S}_\tau} & 0 \\ \Psi_\nu & 0 \\ 0 & \Psi_\nu \end{bmatrix} \tag{2.50}$$

for some $\Theta_{11}, \Theta_{22} \in \mathbf{S}^\nu$*, is satisfied by all* $\Delta \in \{\mathcal{S}_\tau \mid 0 \leq \tau \leq \bar{\tau}\}$ *if* $P(\infty) = 0$.

2.4 Time and Frequency Metrology

2.4.1 Oscillator Models

The core component of any synchronization system is a timekeeping device, i.e. a clock, whose state either poses as absolute time reference or is to be adjusted to match that of an external reference. In engineering, clocks are nearly always built using some realization of the conceptual *oscillator*, e.g. a pendulum, a quartz based electronic resonator circuit, or a coherent light source. In general, an oscillator is a dynamic system undergoing (sustained) motion around an equilibrium state, resulting in a periodic output. An oscillator with particularly stable repetition rate which is used as reference for other systems is also called a *frequency standard*.

Ideal Harmonic Oscillator

We start by considering the ideal *harmonic* oscillator in steady-state. Described by a simple differential equation, its output $v\colon \mathbf{R} \to \mathbf{R}$ is bounded and sinusoidal with some *angular frequency* ω_0, *amplitude* A and *nominal phase* φ_0, given by

$$v(t) = A\cos(\omega_0 t + \varphi_0). \tag{2.51}$$

By virtue of the periodic cosine function, there exists a constant *period* T such that $v(t) = v(t+T)$ for all t, equal to $T = 2\pi\omega_0^{-1}$. The reciprocal of the period is the *fundamental frequency*, $\nu_0 = T^{-1}$, consequently related to the (fundamental) angular frequency by $\omega_0 = 2\pi\nu_0$. The nominal phase φ_0 indicates a relative shift in time of the waveform resulting from initial conditions of the system, which is usually set to zero and omitted from the notation.

Real Harmonic Oscillator

Given a real oscillator, a clock is implemented by inferring, from the oscillators output $v(t)$, the time that has elapsed since an arbitrary starting point by keeping track of the full and fractional cycles the observed waveform has completed. However, any real world system is affected by its environment. When an unknown disturbance alters the behaviour of the oscillator and thus its observed output, the inference of time based on the ideal oscillator model will introduce a timing error. Since the ideal harmonic oscillator is fully characterized by its amplitude and phase, it is sufficient to distinguish between disturbances affecting either parameter, i.e. the amplitude fluctuation and phase fluctuation, given by $a_n(t)$ and $\varphi_n(t)$ respectively. Note that both are noise terms, however we choose to label them slightly more generally as fluctuations to avoid confusion with other conventional uses of the term "phase noise", which will be introduced later. The real harmonic oscillator is therefore given by

$$v(t) = (A + a_n(t))\cos(\omega_0 t + \varphi_n(t)). \tag{2.52}$$

Remark 2.30. We assume here that the real harmonic oscillator is still reasonably well-designed and shielded from the environment such that $|a_n(t)| \ll A$ and $|\varphi_n(t)| \ll 1$. This assures that the spectral integrity of the signal is maintained[2] and that the waveform is still discernable enough to be tracked.

Remark 2.31. Note that we could have also introduced *angular frequency fluctuation* $\omega_n(t)$, which is strictly equivalent to phase fluctuation by

$$\omega_n(t) = \frac{\mathrm{d}}{\mathrm{d}t}\,\varphi_n(t). \tag{2.53}$$

Remark 2.32. While amplitude and phase fluctuations are independent phenomena, they can only be unambiguously discerned in the peaks and zero-crossings of the cosine function, respectively.

General Oscillator

So far we have only considered the harmonic oscillator, which emits a sinusoidal signal. However, clock signals found in many applications have other periodic waveforms. For example, square waves are omnipresent in digital electronics, and mode-locked laser oscillators emit an optical impulse-train with ultra-short pulses, approximating a Dirac comb. Assuming a generic bounded *carrier* waveform $x(t)$, w.l.o.g. normalized to a period of 2π, we may modify the real harmonic oscillator (2.52) as

$$v(t) = (A + a_n(t))\, x(\omega_0 t + \varphi_n(t)). \tag{2.54}$$

Applying the exponential form of the Fourier series expansion to the carrier, we obtain

$$\begin{aligned} v(t) &= (A + a_n(t)) \sum_{k=-\infty}^{\infty} c_k \cdot \mathrm{e}^{k(\omega_0 t + \varphi_n(t))} \\ &= (A + a_n(t)) \sum_{k=-\infty}^{\infty} c_k \cdot \mathrm{e}^{k\omega_0 t}\mathrm{e}^{k\varphi_n(t)}, \end{aligned} \tag{2.55}$$

with $c_{-k} = \bar{c}_k$. This form not only visualizes that any periodic waveform can in fact be recreated by superposing sufficiently many real harmonic oscillators, but also reveals that a phase fluctuation affecting the entire carrier $x(t)$ is reflected by a corresponding, but scaled, phase shift on each harmonic, where the scaling factor is the harmonic order k.

[2] Large phase noise may lead to *carrier collapse*, i.e. significant loss in power concentrated in the carrier frequency, rather than its sidebands [Rub08, p. 9].

Tunable Oscillators

For many applications, in particular communication technology, synchronization and audio synthesis, some sort of control over the oscillator process is required. While the amplitude can be easily adjusted using a suitable amplifier at the output, the modulation of frequency and/or phase is more tightly coupled to the oscillator's method of operation. For example, using temperature control via a heating element is a common way to stabilize the frequency of high-precision crystal oscillator and offers a wide tuning range for optical resonators. In semiconductor-based oscillator circuits, variation of the passive components and biasing voltages also changes the frequency of oscillation. All these are examples of VCOs. The formalism is as before, but with frequency depending on some input signal $u(t)$, such that (using the real harmonic oscillator as an example)

$$v(t) = (A + a_n(t)) \cos(\omega(u(t))t + \varphi_n(t)). \tag{2.56}$$

Usually, the nature of the tuning signal will be irrelevant, and considering a time-varying angular frequency $\omega(t)$ suffices. On the other hand, the phase of the oscillator process can usually only be directly affected by some sort delay line, which scales in size with the desired actuation range and is also limited in speed. For these reasons, the *instantaneous phase* of the oscillator process is controlled via the integral relationship of phase and frequency:

$$\varphi(t) = \int_0^t \omega(\tau)\ \mathrm{d}\tau + \varphi_n(t). \tag{2.57}$$

2.4.2 Phase Noise and Timing Jitter

Depending on the application and the corresponding implementation details, noise induced amplitude and phase fluctuations can both pose a problem to achieving the required timing stability. In practice however, amplitude noise is secondary to phase noise for several reasons. On one hand, amplitude fluctuations are relatively small in modern oscillators, as oscillators have internal amplitude stabilization mechanisms [Rub08]. Also, phase detectors used to determine the relative timing of two oscillators are generally designed to be insensitive to amplitude changes. Phase fluctuations on the other hand are intrinsically related to the temporal stability of the oscillator process and the phase of a stand-alone oscillator moves freely. Being the main adversary, it is therefore essential for the subsequent analysis of synchronization systems to get a fundamental understanding of the properties of phase fluctuations and their impact on the oscillator process.

Traditionally, this discussion is predominantly held in the frequency domain as stable harmonic oscillators are key components of radio communication systems and phase fluctuations cause deterioration of the spectral purity of the carrier, leading to a widening of the frequency line as seen on a spectrum analyser [Rub08]. This also gave raise to the classical definition of "phase noise". When digital clock signals, where positive flanks on a square wave conventionally trigger a step in the underlying discrete architecture, came to be, time-domain specifications, i.e. several variants of "timing jitter", proved more natural to describe noise related performance degradations [DS18]. Inspecting (2.54), it is clear however that all these view-points and specifications are interrelated and can, to a large extend, be converted from one to another. Yet, most literature focusses on either view-point and, if at all, uses simplifications to discuss the interrelation, e.g. just considering the phase of the fundamental harmonic.

Being an important engineering problem, a sizeable amount of literature can be found on the topic of phase noise and timing jitter. A selection is given below, categorized into three focus areas:

1. statistical properties of phase fluctuations and phase noise impaired signals [RH08; Rub08; Dena; HRS02],
2. phase noise models for oscillators [Shi14; DMR98; Dem06; Loh+13a; Loh+13b], and
3. measurement and mitigation practice of phase noise [Denb; BT58; PW09; Hew85; Gru+20; Sch24].

Phase Noise

The IEEE standard 1139 [IEEa] defines *phase noise*, denoted $\mathcal{L}(f)$, only for harmonic oscillators. For general oscillators, the definition is applied to any single harmonic component, often simply the fundamental. With respect to (2.52) and (2.54), it is given as

$$\mathcal{L}(f) := \frac{1}{2} S'_{\varphi_n}(f), \tag{2.58}$$

where $S'_{\varphi_n}(f)$ as in (2.3) is the (estimated) one-sided PSD of $\varphi_n(t)$. Despite existing in this form for now more than 20 years, an outdated definition of $\mathcal{L}(f)$ is still widespread enough to warrant caution when coming across a specification of said quantity in datasheets or scientific literature. Unfortunately, this confusion is also upheld by practitioners continuing to specify the quantity $\mathcal{L}(f)$ logarithmically in units dBc/Hz in both definitions. Here "c" indicates "carrier" and stems from the original definition based on phase noise quantification using RF spectrum analysers but is given as rad^2/2 in the current definition. For completeness, we state the original definition:

$$\mathcal{L}_{\mathrm{RF}}(f) := \frac{S_v(\nu_0 + f) \text{ in 1 Hz bandwidth}}{\text{carrier power}} \quad \text{(obsolete)}. \tag{2.59}$$

While well-defined and of interest in some theoretical contexts, see e.g. [Gar05], $\mathcal{L}_{\mathrm{RF}}(f)$ is inadequate to describe phase stability of oscillators for several reasons pointed out e.g. in [Rub08]:

- Being more closely related to the equivalent additive noise representation $v(t) = A\cos(\omega_0 t) + n(t)$, it does not discriminate between amplitude and phase fluctuations.
- It is only suitable for small phase fluctuation induced angular changes. Once the small angle approximation fails to hold, the relationship of $\mathcal{L}_{\mathrm{RF}}(f)$ and $S'_{\varphi_n}(f)$ breaks down.
- This is in particular the case at sufficiently small offset frequencies, where the small angle approximation is never justified.

Further, it has to be noted that the IEEE standard 1139 [IEEa] bases even the new definition (2.58) on practical measurement techniques and convention rather than sound mathematical support. This is mostly because, as noted in several sources [Gar05; DMR98], $\varphi_n(t)$ is a random process that is notoriously difficult to model correctly. The perhaps most rigorous account to date is given in [Dem06]. The author discusses additive white and coloured noise sources acting on the oscillator dynamics and arrives at the conclusion that the resulting phase fluctuation $\varphi_n(t)$ is best modelled as passing stationary macro noise sources, aggregating all individual noise sources of a particular kind, through an ideal integrator. However, the integrator dynamics are not stable, which means that $\varphi_n(t)$ at the output is strictly non-stationary and *at best* described by a Wiener process, which is accurate, though unrealistic in practice, if the only macro noise source is white. As a consequence, the PSD $S'_{\varphi_n}(f)$ used in the definition (2.58) is actually not well-defined!

The remedy that again gives meaning to $\mathcal{L}(f)$ is to take into account the observation time of the process, that is finite in any real-world scenario. A clock may diverge infinitely from its reference if let run freely indefinitely, indicated by a simultaneously increasing variance σ_{φ_n}, but it will need an infinite amount of time to do so [Gar05]. In [Dem06], the unbounded phase fluctuation $\varphi_n(t)$ for that reason is disregarded in favour of a bounded, self-referencing variant $\varphi_{n,\tau}(t) = \varphi_n(t) - \varphi_n(t-\tau)$, which is provably stationary and has finite variance that scales with the observation time τ. Despite being mathematically well-behaved, we note that:

1. Due to the self-referencing, the PSD $S_{\varphi_{n,\tau}}(f)$ is forced to zero at all frequencies where $f\tau$ is integer. While disturbing at first, it can be argued that this reflects the common use of characteristic points in the waveform, e.g. the positive zero crossing in a sinusoid, for evaluating the phase, which will lead to corresponding blind spots in the frequency domain. This scenario of course corresponds to $\tau = \nu_0^{-1}$.
2. Neither $\varphi_{n,\tau}(t)$ or $S_{\varphi_{n,\tau}}(f)$ are actually measured by state-of-the-art phase noise analysers, which use spectral estimation techniques to compute an estimate $\hat{S}_{\varphi_n}$ from a finite amount of $\varphi_n(t)$ samples, assuming the process is stationary and all operations well-behaved. Hence, $\varphi_{n,\tau}(t)$ is not a practical tool for system design.

An alternative to using $\varphi_{n,\tau}(t)$, that is in agreement with common practice, is to assume that arbitrarily slow frequency drifts (those with time constants much larger than the observation time) are removed using an ethereal regulator, which ensures stationarity and boundedness of $\varphi_n(t)$ and hence an unproblematic $S_{\varphi_n}(f)$. While inaccurate for free-running oscillators in its own right, we note that phase noise analysers essentially work in the same way, i.e. the analysers internal reference oscillators track the frequency of the device under test (DUT), and all applications considered in this work fit this assumption.

An example of a phase noise spectrum is given in Figure 2.7. It is often well-described using a power law as a sum of f^0 white phase, f^{-1} flicker phase, f^{-2} white frequency, f^{-3} flicker frequency and f^{-4} random walk frequency noise [Rub08]. Slowly varying components, usually those below 10 Hz are also categorized as (random) *drift* (as opposed to “noise”). Note that, depending on the oscillator design, it is very well possible to have excellent long-term stability (i.e. low drift), but be very noisy at higher frequencies, and vice-versa. The best oscillators seen in practice are often hybrids created from multiple oscillators types, see e.g. [Gas+20]. The phase noise curve of Figure 2.7 exhibits spurs, for example at 50 Hz (power grid frequency), but also around e.g. 500 kHz. Spurious noise contents like this stem from deterministic disturbance sources and are easily traced back to the power supply, mechanical modes, etc. Due to their deterministic nature, it is inappropriate to account for their contribution to the phase noise in terms of a power spectral density, but should be (vocally) omitted from the spectrum before further processing [IEEa].

Timing Jitter

In digital applications, statistics on the discrete timing of events like the previously discussed positive-flank trigger, play a more dominant role. Several possible variants are described in the IEEE Standard 2414 [IEEb]. Despite a certain degree of discrete-time nature in the main application under consideration in this work[3], continuous-time metrics are more useful with respect to the control design studies which are to follow. The only relevant definition then is that of RMS timing jitter J_t, which can

[3] FELs are fundamentally designed around the concept of particle and optical pulse trains.

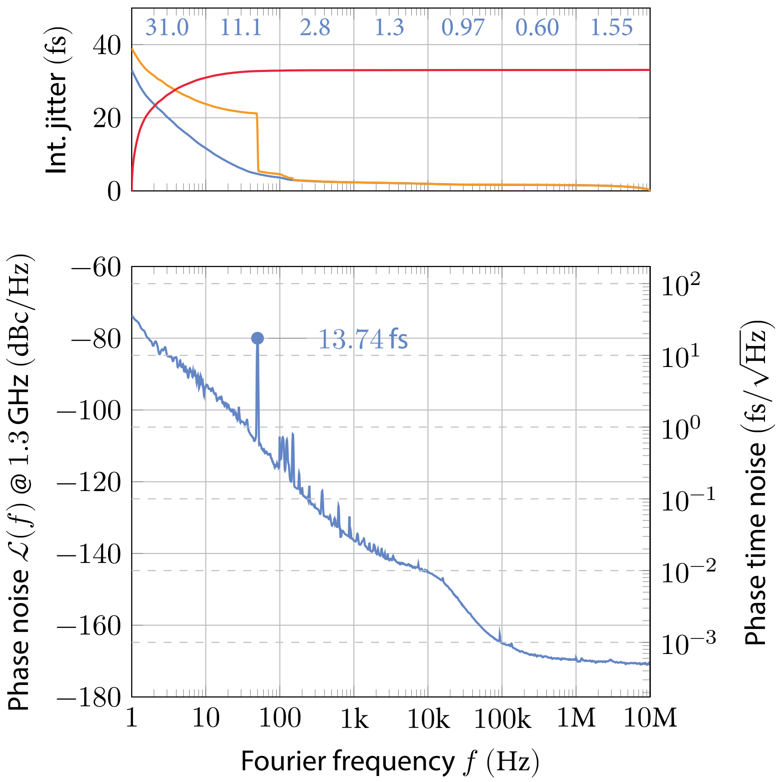

Figure 2.7: Example phase noise spectrum of a 1.3 GHz carrier with disturbance induced spurs and a white noise-floor beyond 1 MHz. The upper plot illustrates differences in the cumulative integrated jitter when integrated right-to-left without spurs (——), right-to-left with spurs (——) and left-to-right without spurs (——). Top-side numbers indicate the integrated jitter per decade.

be calculated via Parseval's theorem from any suitable PSD. In case of the phase noise spectrum $\mathcal{L}(f)$ we get

$$J_t = \theta_{n,\mathrm{rms}} = \lim_{T\to\infty} \sqrt{\frac{1}{2T}\int_{-T}^{T} \theta_n(t)^2 \,\mathrm{d}t} = \frac{1}{\omega_0}\sqrt{\int_0^\infty 2\mathcal{L}(f)\,\mathrm{d}f} = \frac{\varphi_{n,\mathrm{rms}}}{\omega_0}. \tag{2.60}$$

Here, we have implicitly defined the timing fluctuation $\theta_n(t) = \varphi_n(t)\omega_0^{-1}$, cf. (2.54). The corresponding one-sided power spectral density S'_θ will be denoted *timing jitter spectral density*. Note that, for reasons discussed previously, the integrals in (2.60) do not converge for free-running oscillators. It is common, also with consideration for limited bandwidth in phase noise analysers, to adjust the limits of the second integral to a practical range of interest, as e.g. 10 Hz to 10 MHz. In this case, we usually refer to *integrated timing jitter* and state the integration bounds [IEEa]. Deterministic disturbances appearing as spurs in the phase noise spectrum should be included in the timing jitter using root sum square (RSS) of integrated phase noise over the phase noise spectrum with spurs omitted and the spur jitter contribution

$$J_{t,f_{\mathrm{spur}}} = \frac{\mathcal{L}(f_{\mathrm{spur}})}{\sqrt{2}\pi\nu_0}, \tag{2.61}$$

where f_{spur} indicates the offset frequency at which the spur occurs [Smi22]. Unlike the phase, the timing jitter J_t is independent of the fundamental frequency of the system and hence is more suitable when comparing timing stability across heterogeneous systems.

Remark 2.33. In the characterization of oscillator stability, time-domain specifications, so-called *variances*, like the Allan variance, are often used [IEEa]. For the discussion in this work, they provide limited additional insight and therefore will not be explored further.

Chapter 3

Optimal Control for Phase Tracking and Distribution

3.1 The Phase Locked-Loop

3.1.1 Principle of Operation

Using the oscillator models established in Section 2.4.1, we can describe the synchronization procedure as the combination of measuring the instantaneous phase $\varphi(t)$ of the to-be-synchronized oscillator and applying a corrective action to its dynamic process that eliminates the true error to a (generally random) reference phase $\varphi_r(t)$. Quantitatively, this goal is encoded in the size of the performance variable $z(t)$, defined as

$$z(t) = \varphi_r(t) - \varphi(t). \tag{3.1}$$

While $z(t)$ is primarily driven by the oscillator's phase fluctuation $\varphi_n(t)$, other factors like the purity of the reference and measurement noise propagating into the system via the corrective action, play a role as well.

The automation of this synchronization process is a well-studied control problem with applications in various fields of engineering known as the *phase-locked loop (PLL)*, schematically depicted in Figure 3.1. Extensive discussions of PLLs and their properties can be found in [Gar05]. While the PLL was originally considered a stand-alone topic, it is clear that the design of the phase controller (in early implementations just a low-pass filter) is inherently a problem of feedback control theory.

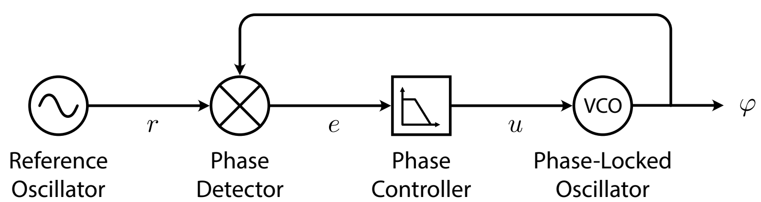

Figure 3.1: Schematic depiction of a phase-locked loop.

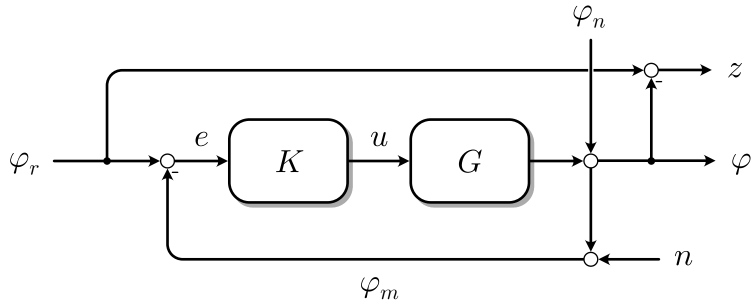

Figure 3.2: Feedback loop abstraction of a PLL.

In practice, phase detectors involve nonlinear operations such as frequency mixing. However, discussing the problem in baseband, i.e. eliminating the carrier signal and just considering dynamics directly affecting the phase of the oscillator, allows us to exclude these unnecessary complications from the analysis without introducing significant conservatism. The resulting feedback loop, shown in Figure 3.2, is standard and has been extensively analysed in the literature, see e.g. [FPE19]. The equations of this dynamic system are extended from (3.1) as:

$$z = \varphi_r - \varphi \tag{3.2a}$$
$$\varphi = Gu + \varphi_n \tag{3.2b}$$
$$u = Ke \tag{3.2c}$$
$$e = \varphi_r - \varphi_m \tag{3.2d}$$
$$\varphi_m = \varphi + n. \tag{3.2e}$$

The meaning of the variables is summarized as follows. The instantaneous oscillator phase $\varphi(t)$ is the sum of the response of the oscillator plant dynamics G to the control input $u(t)$ and the random phase fluctuation $\varphi_n(t)$, which encodes the dynamic response to disturbances acting on the plant and is characterized by its PSD $S_{\varphi_n}(f)$ (equiv. to the phase noise $\mathcal{L}(f)$). The control input is the response of the controller K to the measured phase error $e(t)$, which is computed from the reference phase $\varphi_r(t)$ and the measured oscillator phase $\varphi_m(t)$, which, as opposed to the true phase error $z(t)$ defined in (3.1), is corrupted by measurement noise $n(t)$. Additionally, let

$$K(s) = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array}\right]. \tag{3.3}$$

The main condition for the LTI model to be sufficient is that the phase detector operates in a linear regime once the loop is closed, or, using domain nomenclature, when the PLL is *locked*. For this reason, achieving initial lock can be tricky, as the error $z(t)$ has to be sufficiently small precisely when the controller is being engaged, so that the loop dynamics are indeed stable. Consequently, nonlinear analysis can be useful in very noisy systems. A method to deal with this complication is to use outer and inner loops of complementary bandwidth, dynamic range and resolution. The outer loop need not be particularly optimized, as its primary purpose is to "catch" the PLL, at which point the inner loop can be engaged, which ensures performance. Hence, the following discussion will focus on the inner loop. Vice-versa, strong disturbances can also lead to the PLL to lose lock or *slip*, i.e. jump one or multiple full revolutions in phase. Sometimes, a secondary task of the outer-loop is to avoid actuator saturation of the fast inner-loop by using a slow, high-dynamic-range actuator that slowly tunes the oscillator in the same direction as the inner-loop, thus keeping the inner-loop output centred.

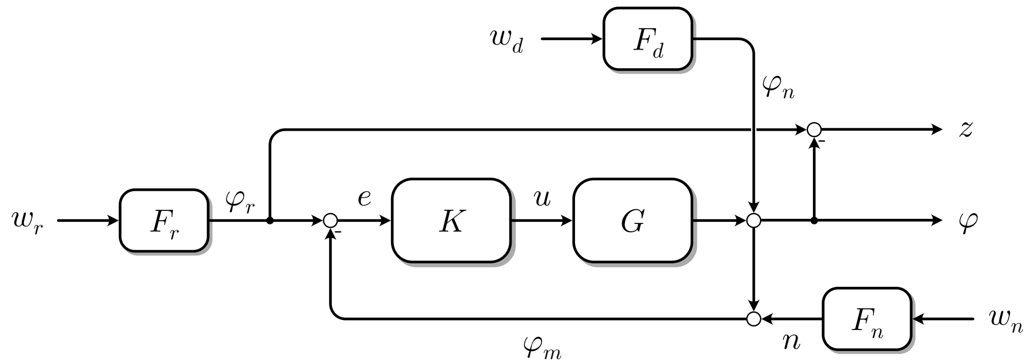

Figure 3.3: Generalized PLL feedback loop with noise shaping filters.

3.1.2 Noise Shaping

Most of the standard theory that considers performance of a control loop in a stochastic setting assumes white noise entering at the exogenous inputs. However, as outlined in Section 2.4.2, φ_n and φ_r certainly cannot be assumed white, and simply neglecting this fact will lead to a suboptimal controller design and poor jitter performance. We therefore employ the use of virtual colouring filters $F_{(\cdot)} \in \mathbf{RH}_2$ to shape the spectral content of the non-white random variables [1] from independent unitary white noise inputs $w_{(\cdot)}$ [BJ65; ZDG96], i.e.

$$\varphi_r = F_r w_r, \qquad \varphi_n = F_d w_d, \quad \text{and} \quad n = F_n w_n.$$

Indeed, we have from Lemma 2.2 that, when fed with white noise, the power spectral density at the output of the noise shaping filters matches the magnitude response of the filter itself.

Remark 3.1. The use of colouring filters introduces conservatism, as perfectly matching a measured noise spectrum would require filters of infinite order, which are non-practical. Also, as shown in [Heu18], the slope of the measured spectrum rarely fits multiples of 20 dB/dec. Fractional order systems are suitable to model non-spurious spectra [Heu18], but lack rigorous support in most numerical tools today and hence are not explored further. Instead, as supported by a real-world case study in Section 5.2.8, a reasonable trade-off between filter order and fitting accuracy generally yields satisfying results.

The resulting system is depicted in Figure 3.3. System (3.2) then adapts to

$$z = \varphi_r - \varphi \tag{3.4a}$$
$$\varphi = Gu + F_d w_d \tag{3.4b}$$
$$u = Ke \tag{3.4c}$$
$$\varphi_m = \varphi + F_n w_n \tag{3.4d}$$
$$e = F_r w_r - \varphi_m. \tag{3.4e}$$

While the form (3.4) nicely illustrates the structure of the PLL loop from a modelling perspective, the results from optimal control theory are applied easiest to a more abstract representation which separates the known plant dynamics from the sought-after controller dynamics and collapses the internal interconnection structure. The LFR shown in Figure 2.2, is easily constructed by collecting external

[1] The subscript d is chosen for the filter and white noise input corresponding to $\varphi_n(t)$ due to its significance as the disturbance input.

inputs, control inputs, measured outputs and performance outputs into vectors and separating all fixed dynamics collected in the so-called *generalized plant* P from the to be designed controller dynamics K. We let $w = \begin{bmatrix} w_r & w_d & w_n \end{bmatrix}^\top$ and write

$$P = \left[\begin{array}{ccc|c} F_r & -F_d & 0 & -G \\ \hdashline F_r & -F_d & -F_n & -G \end{array}\right]. \tag{3.5}$$

In line with analogous definitions made for G and K in Section 3.1.2, let a state-space realization for P be given by

$$P(s) = \left[\begin{array}{c|cc} A & B_w & B_u \\ \hline C_z & D_{zw} & D_{zu} \\ C_e & D_{ew} & D_{eu} \end{array}\right]. \tag{3.6}$$

Unsurprisingly, as will be discussed in more detail in Section 3.2, the properties of the weighting filters $F_{(\cdot)}$ bear consequences on the optimization problem for a suitable phase controller K, in particular existence and type of the solutions. Having in mind the discussion on phase noise from Section 2.4.2, what properties can be inferred when using a measurement from a phase noise analyser such as Figure 2.7 to fit the frequency response of F_d and F_r? First and foremost, the power law approximation for phase noise postulates non-finite steady-state gain and a non-zero feedthrough. Both features are problematic in our eventual formalism (cf. Section 3.2.2). Luckily however, both can be relaxed without sacrifice. Since the infinite bandwidth property of white noise is purely theoretical[2], the same holds for the white phase noise component and $\mathcal{L}(f)$ is eventually band-limited. Even though this limit will be beyond any reasonable measurement bandwidth, experience shows that the integral over the phase noise spectrum will have essentially converged at this point, due to the low magnitude of the white phase noise floor in today's oscillators. Adding an extra pole at a sufficiently high frequency when modelling the shaping filters F_r and F_d to ensure that P is strictly proper is therefore justified. Note that the same argument can be made for the measurement noise F_n, however the extra pole may be omitted for simplicity if K turns out to be strictly proper. We will discuss the conditions when this is the case in Section 3.2.2. At the other end of the frequency response, the marginally stable integrators of a power law model are not stabilizable by K in the given configuration. Indeed, it is a well known fact that any PLL will eventually lose lock because of actuator saturation [Gar05]. Two scenarios can be considered to motivate the same remedy. Either the actuator is not likely to saturate for the (finite) duration the system is operated, or the inner / outer loop configuration described in Section 3.1.1 preventing the fast inner compensator from saturating can be implicitly taken into account in the model for the inner loop. In both cases, a left-hand side roll-off of the noise spectrum starting at the expected operation time, respectively the closed-loop bandwidth of the outer loop, may be introduced. In the latter case, the standing assumption is that the outer loop will itself not saturate for the operation time of the system.

3.2 Optimal Phase Tracking

Given system (3.2), the optimal choice for the control action $u(t)$ depends on the control goal (usually a cost function to be optimized), the plant dynamics, the properties of output disturbance and measurement noise and possibly on other constraints, e.g. an imposed limit on signal sizes, the controller order or, more general, solution structure. In this section, these open points will be addressed by formulation of the optimization problem to be solved. We provide solution methods and discuss the impact of additional constraints on the achievable performance by means of numerical examples.

[2]No real (noise) generation process can produce arbitrarily high frequencies.

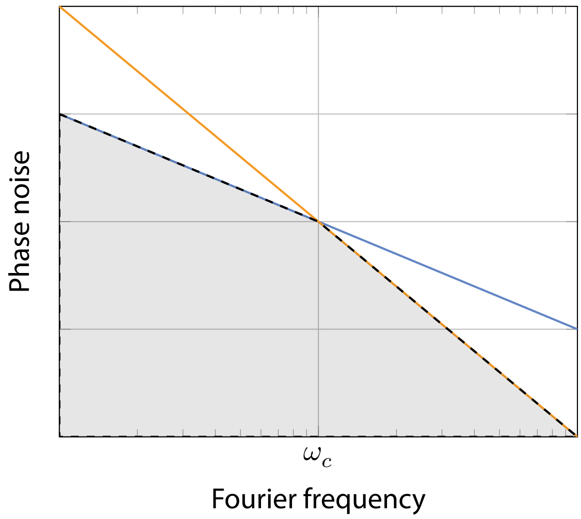

Figure 3.4: The best absolute timing stability (area under the curve) is achieved by setting the locking bandwidth to the phase noise crossover frequency ω_c of the reference (——) and the client (——), resulting in (- - -).

3.2.1 Problem Formulation

The conventional goal of a PLL is infinite-time optimal reference tracking, which ensures good synchronization, i.e. good timing replication at the controlled oscillator. In other words, $z(t)$ as defined in (3.1) is to be kept as small as possible over all time once the controller is engaged at time t_0. This is captured by the standard integral square error (ISE) cost function

$$J = \lim_{T\to\infty} \int_{t_0}^{T} |z(t)|^2 \, \mathrm{d}t. \tag{3.7}$$

Incidentally, this performance index is equivalent to the expected *relative* RMS timing jitter, cf. (2.23), which is easily extended from (2.60) as

$$J_t = \lim_{T\to\infty} \sqrt{\frac{1}{2T} \int_{-T}^{T} (\theta_r(t) - \theta(t))^2 \, \mathrm{d}t} = \frac{1}{\omega_0} \sqrt{\int_0^{\infty} S_z(f) \, \mathrm{d}f} = \frac{z_{\text{rms}}}{\omega_0}, \tag{3.8}$$

with $\theta_{(\cdot)}(t) = \varphi_{(\cdot)}(t)\omega_0^{-1}$.

Remark 3.2. If one is interested in minimizing the *absolute* RMS timing jitter, one would consider J_t with $\varphi_r(t) = 0$. There is however no phase detection scheme that can perfectly determine absolute timing jitter, a real and therefore itself noisy reference is always necessary. In that sense φ_m is in fact an artificial measurement quantity that is not accessible in the real implementation. Assuming independent noise, one could use multiple reference signals and cross-correlate them to isolate the absolute phase $\varphi(t)$, however the rejection scales with the square root of the number of independent reference sources, which does not practically scale, or requires large integration times, which would be unsuitable for feedback. With a single reference, one might be tempted to absorb the undesired random phase fluctuations of the reference into the measurement noise $n(t)$ and aim for rejection. However, this of course only makes sense if the controlled oscillator has superior short-term stability. The best absolute timing stability is then achieved by tuning the control bandwidth to the phase noise crossover point, illustrated in Figure 3.4. For the relative timing, no such limit exists, a higher tracking bandwidth will always further reduce the residual error. Consequently, in applications where absolute timing is *not* required, the best performance is achieved by defining the reference oscillator as "ground-truth" and tracking its fluctuations as closely as possible.

While any control input $u(t)$ can be expected to minimize the cost function (3.7), we are only interested in those that can be reasonably implemented, i.e. those generated from the measured output $e(t)$ by a finite-order and proper LTI control law K. Non-linear control is beyond the scope of this work. It's noteworthy however, that for the basic regulation and tracking problems, the optimal control action is in fact given by the response of such a suitable linear controller [Doy+89]. Furthermore, even though most controllers today are implemented using digital electronics, this work exclusively considers continuous-time linear system analysis and synthesis. Even though analogous results can be expected to exist for the respective problems in discrete-time, the discretization of the continuous-time plant dynamics or other variants of hybrid system control seem to offer no advantage over discretizing the continuous-time control solution - especially given the extremely high clock frequencies of modern day integrated circuits, which significantly surpass the available actuation bandwidth of most plants. Finally, as already motivated in Section 3.1, we assume that the system dynamics considered in this work are sufficiently linear and time-invariant for nonlinear and LTV control to not be of immediate importance. This holds true in particular for the system discussed in Chapter 5, but also for most low-noise phase control applications. Later on, some nonlinearity (specifically LTV dynamics) may be accounted for in the form of uncertainty in the system model.

Problem 3.3 (Optimal Infinite-Time Phase Tracking Problem)**.** Given the system (2.11), find a proper controller $K \in \mathbf{RH}_\infty$ which (by closing the control channel $e \to u$) internally stabilizes P from (3.5) and minimizes the cost function (3.7).

Remark 3.4. Note that Problem 3.3 is slightly non-standard in the sense that the reference to be tracked is not deterministic, but rather the stochastic phase of the (itself imperfect) reference oscillator, $\varphi_r(t)$. Albeit superfluous at this point, realize that the reference can be divided into three components: the control *setpoint*, i.e. a fixed phase offset to be kept w.r.t. the source, a deterministic phase movement of the reference, and the aforementioned reference stochastic phase fluctuation. The first two components however are rarely non-zero and do not warrant or even need special treatment, and the second and third component constitute $\varphi_r(t)$.

The field of modern and optimal control theory enjoys rich literature on the controller design problem associated with the minimization of performance indices such as (3.7) or variants thereof. We direct the interested reader in particular to geometric theory [Won85; Mar+10] and the disturbance decoupling problem [TSH01], and also LQG control [AM89; Wer21], which holds the stochastic nature of the disturbance input signals at its core. $\mathbf{H}_2$-control removes the stochastics from the mathematical formulation by showing equivalence to minimization of the impulse response energy of the system, which in turn is measured by the $\mathbf{H}_2$-norm of the system (cf. Section 2.3.2). Even though more abstract and perhaps less intuitive, owing to its mathematical elegance while still being perfectly well motivated from an engineering viewpoint, the $\mathbf{H}_2$-control variant of Problem 3.3 will be the focus from here on.

Problem 3.5 ($\mathbf{H}_2$-Optimal Dynamic Output Feedback Phase Control)**.** Given the system (2.11), find a proper controller $K \in \mathbf{RH}_\infty$ which (by closing the control channel $e \to u$) internally stabilizes P and minimizes $\|\mathcal{F}_l(P, K)\|_2$, i.e. the $\mathbf{H}_2$-norm of the closed loop system.

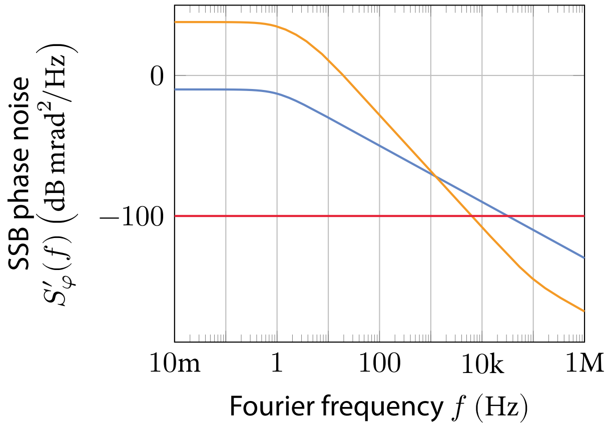

Figure 3.5: Magnitude responses for PLL case study noise colouring filters F_r (——), F_d (——) and F_n (——).

	Gain	Pole	Zero
G_{vco}	250	0 Hz	
G_{drv}	40	25 kHz	
F_r	2	1 Hz	
F_d	500	1 Hz	100 kHz
		5 Hz	
F_n	$1 \cdot 10^{-5}$		

Table 3.1: Parameter summary for a case study PLL system.

Despite the straight-forward formulation, Problem 3.5 is non-trivial. This is mostly because the system (2.11), as posed, does not fulfil certain regularity assumptions [SCS95]. In the following section, existence and uniqueness conditions for the unconstrained (full-order) controller problem and for proportional-integral (PI) control will be discussed. In cases where the infimal $\mathbf{H}_2$-norm cannot be achieved by a suitable controller, suboptimal controllers and sequences of suboptimal controllers converging to the optimal solution are provided. Controller design considerations beyond $\mathbf{H}_2$-norm optimization will be motivated and discussed, leading to the discussion in Section 3.2.3 on reduced-order control.

To illustrate key points in the discussion, we provide numerical examples for a simple 1.3 GHz case study PLL in the generalized formulation of system (3.4), adapted from the application in Chapter 5. The noise spectra of the reference and controlled oscillators and the phase detection chain are as shown in Figure 3.5. We consider the integral behaviour of a frequency controlled oscillator,

$$G_{\text{vco}} = g_{\text{vco}} s^{-1}. \tag{3.9}$$

Additionally, we consider a driver G_{drv} with first-order low-pass characteristics,

$$G_{\text{drv}} = \frac{g_{\text{drv}}}{\frac{s}{\omega_{c,\text{drv}}} + 1}. \tag{3.10}$$

Its gain g_{drv} is chosen such that the physical input saturation of the plant

$$G = G_{\text{drv}} G_{\text{vco}} \tag{3.11}$$

is normalized to the range $[-1, 1]$. The system parameters are summarized in Table 3.1.

3.2.2 Full-Order Dynamic Output Feedback

It is a well established result in control theory that, in general, the optimal full-information state feedback controller is static and the optimal output feedback controller is dynamic with order no larger than the order of the generalized plant to be controlled [SW17]. In particular, for linear quadratic problems, to which we count the LQG and the equivalent $\mathbf{H}_2$-optimal control problems, the optimal output feedback controller assumes an observer structure, i.e. is a combination of the optimal dynamic state observer and the optimal full-information static state feedback law. This is known as the renowned *separation principle* [Won68; HWS18]. In this section we therefore aim to characterize the infimum of the achievable performance and establish conditions (existence and uniqueness) for when the infimum can be achieved by a full-order dynamic output feedback controller.

Any particular $\mathbf{H}_2$-optimal control problem is categorized as either regular or singular.

Definition 3.6 (Regular $\mathbf{H}_2$-Optimal Control Problem)**.** The $\mathbf{H}_2$-optimal control problem (cf. Problem 3.5) with output feedback is called *regular* if the state-space form of the generalized plant P satisfies the following conditions:

1. The feedthrough matrix D_{zu} has full column-rank, i.e. is injective.
2. The feedthrough matrix D_{ew} has full row-rank, i.e. is surjective.

If the problem is regular, we obtain from e.g. [TSH01] sufficient existence and uniqueness conditions.

Theorem 3.7. *Consider the regular $\mathbf{H}_2$-optimal control problem with output feedback. There exists a unique, strictly proper controller K that stabilizes P and achieves infimal $\mathbf{H}_2$-norm over all controllers if*

1. *(A, B_u) is $\mathbf{C}_-$-stabilizable,*
2. *(C_e, A) is $\mathbf{C}_-$-detectable,*
3. *A has no (A, B_w)-uncontrollable eigenvalues on the imaginary axis, and*
4. *A has no (C_z, A)-unobservable eigenvalues on the imaginary axis.*

The necessary and sufficient conditions, relaxing the conditions on unobservable or uncontrollable eigenvalues on the imaginary axis, require details from geometric control theory which are beyond the scope of this work [TSH01, thm. 11.14]. In all cases, the unique minimizing controller, if it exists, is recovered from the solutions of two AREs. To keep the following theorem as comprehensible as possible, we additionally and without loss of generality consider the regular problem to be in *standard form*, i.e. $D_{zu}^\top C_z = 0$, $D_{zu}^\top D_{zu} = I$, $D_{ew} B_w^\top = 0$ and $D_{ew} D_{ew}^\top = I$. The general problem can be reduced to standard form by preliminary state feedback transformation [TSH01].

Theorem 3.8. *Consider the regular $\mathbf{H}_2$-optimal control problem with output feedback and in standard form. If there exists a controller K that stabilizes P and achieves infimal $\mathbf{H}_2$-norm over all controllers, it is unique and given by*

$$K(s) = \left[\begin{array}{c|c} A - B_u B_u^\top R^+ - Q^+ C_e^\top C_e & Q^+ C_e^\top \\ \hline -B_u^\top R^+ & 0 \end{array}\right], \tag{3.12}$$

where R^+ is the largest solution to the ARE

$$A^\top R + RA - RB_u B_u^\top R + C_z^\top C_z = 0 \tag{3.13}$$

and Q^+ is the largest solution to the dual ARE

$$AQ + QA^\top - QC_e^\top C_e Q + B_w B_w^\top = 0. \tag{3.14}$$

We assess the regularity of the problem for the generalized PLL plant (2.11) as well as the sufficient existence conditions for an $\mathbf{H}_2$-optimal controller. For the given structure, we immediately see that D_{ew} has full row-rank if and only if any of the noise colouring filters $F_{(\cdot)}$ has non-zero feedthrough. If either F_d or F_r has non-zero feedthrough, this implies $\|\mathcal{F}_l(P, K)\|_2$ non-finite by (3.5), however this case can be ruled out by arguments on infinite-bandwidth white noise made in Section 3.1.2. This leaves F_n, which can be assumed to have non-zero feedthrough in spite of the same argument, owing to the large noise-floor bandwidth compared to the control bandwidth. On the other hand D_{zu} has full column-rank if and only if P has non-zero feedthrough. This cannot be assumed in general. In fact, as discussed in Section 2.4.1, the oscillators most commonly found in PLLs have voltage controlled *frequency*, indicating integral behaviour by (2.57). In this case, the problem is certainly singular. An integrator-like plant would also irredeemably contribute an imaginary (A, B_w)-uncontrollable eigenvalue to A. Possible (C_e, A)-unobservable eigenvalues on the imaginary axis are in general of no concern.

Despite occurring frequently in applications, singular problems are treated secondary to their regular counterpart in the literature. Nevertheless, essential results on existence and uniqueness of solutions as well as how to compute them have been developed. Noteworthy contributions include early works on linear quadratic problems by Clements & Anderson [CA78], surrounded by attempts to regularize the singular problem by perturbing the system dynamics, see e.g. Jameson & O'Malley [JO75] and Francis [Fra79]. A mathematically more rigorous though less practical solution followed by allowing distributions as input signals, investigated by e.g. Hautus & Silverman [HS83], Willems et al. [WKS86] and Geerts [Gee89]. The singular $\mathbf{H}_2$ problem was eventually solved by Stoorvogel [Sto92], bar a conclusive result for the case when the controller is not required to be strictly proper. Central observations of this paper are that the separation principle does not necessarily hold for singular problems and that, as before with the linear quadratic problems, the Riccati equations providing the solution to the regular problem are replaced with inequalities, in particular LMIs. The latter fact is picked up in a later contribution by Sato & Liu [SL99], who consider the suboptimal case.

To conclude the discussion on full-order dynamic output feedback, we compare two common solutions methods from theoretical and practical standpoints and support the discussion with selected numerical examples computed for the case study system introduced in Section 3.2.1.

Regularization by Perturbation

The idea of the ε-perturbation method is to find a $\mathbf{H}_2$-norm minimizing series of strictly proper output feedback controllers for a perturbed but regular version of the original system that converge to achieve the infimal $\mathbf{H}_2$-norm of the original system as the perturbation parameter ε tends to zero [SCS95]. The perturbations ensuring that the regularity conditions from Definition 3.6 are satisfied take the form of additional input and output channels with ε-sized weights, such that D_{zu} and D_{ew} have full column- and row-rank, respectively. For example, for system (3.6) we thus get

$$\hat{D}_{zu} = \begin{bmatrix} D_{zu} \\ \varepsilon I \end{bmatrix}, \quad \hat{C}_z = \begin{bmatrix} C_z \\ 0 \end{bmatrix}. \tag{3.15}$$

The additional inputs and outputs are also used to ensure that the existence conditions Items 3 and 4 from Theorem 3.7 are satisfied, by guaranteeing (C_z, A)-detectability and (A, B_w)-stabilizability for any imaginary eigenvalues of A. Similar, to (3.15) we have

$$\hat{B}_w = \begin{bmatrix} B_w & \varepsilon I \end{bmatrix}, \quad \hat{D}_{ew} = \begin{bmatrix} D_{ew} & 0 \end{bmatrix} \tag{3.16}$$

and consequently the regular system

$$\hat{P}(s) = \left[\begin{array}{c|cc} A & \hat{B}_w & B_u \\ \hline \hat{C}_z & 0 & \hat{D}_{zu} \\ C_e & \hat{D}_{ew} & 0 \end{array}\right]. \tag{3.17}$$

While it is obvious that the controller synthesized for the perturbed plant necessarily also stabilizes the unperturbed plant [SCS95], the solution is of course conservative and the performance gap is not easily quantifiable a priori. In the practical setting, the performance gap cannot be reduced arbitrarily, as it is well known that the AREs (3.13) and (3.14) and the controller recovery become numerically ill-conditioned as the central solution is approached for $\varepsilon \to 0$ [Lin+98]. A controller recovered from a numerically unstable problem may also fail to stabilize the system at all. This is linked to several undesirable properties of the central solution outlined in e.g. [GA94; Ste13]. One issue is that the central controller cancels the poles of the plant, including those on the imaginary axis, resulting in unstable closed-loop dynamics. Another problem with the central solution results from controller poles near infinity that induce a quasi feedthrough-like structure, but are numerically problematic.

Remark 3.9. The closed-loop system in a regular problem is strictly proper if and only if the controller is strictly proper [TSH01].

Finally, for cases as discussed here, where $\varepsilon \to 0$ implies $\hat{D}_{zu} \to 0$, the limit case tends to produce a controller with extremely high gains, which physically cannot be implemented due to actuator slope and amplitude restrictions. This is reminiscent of the classical linear-quadratic regulator (LQR) control problem when relaxing the condition on the control input weight from positive-definitive to positive-semidefinite, a problem also known as "cheap control" [JO75; Fra79]. For these reasons, solutions obtained via this method are in many instances actually not useful in application.

LMI Solutions to the Singular Problem

A mathematically more rigorous solution to the singular $\mathbf{H}_2$-optimal synthesis problem that incorporates LMIs is developed in the series of works [Sto92; SL99; SW17]. In contrast to the ε-perturbation approach, the unperturbed singular plant is directly used for synthesis and the obscure parameter ε is replaced by a hard upper bound γ on the closed-loop performance, i.e. $\|\mathcal{F}_l(P, K)\|_2 < \gamma$. This parameter may either be subject to minimization when solving the SDP formed from the synthesis LMIs or can be relaxed by an acceptable amount to admit more numerically stable solutions to the solution space. For completeness, we summarize here the synthesis procedure as provided in [SW17] and refer to the original work for additional details.

Theorem 3.10. *Consider the $\mathbf{H}_2$-optimal control problem with output feedback. There exists a controller that achieves closed-loop $\mathbf{H}_2$-performance level $\|\mathcal{F}_l(P, K)\|_2 < \gamma$ if and only if there exist real matrix variables $X \in \mathbf{S}^{n_x}$, $Y \in \mathbf{S}^{n_x}$, $K \in \mathbf{R}^{n_x \times n_x}$, $L \in \mathbf{R}^{n_x \times n_e}$, $M \in \mathbf{R}^{n_u \times n_x}$, $N \in \mathbf{R}^{n_u \times n_e}$ and $Z \in \mathbf{S}^{n_z}$ satisfying*

$$\begin{bmatrix} (AY + B_u M) + (AY + B_u M)^\top & A + B_u N C_y + K^\top & B_w + B_u N D_{yw} \\ K + (A + B_u N C_y)^\top & (XA + LC_y) + (XA + LC_y)^\top & XB_w + LD_{yw} \\ (B_w + B_u N D_{yw})^\top & (XB_w + LD_{yw})^\top & -I \end{bmatrix} \prec 0, \tag{3.18a}$$

$$\begin{bmatrix} Y & I & (C_z Y + D_{zu} M)^\top \\ I & X & (C_z + D_{zu} N C_y)^\top \\ C_z Y + D_{zu} M & C_z + D_{zu} N C_y & Z \end{bmatrix} \succ 0, \tag{3.18b}$$

$$D_{zu} N D_{yw} = 0, \tag{3.18c}$$

$$\operatorname{trace}(Z) < \gamma^2. \tag{3.18d}$$

In that case, such a controller is given by

$$K(s) = \left(\begin{array}{c|c} A_c & B_c \\ \hline C_c & D_c \end{array} \right) \tag{3.19}$$

with

$$\begin{bmatrix} A_c & B_c \\ C_c & D_c \end{bmatrix} := \begin{bmatrix} U & XB \\ 0 & I \end{bmatrix}^{-1} \begin{bmatrix} K - XAY & L \\ M & N \end{bmatrix} \begin{bmatrix} V^\top & 0 \\ CY & I \end{bmatrix}^{-1} \tag{3.20}$$

where the intermediate square and nonsingular variables U, V satisfy $I - XY = UV^\top$.

Note that in this general case, the synthesized controller may be bi-proper. Strictly-proper solutions can be easily enforced by letting $N = 0$, however unlike for the regular problem, where the infimum is always attainable with a strictly proper controller, this may lead to conservatism in the singular case. In fact, as the γ-minimizing solution is identical to the central solution discussed already for the ε-perturbation method, directly allowing a feedthrough term relaxes the need for the optimal solution to produce a pole-zero cancellation at infinity [GA94]. This additional leeway in the solution may be exploited to improve the numerical stability of the controller recovery computations (3.19) through Scherer's two-step procedure [SW17, p. 97]. Nevertheless, the LMI problem remains effectively ill-posed in the singular case, as the infimal γ is not achievable with a finite solution.

The problematic pole-zero cancellation on the imaginary axis due to e.g. integral dynamics in the plant can also be controlled with the flexible LMI approach by adding additional LMI constraints to the synthesis LMIs. So-called pole-region constraints ensure that the poles of the closed-loop system are contained in a region shaped by intersection of basic geometric limits [CG96]. In particular, a minimum decay rate α can be imposed on all closed-loop poles by defining a half-plane left of α as admissible, which in turn causes the resulting pole-zero cancellation to be stable. The pole-region constraint LMI for this case is given by

$$\alpha \begin{bmatrix} Y & I \\ I & X \end{bmatrix} + \begin{bmatrix} (AY + B_u M) + (AY + B_u M)^\top & (A + B_u N C_y) + K^\top \\ K + (A + B_u N C_y)^\top & (XA + LC_y) + (XA + LC_y)^\top \end{bmatrix} \prec 0. \tag{3.21}$$

In [SGC97] the process of adding additional LMI constraints is interpreted as shaping the closed-loop systems Lyapunov function. Not surprisingly, this adds conservatism to the solution and the gap between the computed closed-loop $\mathbf{H}_2$-norm bound γ and the actually achieved value $\|\mathcal{F}_l(P,K)\|_2$ is no longer tight.

Remark 3.11. The enforceable minimum closed-loop decay is limited by the real part of the stable but uncontrollable modes of the generalized plant unless the method can be refined to control the individual dampings of the closed-loop poles. This is particularly relevant for the presently discussed application, where the shaping filters contain perturbed integrator dynamics close to the imaginary axis, motivated by the discussion in Section 3.1.2.

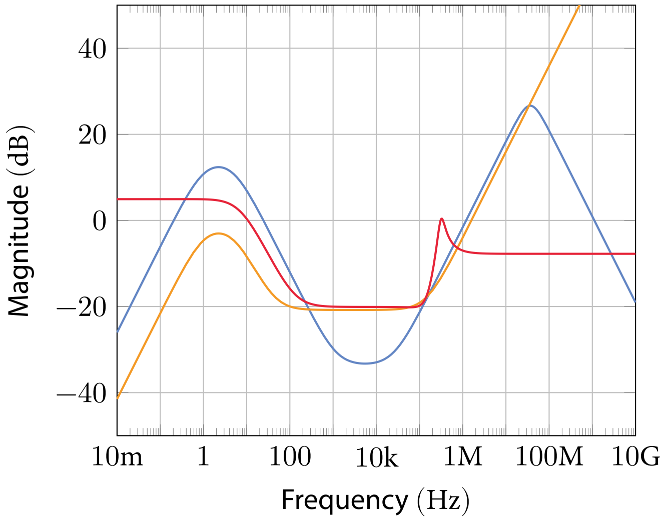

Figure 3.6: Magnitude responses of the optimal controllers obtained using the ε-perturbation method (——), the standard LMI-synthesis SDP (——) and an LMI-synthesis SDP with additional pole-region constraints (——).

Numerical Examples

To illustrate the discussed properties of the full-order feedback solutions, consider Figure 3.6 showing the frequency responses of several suboptimal full-order feedback controllers synthesized for the case study system introduced in Section 3.2.1. Quantitative performance results are summarized in Table 3.2 at the end of this section. The solution obtained by h2syn [The24c], a sophisticated, industry standard implementation of the ε-perturbation method, yields a closed-loop $\mathbf{H}_2$-norm of 4.56 µrad, or equivalently, an integrated relative timing jitter of 0.56 fs. One obvious problem with the synthesized controller is that it features dynamics beyond 10 MHz and very high gains[3] – challenging requirements for implementation. Common practice is to perform order reduction and merely keep poles and zeros below a suitable frequency, however it is not guaranteed that the reduced controller will still stabilize the system. Nevertheless, in this case, eliminating dynamics beyond ~ 120 Hz, still yields a stable closed-loop and achieves a reasonably deteriorated $\mathbf{H}_2$-norm performance of 5.93 µrad, a deterioration of ~ 30 %. Although not a major consideration at this point, we note another problem of the highly performance-optimized solution, which is lack of robustness, e.g. against perturbations in the bandwidth of the driver G_{drv}.

Unlike the h2syn function, which automatically computes a numerically optimized prescaling and as small as possible ε, and therefore offers no tuning knobs besides post-synthesis order reduction and tuning, the quality and properties of the LMI-synthesis solution are subject to several tuning knobs, including strict-inequality tolerances, LMI-region constraints, relaxation and regularization factors, bounds and penalties on the solution size and constraints and bounds on the properness and performance of the solution, respectively. An extensive investigation into the impact of every tuning knob is beyond the scope of this numerical case study. Solving the rudimentary $\mathbf{H}_2$-synthesis inequalities as posed in Theorem 3.8 and grid scanning over the strict-inequality tolerance, we obtain mostly unsatisfactory solutions with extremely fast controller dynamics similar to the h2syn solution but

[3]The state-space matrices contain multiple entries exceeding 10^6.

significantly worse performance. The problems stem from extremely large solutions to the essentially unbounded problem and the inability of the solver to assess the feasibility of the problem and its own progress. The large solution causes numerical problems in the controller recovery step, leading to an unstable closed loop or the observed significantly deteriorated performance.

Imposing the right additional constraints reduces the solution space to solutions that are numerically better behaved. In the case of LMI-region constraints, the constraints have the added benefit of acting as a design tool, allowing concise limits on the closed-loop poles, and thus the poles of the recovered controller. Indeed, enforcing a closed-loop pole damping of greater than $\sqrt{2}/10$ and decay-rate in the range $0.5 < -\operatorname{Re}(s)/(2\pi) < 50 \cdot 10^3$, yields a much more reasonable performance of 6.23 µrad. The cost of the additional constraints here is clearly conservatism, manifested in a worse performance compared to the h2syn method. Also, the performance bound γ guaranteed by the SDP is no longer tight and the gap to the actually achieved performance is significant.

In summary, we have established that, while the ε-perturbation method is intuitive and easy to implement algorithmically, the selection of ε is heuristic and the central solution obtained in the limit $\varepsilon \to 0$ is ridden with problematic features such as fast controller dynamics, high gains and lack of robustness. The LMI approach on the other hand offers great control over the obtained solution in terms of interpretable design parameters, but this comes at the cost of increased conservatism.

3.2.3 Reduced-Order Output Feedback

The previous discussion on full-order solutions, in particular the required order-reduction of the controller obtained from the ε-perturbation method and the sub-optimality of the constrained LMI solutions, immediately raises the question whether controllers of lower, respectively fixed order, can achieve similar performance. Clearly, this question has no useful answer for the general case. For example, 0^{th}-order proportional feedback may be insufficient to even stabilize a sufficiently "difficult" plant, and depending on the shape and order of the noise colouring filters vs. the desired controller order, an arbitrarily large performance gap can be expected. Nevertheless, given the relevance of low-order state-space and proportional-integral-derivative (PID) controllers in PLL applications, a qualitative investigation for a reasonably standard configuration as the one used as numerical case study in the previous section appears worthwhile. With today's abundant computational resources (and with a little knowledge on the expected parameter range), even brute-force controller parameter optimization is viable for low-dimensional search spaces. An extended discussion on reduced-order controller synthesis methods is omitted. We refer the interested reader to references [AM89; ZDG96; GA94] instead and ourselves employ grid search, Nelder-Mead [NM65] or nonsmooth optimization methods [Cla90; AGB14] as convenient.

Fixed-Order State-Space Control

Consider Figure 3.7a, showing the magnitude responses of two optimal fixed-order state-space controllers in relation to the quasi-optimal full-order controller benchmark computed via h2syn. The fixed-order state-space controllers were computed using MATLAB's systune function [The24b; AGB14], implementing a nonsmooth optimization synthesis method [Cla90; AN06b], and therefore come with a local certificate of optimality. The respective solutions for 1^{st} to 4^{th} order were consistently reached from 100 random initial conditions each, adding confidence to the attained optima. The 2^{nd}-order controller attained a closed-loop $\mathbf{H}_2$-norm of 4.56 µrad (identical to the full-order h2syn benchmark), while the 1^{st}-order controller still achieves a just $\sim$29 % higher norm (5.87 µrad). While it may seem

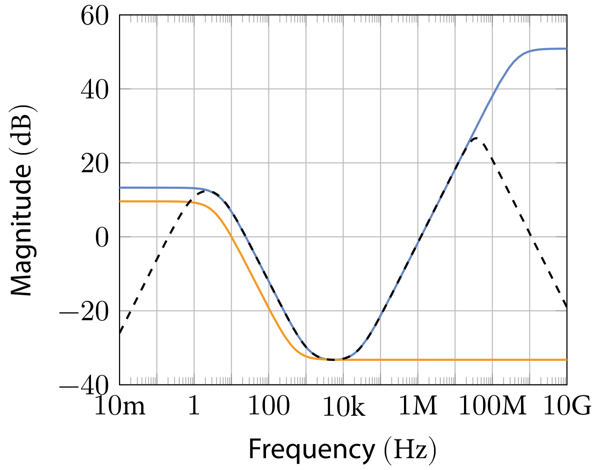

(a) 2^{nd}-order (——) and 1^{st}-order (——) state-space controllers.

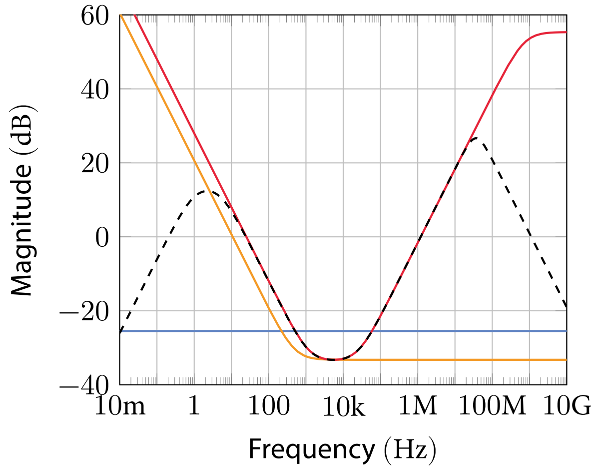

(b) Proportional (——), proportional-integral (——) and proportional-integral-derivative (——) controllers.

Figure 3.7: Magnitude responses of synthesized fixed-order controllers for the case study PLL system in relation to the quasi-optimal full-order controller (- - - .)

surprising that a virtually identical $\mathbf{H}_2$-norm can be achieved with a controller of just 2^{nd} order, recall that the regularization of the problem by ε-perturbation renders the solution marginally conservative and that the state-space solutions obtained by systune are not restricted to be strictly proper. That being said, we also observe that 3^{rd}- and 4^{th}-order state-space controllers with feedthrough cannot visibly improve on the benchmark performance. As such we conclude that full-order output feedback is not in general advantageous. In fact, just the 1^{st}-order solution is immediately feasible to implement, as the 2^{nd}-order solution again suffers from very fast controller dynamics, as visible from the high gain at the far end of the frequency axis.

PID Control

Finally, owing to their widespread use in practice, we remark on the achievable performance for the case study PLL system when using PID-controllers and subvariants thereof, as depicted in Figure 3.7b. Basic proportional output feedback achieved a closed-loop $\mathbf{H}_2$-norm of 7.39 µrad. Integral action is usually added to achieve perfect steady-state tracking. This consideration does not technically apply for $\mathbf{H}_2$-optimal control, where zero-mean disturbances are assumed, however, recalling that phase noise, as primary disturbance source in the PLL, has a power-law shape, adding integral action to the controller in addition to the plant, which is already of type 1, helps to suppress the large low frequency disturbances. In practical terms, unlike proportional feedback, a PI controller is able to track frequency drift in the reference oscillator, respectively perfectly eliminate frequency drift in the controlled oscillator (to the point of integrator saturation). Within the numerical case study PLL, a PI controller achieves a 21 % improvement over the plain P-controller (5.87 µrad). The derivative part of a full PID controller (with suitable derivative filter) boosts the phase of the open-loop towards higher frequencies, achieving better damping in the closed-loop impulse response and further improves the $\mathbf{H}_2$-norm by another 23 %, attaining the benchmark performance of 4.56 µrad.

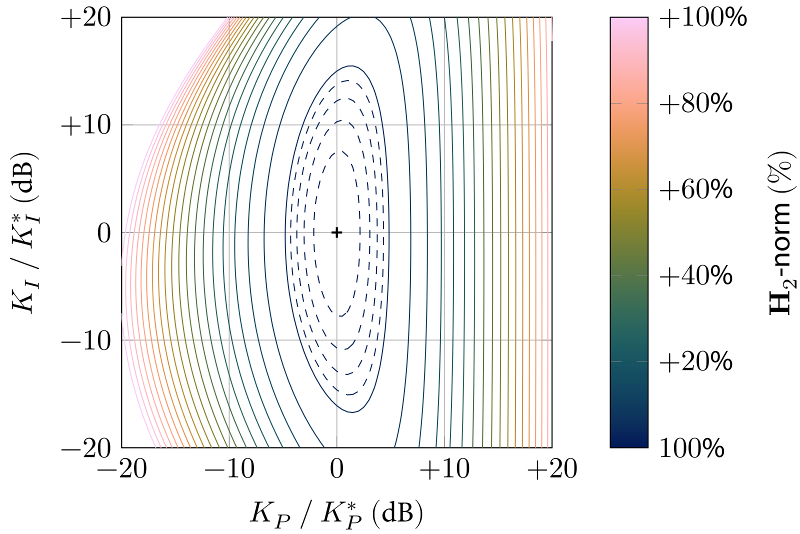

Figure 3.8: Sensitivity of the $\mathbf{H}_2$-norm achieved with a PI controller for the case study PLL system against variation of the optimal controller parameters.

As heuristic tuning of PID-gains is still common in practice, it is worthwhile to assess the sensitivity of the achieved performance against deviations from the optimal controller parameters. To this extend, consider Figure 3.8, showing the field of $\mathbf{H}_2$-norm achieved by a PI controller when the optimal parameters are perturbed by up to a factor ten. The worst $\mathbf{H}_2$-norm seen over the range −6 dB to 6 dB, i.e. half to twice the optimal gains, is less than 10 % higher than the optimal value and the valley around the minimum is reasonably flat, showing low sensitivity. Between the two parameters, the system is much more sensitive to a relative change in proportional feedback gain, shown by the steep flanks on the left and the right of the plot.

In summary, simple PLL systems may be suitably controlled with controllers as simple as plain PI feedback without significant loss of performance, even when the controller parameters are not perfectly optimal. While simple controller structures are in general less problematic as the full-order solutions, the performance gap may not be as low for less straight-forward plants, e.g. ones that feature underdamped resonances at low frequencies.

Method	Implementation	$\mathbf{H}_2$-norm (μrad)	Timing Jitter (fs)	Optimality Gap (%)
ARE + ε-pertb.	h2syn	4.56	0.56	0.00
ARE + ε-pertb. (impl.)	h2syn	5.93	0.73	30.21
LMI	custom	93.06	11.39	1939.34
LMI (pole-region cnst'd)	custom	6.23	0.76	35.26
Nonsmooth (2$^{\text{nd}}$-order)	systune	4.56	0.56	0.00
Nonsmooth (1$^{\text{st}}$-order)	systune	5.87	0.72	28.87
Nonsmooth (P)	systune	7.39	0.91	62.24
Nonsmooth (PI)	systune	5.87	0.72	28.88
Nonsmooth (PD)	systune	5.89	0.72	29.27
Nonsmooth (PID)	systune	4.56	0.56	0.00

Table 3.2: Performance summary of various controller types for a case study PLL system.

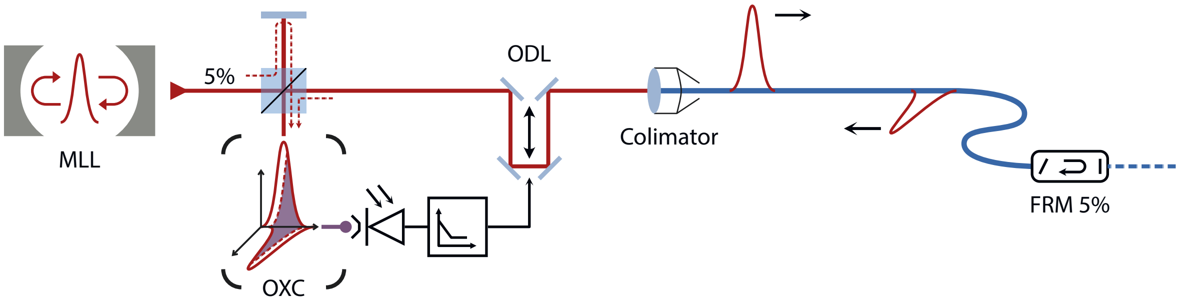

Figure 3.9: Schematic depiction of a pulsed optical fibre timing link based on the auto-correlation principle. The polarized pulses are partially reflected and rotated at the end of the timing link. The remaining light is available to the synchronization client. The returned pulse is compared to a fraction of a new pulse in a balanced optical cross-correlator (BOXC). The error is used to feedback to an optical delay line (ODL).

3.3 Timing Link Stabilization by Auto-Correlation

3.3.1 Principle of Operation

The previous section discussed the problem of synchronizing an oscillator to an available reference (usually another oscillator). We now turn to the problem when that reference is far away, i.e. the task of *phase distribution*. This discussion is motivated by a simple example.

Example 3.12 (Sensitivity of Signal Propagation by Coaxial Cable to Environmental Disturbances)**.** *Consider a phase-locked oscillator that is connected to a reference source via a* 1 km *long coaxial cable. Further, assume that the cable is initially trimmed exactly to an integer multiple of the wavelength of the reference oscillation and neglect any electrical losses or similar imperfections. Using the linear thermal expansion coefficient for copper and the velocity factor for industry-standard dielectrics polyethylene (PE) or polytetrafluorethylen (PTFE), we calculate for a temperature difference of just* 1 K *a timing error of* $\sim$85 ps*. For many applications dealing with dynamics in the few or sub-femtosecond regimes (cf. Section 5.1.1), a timing error of this magnitude is unacceptable. In reality, thermal expansion is only one among many effects deteriorating the so-called phase stability of coaxial cables or other transmission line media. Other factors include mechanical stress, humidity, air pressure and higher-order effects like the PTFE-knee, affecting not only the path length but also propagation speed [CS11; Šer+26].*

Several solutions have been developed to detect and compensate the phase error of communication channels induced by environmental disturbances [Sik+20; Mur+22], among which optical methods as described in [Kim+08] achieve most sensitive detection in applications, see e.g. [Sch+15]. The principle of operation is schematically depicted in Figure 3.9. The timing information is transmitted using an optical pulse train carrier waveform – usually generated by a mode-locked laser (MLL) oscillator and carried on an optical fibre – which is partially reflected at the end of the signal path. Back at the source, spatio-temporal overlap between the reflected and a new forward pulse is detected in a (balanced) optical cross-correlator via second-harmonic generation [Sch+03]. Variations in the round trip time of a single pulse due to disturbances acting on the optical fibre will result in varying amount of overlap and average power at the output of the correlator. In combination with a delay line actuator, this power measurement can be used for feedback stabilization of the timing link's phase. Due to the high energy density of the ultra-short optical pulses, attosecond resolution can be achieved, see e.g. [Sch+03].

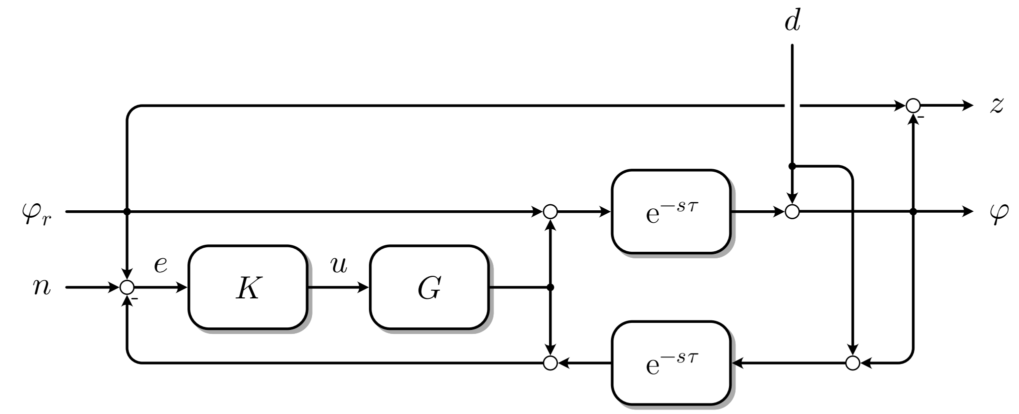

Figure 3.10: Baseband feedback loop of an auto-correlation timing link.

As before with the PLL system, we consider a baseband model of the phase dynamics in the auto-correlation timing link. The corresponding feedback loop is depicted in Figure 3.10. Assuming that all relevant components have linear responses (supported by arguments in [Heu18]), the system equations are

$$z = \varphi_r - \varphi \tag{3.22a}$$
$$\varphi = \mathrm{e}^{-s\tau}\varphi_r + \mathrm{e}^{-s\tau}Gu + d \tag{3.22b}$$
$$u = Ke \tag{3.22c}$$
$$e = (1 - \mathrm{e}^{-s2\tau})\varphi_r - (\mathrm{e}^{-s2\tau} + 1)Gu - 2\mathrm{e}^{-s\tau}d + n \tag{3.22d}$$

where G contains the dynamics of the delay line actuator and τ designates the nominal single-path propagation time. The remaining terms are as in Section 3.1.1.

Remark 3.13. Even though arguments can be made that the controlled variable is in fact the propagation duration of a signal over the physical link and hence should be expressed in units of time, we stick here to units of phase (i.e. a quantity normalized to some carrier frequency), hoping it will cause less ambiguity between the controlled deviation from nominal propagation time and the delay terms expressing the original nominal propagation time. In addition, it will later allow for easier interconnection with the previously formulated PLL models.

Note that, while actuator compensation and the disturbances have significant effect on φ, the relative change in propagation delay $\Delta\tau$ they induce is minuscule, which is why they are taken into account as additive terms exclusively. The model is also simplified in assuming that the disturbance acts on the end of the timing link, rather than as a combination of disturbances at random locations along the path. This simplification is justified as long as the following assumption holds.

Assumption 3.14. The correlation time of the random disturbance d is (much) larger than the single-path propagation time delay τ, i.e. $\mathrm{E}(d(t+\tau)d(t)) \gg 0$.

Similar to Section 3.1.2, we develop the generalized plant model by using the noise shaping filters

$$\varphi_r = F_r w_r, \qquad d = F_d w_d, \quad \text{and} \quad n = F_n w_n.$$

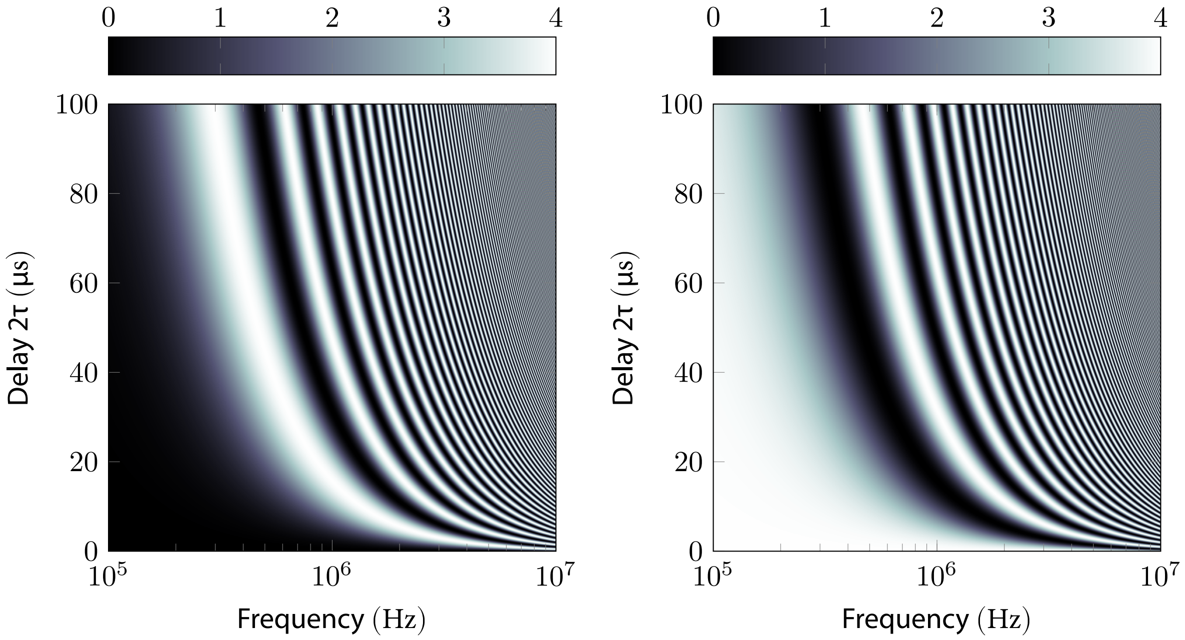

(a) $|1 - \mathrm{e}^{-s2\tau}|^2$ (b) $|1 + \mathrm{e}^{-s2\tau}|^2$

Figure 3.11: Squared magnitude responses of delay interferences.

The system then again fits the LFR form of Figure 2.2 and (2.11) but with

$$P = \left[\begin{array}{ccc|c} (1-\mathrm{e}^{-s\tau})F_r & -F_d & 0 & -\mathrm{e}^{-s\tau}G \\ \hline (1-\mathrm{e}^{-s2\tau})F_r & -2\mathrm{e}^{-s\tau}F_d & F_n & -(1+\mathrm{e}^{-s2\tau})G \end{array}\right]. \tag{3.23}$$

Delay Line Effect

As seen from (3.22d), the auto-correlation based measurement principle is not able to separate phase errors resulting from a jittering reference φ_r and those resulting from disturbances acting on the timing link. Indeed, when detecting a change in overlap of new and returning pulse in the detector, it is ambiguous which of the two pulses is early or late and by how much. All conceivable round-trip measurement methods suffer from an equivalent kind of ambiguity. For example, when using a local clock to measure the time it takes for pulse to return to the origin of the timing link, there imperfect local clock may show a deviation even though the timing link was perfectly stable. The achievable phase stabilization of a timing link with an auto-correlation scheme is thus always limited by the stability of the reference signal. From a control perspective, this is akin to measurement noise, but with an explicit dependency on the link length, via the round-trip time delay. This has been previously investigated in [Heu18]. The squared magnitude response (c.f. (2.5)) of the transfer function from φ_r to e, $1-\mathrm{e}^{-s2\tau}$, depending on the round-trip delay 2τ, evaluates to $4\sin(-\omega\tau)^2$ on the imaginary axis and is depicted in Figure 3.11a. Low frequencies are attenuated significantly while, for $n \in \mathbf{N}$, frequencies in the range $\left[\frac{(6n+1)}{2\tau}, \frac{(6n+5)}{2\tau}\right]$ are amplified up to 6 dB at $\frac{2n+1}{4\tau}$. Full elimination occurs at $\frac{n}{2\tau}$.

Remark 3.15. The above discussion also implies that vice-versa, a very stable timing link can be used to — within limits — measure and stabilize the absolute phase stability of a reference signal [Hew85].

We will not further pursue this idea, as the standing fundamental assumption is that the reference is authoritative and everything else relative to it. Nevertheless, the autocorrelation scheme reveals a rare case of absolute to relative jitter conversion.

A dual effect pertains the plant's response to the control action u, which is scaled by $1 + \mathrm{e}^{-s2\tau}$, the squared magnitude response of which, $4\cos(-\omega\tau)^2$, is depicted in Figure 3.11b. As the delay line actuator is part of the stabilized timing link, any pulse passes it twice per round-trip, separated in time by almost one round-trip delay 2τ. Consequently, there exist critical frequencies $\frac{2n+1}{4\tau}$ at which the phase of the actuation is exactly inverted when the pulse passes for the second time, leading to disturbance frequencies that virtually cannot be compensated. We say "virtually" because, even though the effect of the controller on those frequencies cannot be measured, (3.22b) shows that at the end of the link no such elimination exists for the delay line response.

3.3.2 Classification and Control of Time-Delay Systems

Compared to the PLL system discussed in Section 3.1, the main novel system characteristic introduced in the phase distribution problem is the propagation delay on the stabilized timing link and the switch from finite-dimensional system models to TDS. The goal of this section, rather than an extensive and systematic study of the field (which is vastly beyond the scope of this work), is to classify the type of TDS we have at hand and discuss and compare selected, tried and tested methods to synthesize adequate controllers.

Looking at the delay terms in (3.22), we note that the system is affected by multiple delays of different durations, which renders single-delay methods inapplicable. Still, in the absence of any extra delays (e.g. actuator input delays), the delay terms are homogeneous to the extent that all are integer multiples of the single-path transmission delay τ. Such *commensurate* delays are exploited in some approaches, see e.g. [JVM11]. Furthermore, as the delay terms cannot be factored out in (3.23), the system is affected primarily by state delays instead of merely input delays.

Methods for handling TDSs are classified by their ability to handle time-dependent and possibly uncertain delay terms. It is reasonable to assume that the nominal length of the timing link determining the delay terms is known or can be experimentally determined up to an uncertainty that is negligible for controller design, especially due to the fast propagation speed of the signal. Since all delay terms present in the system model (3.22) are directly tied to the variable being controlled, i.e. the round-trip propagation time, the terms are further sure to be highly time-invariant once the loop is closed.

Having classified the structural properties of the timing link TDS, we briefly discuss possible approaches to the arising control problem. Three main classes are prevalent in the literature:

1. **Rational approximation.** Approximation approaches aim to replace the irrational infinite-dimensional time delay terms with finite-dimensional real-rational approximations. The most notable among many is the Padé approximant. Since this method restores the applicability of standard methods for delay-free systems, this approach is often favoured in practice due to its simplicity [Ric03]. The downsides are that high-order approximations produce high-order systems models and controllers or that low-order approximations are unsuitable to represent and shape the system dynamics at higher frequencies, respectively. Furthermore, to utilize this approach, the delay has, of course, to be known a priori, and even then there is no guarantee that a controller synthesized for the approximated system model will stabilize the plant. The latter issue is especially problematic for time-varying delays.

2. **Delay compensation.** Adding structures such as the well-known Smith predictor [Smi57] or one of its derivatives to the controller stems from the idea of accurately transforming the TDS to a system without delay, designing a control law as usual and thus obtaining an optimally regulated plant. While the exactness is an appealing advantage over the previously discussed rational approximation approach, the trade-off is increased complexity. The classical Smith predictor can only be utilized for systems that are exclusively affected by input delay, and derivatives for state delays TDSs such as [Den+22] no longer yield LTI controllers due to the added integral calculations in the predictor. Again, the delay value needs to be known at the design phase and due to the known sensitivity of the Smith predictor, time-varying delays again are problematic to the point of critical closed-loop instability [Pal80]. Recently, the benefit of the Smith predictor has been questioned [SG18].

3. **Robust control.** The majority of works on analysis and control of TDSs acknowledge the fact that the time delays occurring in real systems are often not an integral part of the system dynamics but are undesired yet unavoidable side-products of the system's realization and as such inherently affected by uncertainty and often also time-dependence [KR07]. Rather than characterizing all occurring delays within the system realization – if at all possible –, methods from the field of robust control attempt to design control laws (generally low-order LTI) that stabilize the loop independent of possible delays or guarantee stability for an upper bound on the delay's magnitude at a given location within the loop. The flexibility of these approaches w.r.t. to the variety of covered scenarios and the guarantees they provide come at the cost of increased conservatism in terms of performance.

From the available approaches, the following discussion will focus on Padé approximations and robust control via IQCs, introduced in Section 2.3.4 and Sections 2.3.3 and 2.3.4, respectively. The Padé approximation will be primarily used for ad hoc analysis and benchmarks, whereas the IQC robust control techniques provide the desired stability and performance guarantees. Predictor compensation approaches are eliminated because of the existence of state delays in the auto-correlation timing link model and the resulting complexity of the suitable predictor structures in the controller, which is prune to limit the closed-loop bandwidth in a digital control implementation.

3.4 Optimal Phase Distribution

3.4.1 Problem Formulation

Consider the case when the delays affecting the system are known and fixed. The derivation of the formal optimization problem to be solved is in large parts identical to that of Problem 3.5 in Section 3.2.1. Just as the goal of a PLL is to replicate a reference phase at the output of a controlled oscillator, expressed in terms of the minimization of an ISE, the goal of the auto-correlation timing link is replication of the input signal phase at the output of the link. Again the objective can be expressed as the minimization of the closed-loop $\mathbf{H}_2$-norm. The only difference then is the internal structure of the plant to be controlled, which is secondary to the problem formulation. Hence, for completeness, we state

Problem 3.16 ($\mathbf{H}_2$-Optimal Dynamic Output Feedback Phase Transport)**.** Given the system (2.11) with P as in (3.23), find a proper controller $K \in \mathbf{RH}_\infty$ which (by closing the control channel $e \to u$) internally stabilizes P and minimizes (an upper bound on) $\|\mathcal{F}_l(P, K)\|_2$, i.e. the $\mathbf{H}_2$-norm of the closed loop system.

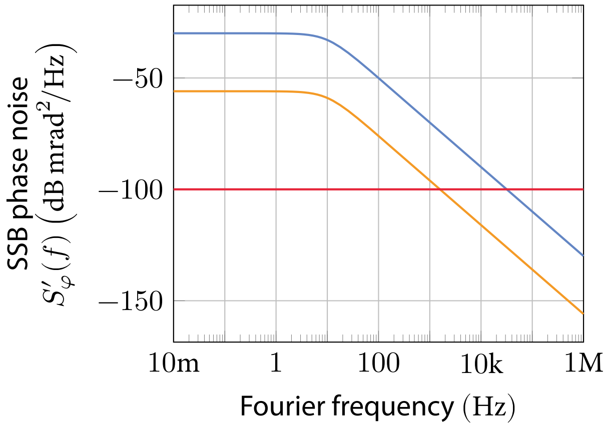

Figure 3.12: Magnitude responses for timing link case study noise colouring filters F_r (——), F_d (——) and F_n (——).

	Gain	Pole	Zero
g_{act}	0.4		
G_{drv}	80	1 kHz	
F_r	2	0.1 Hz	
F_d	$1.6 \cdot 10^{-3}$	10 Hz	
F_n	$1 \cdot 10^{-5}$		

Table 3.3: Parameter summary for a case study auto-correlation timing link system.

The presence of delays in the system to be controlled adds new aspects to possible problem solutions beyond those considered in Section 3.2, such as

1. computation of the $\mathbf{H}_2$-norm of a TDS [JVM11],
2. stabilizability of the system with respect to the size of the delay terms,
3. limits on achievable performance with respect to the size of the delay terms,
4. conservatism of the synthesized solutions,
5. existence of non-conservative problem solutions, and
6. robustness against uncertainty in the delay terms.

Some of these aspects will be investigated in the subsequent sections. As before, a case study model derived from the system discussed in Chapter 5 will be used to support the stated conclusions. Figure 3.12 shows the noise spectra of an input signal to be transmitted, a typical environmental disturbance noise spectrum for a 1 km optical fibre with spurs removed, and the typical noise floor of the round-trip time detector. The corresponding single path propagation time is $\tau = 4.9\,\mu\text{s}$. For G, an optical fibre stretcher translating an applied high-voltage signal to strain on the timing link is considered, modelled by a scalar conversion factor $g_{\text{act}} = 0.4$, in combination with a bandlimited high-voltage amplifier G_{drv}. The driver's gain is again chosen such that the physical input saturation of the plant $G = G_{\text{drv}} g_{\text{act}}$ is normalized to $[-1, 1]$. The system parameters are summarized in Table 3.3.

3.4.2 Synthesis Using Real-Rational Approximations

When applying the Padé approximant of order n to a system model, the resulting system will have n additional states for each independent delay term in the system dynamics, of which, in the case of (3.22), there are two when not considering additional input delays. This is particularly relevant when using the full-order output feedback methods introduced in Section 3.2.2 with the real-rational system approximation, as the synthesized controller size scales accordingly. We also note that high-order Padé

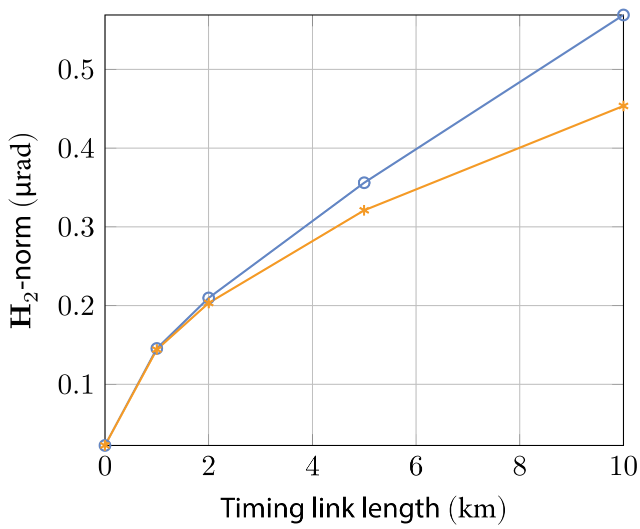

Figure 3.13: Approximate closed-loop $\mathbf{H}_2$-norms achieved by (3.24) (—o—), and the respective optimal controller for a Padé-approximated plant (—*—), when applied to the Padé-approximated plant. Padé approximation of 4$^{\text{th}}$-order. Parameters for case study system from Table 3.3.

approximants and those approximating small delays τ will lead to high-frequency poles and zeros in the approximated plant, which the central full-order output feedback controller will cancel (cf. Section 3.2.2), which is equally undesirable and at the same time limits the usefulness of pole-placement constraints in the LMI synthesis method.

Numerical Examples

We start the case study discussion by investigating the conservativeness of neglecting the transmission delay entirely during controller synthesis. The motivation for this scenario is the heavily simplified system identification and controller synthesis process. Applying the ARE-based h2syn method introduced in Section 3.2.2 to the minimal 2$^{\text{nd}}$-order plant (3.23) with $\tau = 0$ yields a 2$^{\text{nd}}$-order controller with one pole virtually at infinity, inducing a quasi-feedthrough, and one close to the origin, approximating integral action. For implementation, the controller can be very well approximated by the PI form

$$K = 0.05 + 300s^{-1}. \tag{3.24}$$

The performance achieved using this controller for varying link lengths is shown in Figure 3.13. The performance has been evaluated using 4$^{\text{th}}$-order Padé approximations of the plant and shows a 25.6 % decrease in performance for a 10 km long timing link when compared to the optimal controller synthesized for the Padé approximated plant. The divergence shows that for sufficiently large delays τ, neglecting the delay leads to significant sub-optimality. We also infer from this analysis that a greater delay is strictly linked to a higher infimal closed-loop $\mathbf{H}_2$-norm for transmission distances in that order of magnitude.

Remark 3.17. We acknowledge that the optimal controllers synthesized for Padé approximated plants might be challenging to implement. However, the larger the transmission delay τ, the slower the required dynamics of the optimal controller become while at the same time offering increasing performance benefits. Figure 3.14 shows a realistically implementable optimal controller synthesized using a 4$^{\text{th}}$-order Padé approximated plant with 10 km timing link.

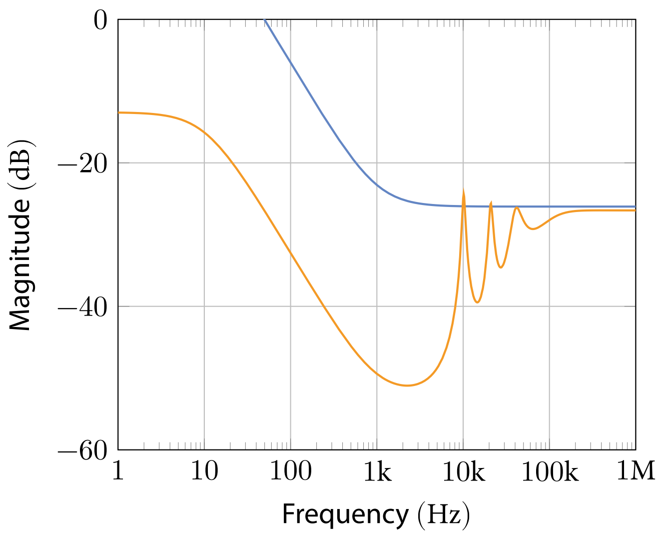

Figure 3.14: Magnitude response of (3.24) (——) and the optimal controller for a 4^{th}-order Padé approximation of a 10 km timing link (——).

While no hard robustness guarantees can be derived from the Padé approximant directly, the delay root locus (——) for the simple PI controller (3.24) shown in Figure 3.15 indicates infinite time delay margin. While the limit $\tau \to \infty$ is clearly always stable, this is not necessarily the case for all time delays, as shown by the second root-locus (——) for a controller with ten times higher integral gain.

3.4.3 Synthesis Using Integral Quadratic Constraints

Another scenario arises when the delays affecting the system are time-varying or not known well enough a priori for the methods explored for the nominal case to be feasible, or when hard stability and performance guarantees are required, which cannot be obtained from approximations. The natural extension is to pursue guaranteed closed-loop stability for uncertain parameters within a confidence interval or set. In general, the extended robustness to parameter variation in terms of closed-loop stability is paid for with a penalty in performance.

Following up on the introduction to the IQC framework in Section 2.3.3, we first provide an LFR for the interconnection of the delay-less plant with a suitable time-delay operator. Then, using the time-delay multipliers from Section 2.3.4, the robust stability of the controllers synthesized for nominal Padé-approximated plants from Section 3.4.2 is verified. Further, we discuss robust $\mathbf{H}_2$-performance, i.e. the guaranteed worst-case upper bound on the $\mathbf{H}_2$-norm over all possible realizations of the uncertainty, and finally outline IQC synthesis methods directly yielding robust stability and robust $\mathbf{H}_2$-performance without any finite-order approximation involved. As explained above, the controller effectively renders the transmission delay time invariant. Hence, the treatment of the time-varying case will be limited to literature references.

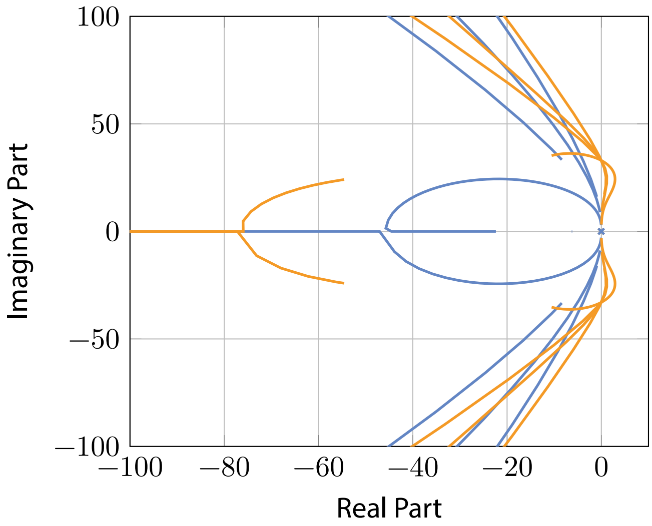

Figure 3.15: Closed-loop delay root loci from a 4th-order Padé approximation of plant (3.23) with controller (3.24) (——), and for an example controller achieving only finite delay margin (——), respectively. Lines indicate the positions of the closed-loop system poles (including poles from the delay approximant) for varying τ corresponding to 1 km to 100 km timing link length.

For the shifted configuration illustrated in Figure 2.4b and given by (2.12) and

$$\begin{bmatrix} q \\ z \\ y \end{bmatrix} = \tilde{P} \begin{bmatrix} \tilde{p} \\ w \\ u \end{bmatrix} \tag{3.25a}$$

$$\tilde{p} = p + q \tag{3.25b}$$

one easily rearranges (3.22) to obtain

$$\tilde{P} = \left[\begin{array}{cc:ccc:c} 0 & 0 & F_r & 0 & 0 & G \\ 1 & 0 & 0 & 2F_d & 0 & 0 \\ \hdashline -1 & 0 & F_r & -F_d & 0 & 0 \\ \hdashline 0 & -1 & F_r & 0 & F_n & -G \end{array}\right]. \tag{3.26}$$

The corresponding uncertainty block is given by

$$\Delta = \begin{bmatrix} \mathcal{S}_\tau & 0 \\ 0 & \mathcal{S}_\tau \end{bmatrix} = \mathcal{S}_\tau I_{n_q}. \tag{3.27}$$

Recall that $\mathcal{L}\{\mathcal{S}_\tau\} = \mathrm{e}^{-s\tau} - 1$.

Remark 3.18. Even though we have previously postulated that the propagation delay τ incurred on the timing link is a) effectively stabilized and b) experimentally well quantifiable, we remind the reader that the IQC framework is used here to provide hard robustness guarantees for a system including infinite-order dynamics from time-delays, whereas the other considered approaches can only provide soft guarantees in terms of approximations.

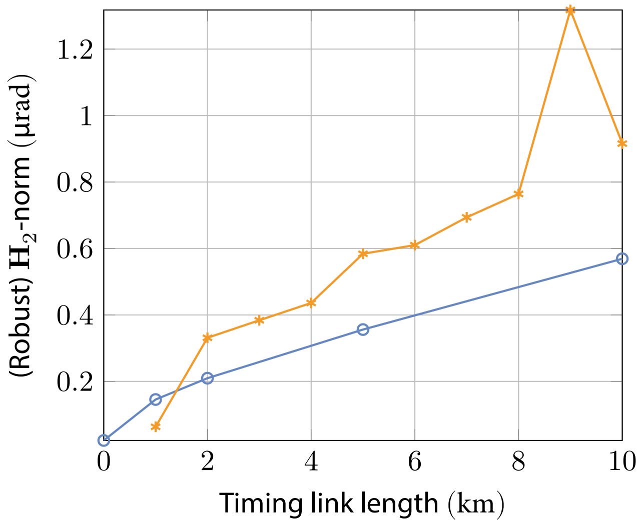

Figure 3.16: Upper-bound on the $\mathbf{H}_2$-performance guaranteed by Lemma 2.25 and Π_4 (2.50) ($\nu = 3, \rho = -1, \varepsilon = 10^{-3}$) for the plant (3.23) with controller (3.24), i.e. the quasi-optimal controller for $\tau = 0$, (—∗—), in comparison to the approximate closed-loop $\mathbf{H}_2$-norm computed by 4[th]-order Padé approximation (—○—).

Numerical Examples

Using the static time-delay multiplier Π_1 (2.41) and applying the KYP lemma to the system $\mathcal{F}_l(\tilde{P}, K)$ with (3.24) and (3.26), one easily shows feasibility of the resulting LMIs and hence infinite delay-margin using a basic SDP solver. For this robust stability analysis, the system (3.26) can in fact be drastically simplified by neglecting the disturbance inputs and the corresponding (a-priori stable) noise shaping filters. The block-diagonally repeated uncertainty (3.27) also reduces to a single time-delay operator $\Delta = \mathcal{S}_{2\tau}$ for the full round-trip time delay.

When considering $\mathbf{H}_2$-performance via Lemma 2.25, the SDP solver fails to assert feasibility for any conic combination of $\Pi_1, \bar{\Pi}_2, \bar{\Pi}_3$. One reason is the increased conservatism introduced by having to consider the original block-diagonally repeated uncertainty (3.27) for the forward- and backward-path delays, instead of just the round-trip delay, because of the additional performance channel signals. It is also possible that the real-rational approximations used in $\bar{\Pi}_2, \bar{\Pi}_3$ are overly conservative without additional frequency scaling. For this reason, Π_4 (2.50) is employed with $\nu = 3$, $\rho = -1$ and $\varepsilon = 10^{-3}$. The guaranteed robust $\mathbf{H}_2$-performance bound γ for varying values of $\bar{\tau}$ is shown in Figure 3.16 in comparison to previously obtained approximate $\mathbf{H}_2$-norm values using 4[th]-order Padé approximations (cf. Figure 3.13). Notably, for 1 km link length, the Padé approximation is more conservative than the IQC-based bound. While surprising, we remark that the LMI problem corresponding to the IQC robust performance test was solved with sufficiently large residuals on all constraints and good solver diagnostics. Other outliers on (—∗—) however indicate that numerical limitations can have severe consequences on the conservatism of the robust performance bound, though likely can be mitigated by numerical tuning of the model and multiplier.

Chapter 4

Optimal Control for Distributed Synchronization Systems

4.1 Architecture and Challenges of Large-Scale Synchronization Systems

Chapter 3 introduced the two core building blocks of any timing-synchronization system. The phase-locked loop was introduced as general tool for synchronizing a tunable oscillator (clock) onto a given reference. The means of phase error detection and frequency tuning of the synchronized oscillator are secondary to the theoretical analysis and merely provide the actual model parameters, while saturation constraints and higher-order effects can usually be neglected. The same holds true for the phase distribution problem, that was solved using the auto-correlation timing link principle. In this chapter we turn to the central part of this work, namely the study of an interconnection of multiple such building blocks that yields a full high-precision timing-synchronization system as it is needed in many large-scale science facilities.

Motivating this research is the key insight that individual optimal controller design for the composing subsystems does not yield satisfactory overall system performance. This fact is easily demonstrated by the following examples.

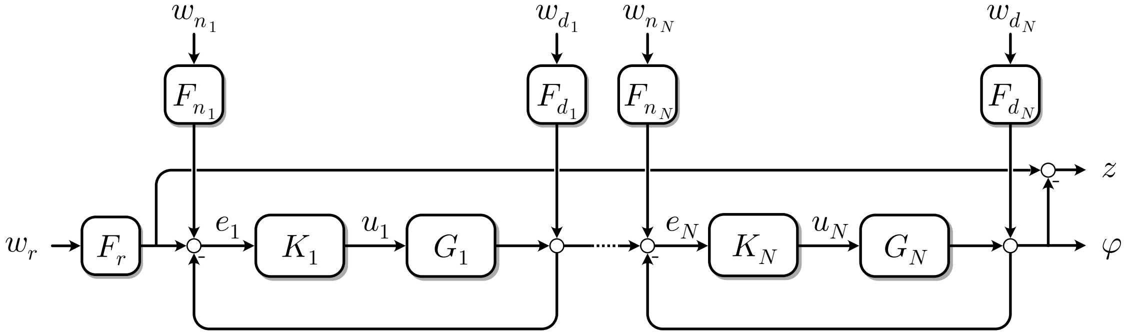

Figure 4.1: Block diagram of a phase reference tracking system consisting of daisy-chained PLLs.

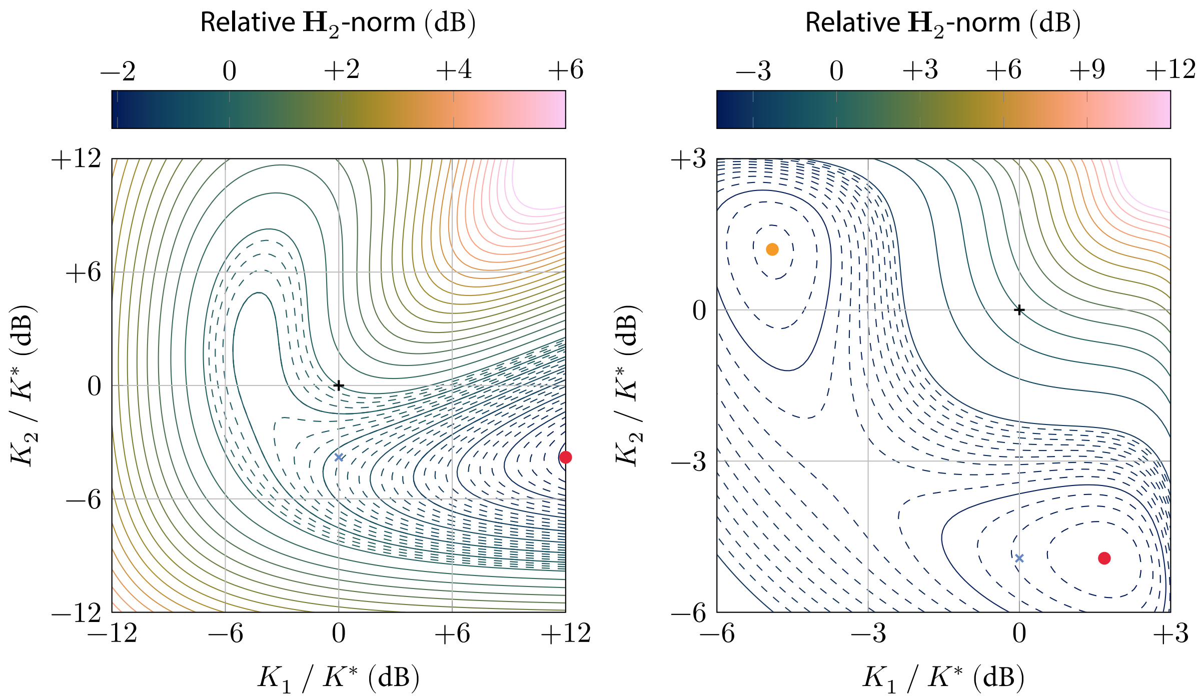

(a) Case study system according to Table 3.1.

(b) G contains an additional pole at 10 kHz.

Figure 4.2: Relative $\mathbf{H}_2$-norm achieved over a daisy-chain of two case study PLLs against variation of the proportional feedback gains. (+) marks the baseline performance achieved when using the optimal feedback gain K^* of a single PLL for both subsystems, (×) marks successive tuning, and (●) marks the optimum on the grid. (●) marks a second, local minimum.

Example 4.1 (Optimal Control of Daisy-Chained PLLs (a)). *Consider N daisy-chained but otherwise identical PLLs as depicted in Figure 4.1 for $N = 2$. To ease the visualization of the simulation result, we just consider a proportional feedback gain for each PLL controller and consider as performance index the synchronization of the trailing oscillator's phase to the source reference phase, quantified by the overall system's $\mathbf{H}_2$-norm. For a small system $N = 2$ with just two decision variables, parameter gridding is computationally feasible. The simulation results are depicted in Figure 4.2a, showing the $\mathbf{H}_2$-performance field against variation of the proportional feedback gains K_1, and K_2. This example uses the same model that was used for the numerical examples in Section 3.2 based on Table 3.1. The origin marked by (+) signifies the baseline performance achieved when tuning both PLLs in the chain to track the reference signal, i.e. when neglecting the interconnection structure and assuming for the trailing oscillator that the intermediate oscillator tracks perfectly. The shown cost function is clearly non-convex and the optimal tunings marked by (●) are clearly separated from the baseline. The performance gap is almost* 30 %.

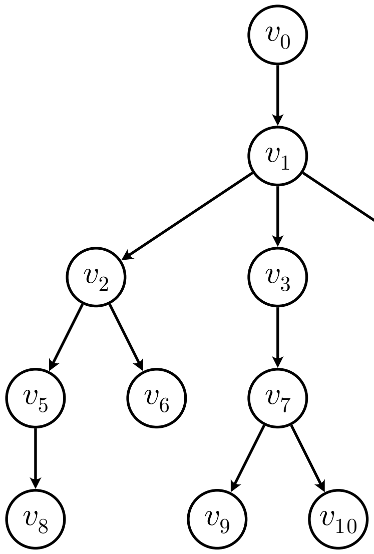

Figure 4.3: Example of an ordered tree representing a timing-synchronization network.

Example 4.2 (Optimal Control of Two Daisy-Chained PLLs (b))**.** *As Example 4.1 but assumes an additional pole in G at* 10 kHz*, inducing finite gain and phase margin. The new $\mathbf{H}_2$-performance field for $N = 2$ is depicted in Figure 4.2b, is again clearly non-convex and shows two distinct local minima. The performance gap is approximately* 40 %*.*

The results of Examples 4.1 and 4.2 are ultimately explained by Bode's sensitivity integral, commonly known as the *waterbed effect* [ÅM21]. It coarsely states that suppression of low-frequency disturbances well within the locking bandwidth inevitably leads to amplification of noise in another frequency band shortly beyond the locking bandwidth. If now both loops in the daisy-chain are tuned identically, this will lead to excessive over-amplification of noise in the waterbed regime, compared to when both regimes are separated. In the former case of Example 4.1, the feedback gain K_1 of the first PLL can be increased arbitrarily, as the second PLL acts as a low-pass filter, suppressing the high-frequency noise added by the first PLL. With the additional pole considered for Example 4.2 on the other hand, the first loop will eventually become unstable, leading to two distinct local minima. While in this example it is reasonably straight forward to arrive at (×) *close to* the local minimum (●) by first tuning the intermediate PLL and then tuning the trailing PLL based on the simulated output spectrum of the intermediate oscillator, in a less homogeneous system, the other local minimum might be significantly superior to the first, and is not so straight forward to attain due to the dependency inversion. When more complex interconnections are considered, as is usually the case in practice, other controller synthesis approaches are needed.

The remainder of this section is used to develop a general architecture of timing-synchronization systems and ways to express control goals on this architecture. This concludes in the formulation of a new high-level optimization problem to be solved, which will be refined in the subsequent sections.

Being an interconnection network of two kinds of subsystems, it comes to no surprise that the architecture of such a system is best described using graph theory. The constraints of a graph describing such a network are summarized in the following definition. For basic definitions on graph theory, the reader is directed to Appendix A.3.

Definition 4.3 (Timing-Synchronization Network)**.** A timing-synchronization network is represented by a directed graph (V, E) with the following properties:

1. The graph is connected, acyclic and rooted, i.e. an ordered tree.
2. The dedicated root node $v_0 \in V$ represents the reference source.
3. All nodes other than the root, $v_i = V \setminus \{v_0\}, i = 1, ..., k$, are control loops representing either a PLL (3.2) or a stabilized timing link (3.22) and have exactly one incoming edge[1], i.e.

$$\forall\, y \in V \setminus \{v_0\}\, \exists!(x, y) \in E, x \in V.$$

4. The (directed) edges of the graph represent input-output connections of the joined control loops.

An example of such a network is depicted in Figure 4.3. Of course one can construct cases where these constraints are violated, e.g. a purely cyclic network (which poses an interesting problem for stability analysis) or a system where a controlled oscillator should track two sources simultaneously as well as possible, however these are mostly academic problems with little practical application.

Remark 4.4. Note that some existing systems, such as the optical synchronization system for the FLASH and European XFEL facilities additionally feature special “resynchronization” units that correct the reference phase received via a drifting RF coaxial cable against a stabilized optical counterpart, see Section 5.2.5. Including such systems requires a slight adaption of Definition 4.3 to allow for nodes with multiple (ordered) inputs, but is ultimately fully compatible with the theory of this work.

We now ask how to define a performance index on this network. The simplest conceivable performance index is the uniformly weighted average phase ISE of all nodes $v_i, i = 1, ..., k$ with respect to the root v_0. In many practical applications however, a non-uniformly weighted average is the more suitable choice. Recall that the purpose of a timing-synchronization system is to enable several spatially distributed “client” systems to synchronously interact with a shared dynamic process, e.g. a particle beam in an accelerator or a radio wave received from space. Naturally, depending on their task, some client systems require tighter synchronization than others. Also note that client systems need not only occur at the leafs of the timing-synchronization network tree, although some intermittent nodes might indeed not be attached to a client system and just are required as repeater or fan-out stages within the network.

We hence formulate the optimization problem to be studied going forward:

Problem 4.5 (Optimal Distributed Timing-Synchronization)**.** Given a timing-synchronization network (V, E) and weights $w_i \geq 0$, find proper controllers $K_i \in \mathbf{RL}_\infty$ for each control loop $v_i \in V$, $i = 1, ..., k$, such that the weighted average phase ISE with respect to the root v_0 is minimized:

$$J = \lim_{T\to\infty} \sum_{i=1}^{k} w_i \int_0^T |\varphi_r(t) - \varphi_i(t)|^2 \;\mathrm{d}t. \tag{4.1}$$

4.2 Distributed Systems

4.2.1 Modelling Distributed Systems

In this section we build on Definition 4.3 and define general interconnected systems from a control theoretic point of view. Given the broadness of applications with interconnected systems and distributed

[1] This is in fact implied by the acyclic property.

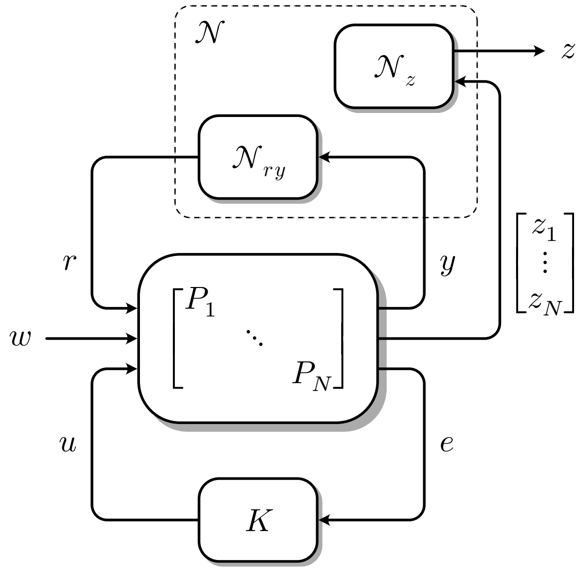

Figure 4.4: Linear fractional representation of a distributed system.

control, the theory in this section does not claim generality and instead is tailored to what is required for the discussion ahead.

Definition 4.6 (Distributed System [in the style of DD03; Eic16])**.** A distributed system is characterized by an interconnection of (spatially) separated subsystems (nodes), each with local dynamics, measurement outputs and actuation inputs.

Consider the causal and bounded distributed system $P\colon \mathbf{L}_{2_e}^{n_w+n_u} \to \mathbf{L}_{2_e}^{n_z+n_e}$ given by an interconnection of causal and bounded subsystems $P_1, \ldots, P_k : \mathbf{L}_{2_e}^{n_{r_i}+n_{w_i}+n_{u_i}} \to \mathbf{L}_{2_e}^{n_{y_i}+n_{z_i}+n_{e_i}}$. Derived from (2.11), we tactically dedicate n_{r_i} exogenous inputs and n_{y_i} regulated outputs for interconnection with other subsystems, while the remaining n_{w_i} exogenous inputs and n_{z_i} regulated outputs continue to be used for disturbances and performance signals, respectively. To avoid having to ensure compatible input and output sizes and to keep notation manageable, we assume w.l.o.g. that all subsystems have homogeneous input-output configurations, i.e. $n_{r_1} = n_{y_1} \ldots = n_{r_k} = n_{y_k}$ etc. It is then straight forward to write the distributed system P as an upper LFT of the concatenated subsystems $\mathrm{daug}(P_1, \ldots, P_k)$ and a special *network operator* $\mathcal{N}$. This is depicted in Figure 4.4.

$$P = \mathcal{F}_u(\mathrm{daug}(P_1, \ldots, P_k), \mathcal{N}) \tag{4.2}$$

The diagonal augmentation operator (2.29) ensures that the inputs and outputs of the concatenated subsystems are in order (i.e. not interleaved). The network operator $\mathcal{N}$ is given by

$$\mathcal{N} = \begin{bmatrix} 0 & \mathcal{N}_z \\ \mathcal{N}_A & 0 \end{bmatrix}, \tag{4.3}$$

where $\mathcal{N}_z \in \mathbf{RH}_\infty^{n_z \times (n_{z_1}+\cdots+n_{z_k})}$ aggregates the network performance signal z from the (weighted) subsystem performance signals and $\mathcal{N}_A \in \{0,1\}^{(n_{r_1}+\cdots+n_{r_k}) \times (n_{y_1}+\cdots+n_{y_k})}$ is the adjacency matrix of the interconnection graph.

Remark 4.7. When the interconnection satisfies the conditions of Definition 4.3, common system properties of the distributed system P such as stability, controllability, stabilizability, observability, detectability, etc. are easily inferred from the corresponding properties of the composing subsystems. For example, P is necessarily stabilizable if all P_i, $i = 1, ..., k$, are stabilizable.

4.2.2 Control of Distributed Systems

The problem of optimal control for distributed systems is not trivial. Indeed, it was already shown in Examples 4.1 and 4.2 that the optimal controller synthesis problem for a network performance index cannot be broken down into (smaller) optimal synthesis problems for the comprising subsystems. If we can afford to neglect the structure of P and apply the optimal controller synthesis methods introduced in e.g. Section 3.2.2, we arrive at the *centralized control* case.

Problem 4.8 (Optimal Centralized Control)**.** Given a distributed system P as in (4.2), find a proper controller $K \in \mathbf{RL}_\infty$ which (by closing the control channel $e \to u$) internally stabilizes P and minimizes some performance index on the closed-loop channel $w \to z$.

The characterizing property of the centralized solution is that the block matrix obtained when partitioning the solution K according to the signals of the comprising subsystems is potentially fully populated, i.e. the measured information of all subsystems is available (and used) to compute the optimal control input for any one subsystem in the network.

This degree-of-freedom is however not given in many practical applications, where often each subsystem is operating in informational isolation from the other subsystems, i.e. the locally measured information is available only to the local controller and not transmitted to other subsystems. Thus, a block-diagonal structure needs to be enforced on the solution K,

$$\begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_k \end{bmatrix} = \begin{bmatrix} K_1 & & & \\ & K_2 & & \\ & & \ddots & \\ & & & K_k \end{bmatrix} \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_k \end{bmatrix}, \tag{4.4}$$

so that ultimately $u_1 = K_1 e_1, ..., u_k = K_k e_k$. This other extreme is known as (fully) *decentralized* control.

Problem 4.9 (Optimal Decentralized control)**.** Given a distributed system P as in (4.2), find proper controllers $K_i \in \mathbf{RL}_\infty^{n_{u_i} \times n_{e_i}}$ for all subsystems P_i, $i = 1, ..., k$, such that $K = \operatorname{diag}(K_1, ..., K_k)$ (by closing the control channel $e \to u$) internally stabilizes P and minimizes some performance index on the closed-loop channel $w \to z$.

Naturally, intermediate cases exist as well, covered by the general term *distributed* control, illustrated in Figure 4.5. To formalize which controllers have access to which information, we introduce another binary matrix analogous to the adjacency matrix of the distributed system interconnection from Definition 4.3. This (controller) *structure* matrix $S \in \{0, 1\}^{k \times k}$ encodes which subsystem's measured signals are available for the computation of which subsystem's control signals.

Problem 4.10 (Optimal Distributed Control)**.** Given a distributed system P as in (4.2) and a controller structure matrix $S \in \{0, 1\}^{k \times k}$, find controllers $K_{ij} \in \mathbf{RL}_\infty^{n_{u_i} \times n_{e_j}}$ for $i, j = 1, ..., k$ satisfying $K_{ij} = 0$ if $S_{ij} = 0$ such that $K = [K_{ij}]$ (by closing the control channel $e \to u$) internally stabilizes P and minimizes some performance index on the channel $w \to z$.

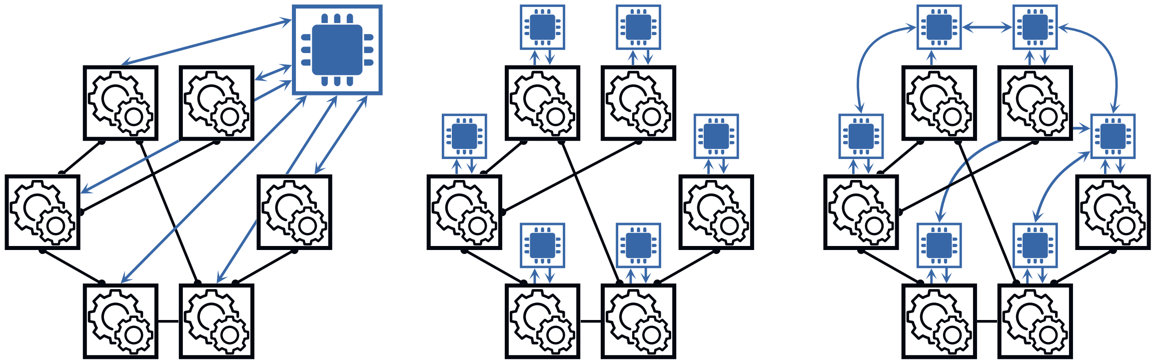

(a) Centralized (b) Decentralized (c) Distributed

Figure 4.5: Controller architectures (blue) for interconnected systems with coupled dynamics (black). Figure adapted from A. Eichler [Eic16].

Remark 4.11. In practice, one will usually find that for three subsystems P_1, P_2, P_3, if P_3 has information from P_2 and P_2 has information from P_1 then a) P_1 also has information from P_2 (and P_2 from P_3) and b) P_3 also has information from P_1, possibly at summed up latency. Formally, S is usually a) symmetric and b) transitively closed.

Remark 4.12. The question may arise if, in addition to the measured subsystem information e_i, the state information of the subsystem's controller can improve the overall network performance, when shared with other subsystems. For information structures where S is not transitively closed, cf. Remark 4.11, the controller state indeed conveys additional information to systems in the communication downstream. In the transitively closed case however, given initial controller state information across the network, each controller can infer ("observe") what the controllers of the connected subsystems are computing and track this information at the cost of higher local controller order.

Problems 4.9 and 4.10 are significantly more challenging than Problem 4.8 from an optimization point of view because, unlike the ARE-based method, where the solution is computed explicitly, or the LMI methods, where the resulting SDP can be solved efficiently using convex optimization, as discussed in Examples 4.1 and 4.2 and shown in Figure 4.2, the cost functions of distributed control problems are in general not convex and may admit multiple local minima. Determining the global optimum is then NP-hard in general, exceptions are mentioned in Section 4.3. We underline that the loss of convexity explicitly stems from the structural constraints on the solution. In order to tackle the general NP-hard problem, two approaches are dominant in the literature. The first is based on classes of convex approximations of the non-convex problem, yielding a related problem that can be solved efficiently but only provides a suboptimal solution. The main challenges here are finding the convex approximation "closest" to the original problem, minimizing the optimality gap, and, for robust control, ensuring that the solution of the approximated problem also guarantees robustness for the original problem [FK18]. The second approach is based on non-linear programming, in particular nonsmooth optimization based on subgradients and bundle methods, providing at least local optimality certificates. Ongoing research in this field attempts to integrate the powerful robust stability and performance analysis capabilities of the IQC framework (and multiplier theory in general) into the nonsmooth problem formulation [AN06a; CS20]. In the following Sections 4.3 to 4.5, state-of-the-art methods of both approaches are presented, discussed and developed further with novel extensions.

4.3 Structured Nominal Control by Sparsity Invariance

4.3.1 Convex Restrictions and the Sparsity Invariance Framework

This section introduces the SI framework developed by Furieri, Zheng et al. in [Fur+19; Fur+20]. [Fur+19] discusses the static state-feedback case, while [Fur+20] covers the more general (full-order) dynamic output feedback case. The following Section 4.3.2 extends on this framework by showing applicability of the method to static and fixed-order output feedback [Sch+21].

We start by introducing additional useful notation for binary and structured (transfer) matrices. The *sparsity subspace* associated with a binary matrix $X \in \{0,1\}^{m\times n}$ is defined as the set of causal transfer matrices

$$\text{sparse}(X) := \left\{ Y \in \mathbf{RL}_\infty^{m\times n} \mid Y_{ij} = 0 \;\; \forall\, i,j \text{ where } X_{ij} = 0 \right\}. \tag{4.5}$$

Analogously, given a (transfer) matrix $Y \in \mathbf{RL}_\infty^{m\times n}$, let $X = \text{struct}(Y)$ be defined by

$$X_{ij} = \begin{cases} 0 & \text{if } Y_{ij} = 0 \\ 1 & \text{otherwise.} \end{cases} \tag{4.6}$$

Given two compatible binary matrices $X, Y \in \{0,1\}^{m\times n}$, we say $X \leq Y$ if and only if $X_{ij} \leq Y_{ij} \;\; \forall\, i,j$ and $X < Y$ if in addition there exists at least some indices i,j such that $X_{ij} < Y_{ij}$.

The SI framework tackles non-convex distributed control problems of the form of Problem 4.10 by systematic construction of *convex restrictions* of the original problem. The appealing feature of convex restrictions over the complementary class of convex approximation by relaxation is that solutions to convex restrictions of a problem are necessarily feasible for the original problem as well. Convex relaxations on the other hand may produce infeasible solutions, but also feasible solutions where the convex restriction has no solution at all. Additional information on convex relaxations can be found in [Faz+17; WLS18; WFS24]. Compared to previous work on convex restrictions for structured control problems, such as block-diagonal Lyapunov functions [e.g. Con+12] and nearest quadratic invariance (QI) subsets [RL06], the SI framework is in fact shown to be the most general class of convex restrictions based on separability arguments (either of the Lyapunov function or the Youla parameter or pendants thereof) [Fur+20]. As such, the SI framework poses the least conservative method of convexification by restriction for non-convex structured control problems to date.

The non-convexity of the structured control problem stems from the structural constraint on the controller matrix, $K \in \text{sparse}(S)$, which is generally recovered from a non-linear mapping of the kind $K = YX^{-1}$, where $Y(s)$ and $X(s)$ are transfer matrices depending on some intermediate decision variable(s) of the convex optimization problem formulation. Recognizing this fact, the SI framework substitutes this offending constraint with two structural constraints $Y \in \text{sparse}(T)$ and $X \in \text{sparse}(R)$, which are known to satisfy the SI property.

Definition 4.13 (Sparsity Invariance [Fur+20])**.** Given a binary matrix S, the pair of binary matrices T and R satisfies a property of sparsity invariance with respect to S if

$$Y \in \text{sparse}(T) \text{ and } X \in \text{sparse}(R) \Rightarrow YX^{-1} \in \text{sparse}(S), \tag{4.7}$$

where X is assumed invertible.

We now guide through the relevant theorems that prove the claims made earlier about the SI framework and summarize methodologies and open questions regarding the choice of the structure matrices T

and R. As in [Fur+20], we will focus on formulations using the Youla parameterization, but the principle applies analogously to static controller design and other parameterizations such as the system-level and input-output parameterizations. The static controller design case will be discussed in additional detail in Section 4.3.2.

Consider again a distributed system given by a stabilizable generalized LTI plant P, w.l.o.g. strictly proper in $u \to e$, and partitioned as

$$\begin{bmatrix} z \\ e \end{bmatrix} = \underbrace{\begin{bmatrix} P_{zw} & P_{zu} \\ P_{ew} & P_{eu} \end{bmatrix}}_{P} \begin{bmatrix} w \\ u \end{bmatrix}. \tag{4.8}$$

To simplify the exposition, we neglect the subsystem structure and Remark 4.11 and assume more generally an arbitrary information structure, leading to the non-convex constraint $K \in \operatorname{sparse}(S)$ with $S \in \{0,1\}^{n_u \times n_e}$. Let $\mathbf{K} \subset \mathbf{RL}_\infty^{n_u \times n_e}$ be the set of stabilizing controllers, but not necessarily satisfying the structure constraint. The feasible set of solutions to Problem 4.10 then is $\mathbf{K} \cap \operatorname{sparse}(S)$. We aim to parameterize the set of stabilizing controllers. For this reason, we introduce a useful system factorization.

Lemma 4.14 (Doubly Coprime Factorization [e.g. ZDG96])**.** *For any $G \in \mathbf{RL}_\infty^{m \times n}$, there exist eight proper and stable transfer matrices defining a doubly coprime factorization of G, that is, they satisfy*

$$G = N_r M_r^{-1} = M_l^{-1} N_l, \tag{4.9}$$

$$\begin{bmatrix} U_l & -V_l \\ -N_l & M_l \end{bmatrix} \begin{bmatrix} M_r & V_r \\ N_r & U_r \end{bmatrix} = I_{m+n}. \tag{4.10}$$

Applying Lemma 4.14 to P_{eu}, we arrive at

Corollary 4.15 (Youla Parameterization [YBJ76])**.** *Given a stabilizable generalized plant P, partitioned as in (4.8) and with $P_{eu} \in \mathbf{RL}_2^{n_e \times n_u}$, the set of all stabilizing controllers $\mathbf{K} \subset \mathbf{RL}_\infty^{n_u \times n_e}$ can be written as*

$$\mathbf{K} = \left\{ (V_r - M_r Q)(U_r - N_r Q)^{-1} \mid Q \in \mathbf{RH}_\infty^{n_u \times n_e} \right\} \tag{4.11a}$$

$$= \left\{ (U_l - Q N_l)^{-1}(V_l - Q M_l) \mid Q \in \mathbf{RH}_\infty^{n_u \times n_e} \right\}, \tag{4.11b}$$

where $V_r, M_r, U_r, N_r, U_l, N_l, V_l, M_l$ form a doubly coprime factorization of P_{eu} per Lemma 4.14 and Q is a free variable called the Youla parameter. *Further, the set of stable closed-loop transfer matrices is given by*

$$\mathcal{F}_l(P, \mathbf{K}) = \left\{ T_1 - T_2 Q T_3 \mid Q \in \mathbf{RH}_\infty^{n_u \times n_e} \right\}, \tag{4.12}$$

where

$$\begin{aligned} T_1 &= P_{zw} + P_{zu} V_r M_l P_{ew} \\ T_2 &= P_{zu} M_r \\ T_3 &= M_l P_{ew}. \end{aligned} \tag{4.13}$$

Now let

$$Y_Q = (V_r - M_r Q) M_l \tag{4.14a}$$

$$X_Q = (U_R - N_r Q) M_l \tag{4.14b}$$

and note that $X_Q = I_{n_e} + P_{eu}Y_Q$ [Fur+20]. Consequently,

$$\mathbf{K} = \left\{ Y_Q X_Q^{-1} \mid (4.14), Q \in \mathbf{RH}_\infty^{n_u \times n_e} \right\}. \tag{4.15}$$

This leads to the following variation of Problem 4.10 in the Youla parameter.

Problem 4.16 (Youla Parameterized Optimal Structured Control Problem)**.** Given a generalized plant P partitioned as in (4.8), a doubly coprime factorization $V_r, M_r, U_r, N_r, U_l, N_l, V_l, M_l$ for $P_{eu} \in \mathbf{RL}_2^{n_e \times n_u}$, and a controller structure matrix S, find a $Q \in \mathbf{RH}_\infty^{n_u \times n_e}$ such that the stabilizing controller $K = Y_Q X_Q^{-1}$ with (4.14) satisfies $K \in \text{sparse}(S)$ and minimizes some performance index on the closed-loop $\mathcal{F}_l(P, K) = T_1 - T_2 Q T_3$ with (4.13).

Applying Definition 4.13, we rewrite the constraint $K \in \text{sparse}(S)$ as $Y_Q \Gamma \in \text{sparse}(T)$ and $X_Q \Gamma \in \text{sparse}(R)$ and yield the related problem

Problem 4.17 (SI Youla Parameterized Optimal Structured Control Problem)**.** Given a stabilizable generalized plant P partitioned as in (4.8), a doubly coprime factorization $V_r, M_r, U_r, N_r, U_l, N_l, V_l, M_l$ for $P_{eu} \in \mathbf{RL}_2^{n_e \times n_u}$, a controller structure matrix S, and a pair of binary matrices $T \in \{0,1\}^{n_u \times n_e}$ and $R \in \{0,1\}^{n_e \times n_e}$ satisfying the property of SI with respect to S, find a $Q \in \mathbf{RH}_\infty^{n_u \times n_e}$ such that, given (4.14) and invertible $\Gamma \in \mathbf{RL}_\infty^{n_e \times n_e}$, $Y_Q \Gamma \in \text{sparse}(T)$ and $X_Q \Gamma \in \text{sparse}(R)$ and the stabilizing controller $K = Y_Q X_Q^{-1}$, by definition $K \in \text{sparse}(S)$, minimizes some performance index on the closed-loop $\mathcal{F}_l(P, K) = T_1 - T_2 Q T_3$ with (4.13).

We now formally relate Problem 4.17 and Problem 4.16.

Lemma 4.18 ([Fur+20])**.** *For every invertible $\Gamma \in \mathbf{RL}_\infty^{n_e \times n_e}$ Problem 4.17 is convex, and every solution is also a feasible solution for Problem 4.16, i.e. Problem 4.17 is a convex restriction of Problem 4.16.*

Next, we state sufficient and necessary conditions for T and R to satisfy the SI property with respect to S.

Theorem 4.19 ([Fur+20])**.** *Let $T \in \{0,1\}^{m \times n}$ and $R \in \{0,1\}^{n \times n}$ be such that $R \geq I$. The following statements are equivalent:*

1. *$T \leq S$ and $TR^{n-1} \leq S$.*
2. *SI as per Definition 4.13 holds.*

Clearly, Theorem 4.19 leaves much room for possible choices of T and R, leading to various degrees of conservatism in the convex restriction. Given the sizes of T and R, iterating over all possible choices is generally not feasible. The optimal choice of T and R (and Γ) leading to the least amount of conservatism is indeed an open problem. However, it is recognized in [Fur+19; Fur+20] that minimizing the degree of separability in the decision variable (i.e. the Lyapunov matrix or the Youla parameter), induced by R, yields a larger candidate space and hence less conservatism. To this extend, they provide the following algorithm and theorem for computing separability maximizing choices of R given some choice of T.

Theorem 4.20 ([Fur+20])**.** *Given a binary matrix $T \in \{0,1\}^{m \times n}$, let*
$\mathbf{R}_T := \{R \in \{0,1\}^{n \times n} \mid R \geq I, TR^{n-1} \leq T\}$. Then:

1. *there exists a unique $R_T^* \in \mathbf{R}_T$ such that $R_T^* \geq R^{n-1} \;\; \forall\, R \in \mathbf{R}_T$,*
2. *R_T^* is computed via Algorithm 1.*

Algorithm 1 Generation of R_T^* [Fur+20]

Require: $T \in \{0,1\}^{m \times n}$
$R_T^* \leftarrow 1_{n \times n}$
for $i \leftarrow 1, ..., n; k \leftarrow 1, ..., m$ **do**
  if $T_{ik} = 0$ **then**
    for $j \leftarrow 1, ..., m$ **do**
      if $T_{ij} = 1$ **then**
        $(R_T^*)_{jk} \leftarrow 0$
return R_T^*

Corollary 4.21 ([Fur+20])**.** *In Problem 4.17, select $T \leq S$ and compute R_T^* using Algorithm 1. Then, for every fixed invertible $\Gamma \in \mathbf{RL}_\infty^{n_e \times n_e}$, Problem 4.17 is the tightest convex restriction of Problem 4.16 among all choices of $R \in \mathbf{R}_T$.*

Remark 4.22. While Corollary 4.21 does not provide hints how to optimally choose S and Γ, the authors in [Fur+20] note that the generic choice $T = S$ and $\Gamma = I$ is sufficient to recover or outperform solutions based on previous, more restrictive convex restrictions of Problem 4.16.

Finally, we summarize notes from [Fur+20] about the relation of the SI method to the previous notion of QI defined below. Even though the QI methodology is redundant at this point, the comparison reveals interesting facts about whether Problem 4.16 has globally optimal solutions and when and how those can be recovered by convex restrictions of the likes of Problem 4.17.

Definition 4.23 (quadratic invariance (QI) [RL06])**.** Let $G \in \mathbf{RL}_\infty^{m \times n}$ and $\mathbf{S} \subseteq \mathbf{RL}_\infty^{n \times m}$. The set $\mathbf{S}$ is called *quadratically invariant* under G if

$$KGK \in \mathbf{S} \;\; \forall\, K \in \mathbf{S}. \tag{4.16}$$

Corollary 4.24 ([Fur+20])**.** *Consider Problems 4.16 and 4.17. The following statements are equivalent:*

1. sparse(S) *is QI with respect to P_{eu}.*
2. *There exists a globally optimal controller K^* recovered from the optimal solution Q^* to Problem 4.16.*
3. *Problem 4.16 and Problem 4.17 are equivalent (and admit the same optimal solution Q^*) when choosing $T = S$, $\Gamma = I$ and computing R_T^* via Algorithm 1.*

When sparse(S) is not QI with respect to P_{eu}, the work [RM12] proposed to find the closest QI subset of S such that QI is given. By constructive examples, the authors in [Fur+20] show that, beyond Corollary 4.24, the SI framework is able to recover globally optimal solutions to Problem 4.16 where convex restrictions based on the closest QI subset fail to do so, and that the SI framework is strictly more performing than the closest QI subset method in all cases.

4.3.2 Static and Fixed-Order Output Feedback by Sparsity Invariance

Having introduced the SI framework, this section will present theory that extends the applicability of the framework from structured dynamic output feedback and structured static state-feedback to static and fixed-order output feedback. Note that the case of fixed-order output feedback is indeed not covered by the dynamic output feedback by Youla parameterization discussed in Section 4.3.1, because even finite-order restrictions on the Youla parameter Q do not necessarily yield controllers of a desired order. Fixed-order output feedback is particularly interesting for applications where limited resources are available for controller implementation and those where reduced-order controllers achieve approximately the same performance as full-order controllers, e.g. those discussed in Sections 3.2.3 and 3.4. The proposed extension of the SI framework was published in [Sch+21] and builds on the static state-feedback case [Fur+19]. It will therefore be helpful to re-pose some theorems of Section 4.3.1 in the static state-feedback setting.

Static State-Feedback

First, consider Problem 4.10 in the state-feedback setting (analogously to Problem 4.16). We will assume the following state-space realization for the generalized plant P to be controlled, as it is common with no feedthrough in the channel $u \to e$.

$$P(s) = \left[\begin{array}{c|cc} A & B_w & B_u \\ \hline C_z & D_{zw} & D_{zu} \\ I & 0 & 0 \end{array}\right]. \tag{4.17}$$

Remark 4.25. An apparent restriction is the additional assumption $D_{ew} = 0$. While it is likely possible to reformulate the problem to be convex even for $D_{ew} \neq 0$, one can also just introduce appropriate high-frequency poles in the measurement noise channel, realistically modelling band-limited measurement noise.

Problem 4.26. Given a generalized plant P with realization given by (4.17), and a controller structure matrix S, find a static feedback law $K \in \text{sparse}(S) \cap \mathbf{R}^{n_u \times n_x}$ that (by closing the control channel $e \to u$) internally stabilizes P and minimizes some performance index on the closed-loop channel $w \to z$.

A notable feature of the state-feedback synthesis problem is that, independent of the chosen performance index, the respective optimal static feedback law K and a suitable Lyapunov matrix $\mathcal{X}$ for the closed-loop system are obtained from two optimization variables $M \in \mathbf{R}^{n_u \times n_x}$ and $Y \in \mathbf{S}^{n_x}$ as

$$K = MY^{-1}, \tag{4.18}$$

$$\mathcal{X} = Y^{-1}. \tag{4.19}$$

See [SW17] for a derivation. From this, it becomes immediately clear how to apply the theory of the SI framework to the static state-feedback setting: Given a controller structure S, formulate the closed-loop LMI problem for the given performance index and constrain the optimization variables $M = KY$ and Y such that $M \in \text{sparse}(T) \cap \mathbf{R}^{n_u \times n_x}$ and $Y \in \text{sparse}(R) \cap \mathbf{S}^{n_x}$, where binary matrices T and R are SI with respect to S. For illustration, we will formulate the following theory in the $\mathbf{H}_2$ performance criterion.

Problem 4.27. Given a generalized plant P with realization given by (4.17), a controller structure matrix S, an invertible matrix $\Gamma \in \mathbf{R}^{n_x \times n_x}$, and a pair of binary matrices $T \in \{0,1\}^{n_u \times n_x}$ and $R = R^\top \in \{0,1\}^{n_x \times n_x}$ satisfying the property of SI with respect to S, find real matrices $M, Y = Y^\top \succ 0$

and $Z = Z^\top \succ 0$ minimizing $\gamma > 0$ while satisfying

$$\begin{bmatrix} (AY + B_u M)^\top + (AY + B_u M) & B_w \\ B_w^\top & -I \end{bmatrix} \prec 0, \tag{4.20a}$$

$$\begin{bmatrix} Y & (C_z Y + D_{zu} M)^\top \\ (C_z Y + D_{zu} M) & Z \end{bmatrix} \succ 0, \tag{4.20b}$$

$$\operatorname{trace}(Z) < \gamma^2, \tag{4.20c}$$

$$M\Gamma \in \operatorname{sparse}(T), \tag{4.20d}$$

$$Y\Gamma \in \operatorname{sparse}(R), \tag{4.20e}$$

$$D_{zw} = 0. \tag{4.20f}$$

Analogously to Lemma 4.18 we state

Lemma 4.28 ([Fur+19; SW17]). *For every invertible $\Gamma \in \mathbf{R}^{n_x \times n_x}$,*

1. *Problem 4.27 is convex,*
2. *the static state-feedback law $K = MY^{-1}$ ensures a closed-loop $\mathbf{H}_2$-norm less than γ, and*
3. *every solution is also a feasible solution for Problem 4.26 in the $\mathbf{H}_2$-performance case.*

Hence, Problem 4.27 is a convex restriction of Problem 4.26.

Regarding the choice of suitable matrices T, R, and Γ, the same arguments as in Section 4.3.1 apply. However, Algorithm 1 needs to be extended with an additional step in order to ensure that R_T^* is symmetric. This additional step involves setting all offending off-diagonal elements to zero.

Static and Fixed-Order Output Feedback

To proceed from static state-feedback to static and fixed-order output feedback, we utilize the fact that the controller structure matrix S can be used to restrict the controller to partial state information and that the state-space form of the generalized plant P is arbitrary in the previously stated theory. Assume now an output feedback realization for P,

$$P(s) = \left[\begin{array}{c|cc} A & B_w & B_u \\ \hline C_z & D_{zw} & D_{zu} \\ C_e & 0 & 0 \end{array}\right]. \tag{4.21}$$

The first step, transforming the static state-feedback to a static output feedback problem, is based on the following idea: if we can find a state-space form in which each output is identical to one of the system states, S can be chosen in such a way that the controller only has access to those states that represent a system output, ultimately yielding a static output feedback controller. The proposed state-space form is termed *state-output* form and defined next.

Definition 4.29 (State-Output Form). A system (2.6) with linearly independent outputs is said to be in state-output form if

$$C = \begin{bmatrix} 0_{n_x \times (n_x - n_y)} & I \end{bmatrix}. \tag{4.22}$$

Remark 4.30. For systems with $n_y > n_x$, the linearly dependent outputs can be dropped without loss of generality.

In order to compute such a state-output form, we provide

Lemma 4.31 (State-Output Transformation)**.** *For every system Σ as in (2.6) with linearly independent outputs, there exists an invertible transformation matrix $Q \in \mathbf{R}^{n_x \times n_x}$, such that the system in the transformed state $\bar{x} = Q^{-1}x$ is in state-output form. The transformation is given by*

$$Q = \begin{bmatrix} C_0 & C^\top (CC^\top)^{-1} \end{bmatrix}, \tag{4.23}$$

where C_0 is any matrix whose columns span the kernel of C.

Proof. Since Σ has linearly independent outputs, C has full row rank, guaranteeing the existence of a right inverse $C^T(CC^\top)^{-1}$ [Rao72]. The rank-nullity theorem ensures that Q is square of size n_x, and the orthogonality of the kernel and row space of C guarantees that Q is invertible. It is then easy to see that

$$CQ = \begin{bmatrix} 0_{n_x \times (n_x - n_y)} & I \end{bmatrix}, \tag{4.24}$$

as desired. □

The matrix C_0 in (4.23) is readily computed using singular value decomposition. The $\mathbf{H}_2$-optimal static structured output feedback controller is then obtained in the following way.

Proposition 4.32. *Given a generalized plant P with realization given by (4.21) and an output feedback controller structure matrix S, then a static output feedback law $K \in \mathbf{R}^{n_u \times n_e}$ guaranteeing a closed-loop $\mathbf{H}_2$-norm less than γ can be recovered from the solution of Problem 4.27 when first transforming P to state-output form in e and assuming the state-feedback structure matrix*

$$\bar{S} = \begin{bmatrix} 0_{n_x - n_e} & 0 \\ 0 & S \end{bmatrix}. \tag{4.25}$$

The static output feedback law K is recovered from the solution of Problem 4.27 – according to the position of S in $\bar{S}$ – as the lower-right $(n_u \times n_e)$-sized block of MY^{-1}, i.e.

$$K = \begin{bmatrix} 0 \\ I_{n_u} \end{bmatrix}^\top MY^{-1} \begin{bmatrix} 0 \\ I_{n_e} \end{bmatrix}. \tag{4.26}$$

Proof. It follows from (4.25) and T and R being SI with respect to $\bar{S}$ that

$$MY^{-1} = \begin{bmatrix} 0_{n_x - n_e} & 0 \\ 0 & K \end{bmatrix}, \tag{4.27}$$

with $K \in \text{sparse}(S)$. Let x be the state vector of P. Because P is in state-output form in e, we have

$$MY^{-1}x = KC_e x = Ke, \tag{4.28}$$

which shows that K is the desired structured static output feedback law guaranteeing a closed-loop $\mathbf{H}_2$-norm less than γ. □

Remark 4.33. It may seem invalid to apply the state-output transformed plant to Problem 4.27, given that it is incompatible with (4.17). However, there is no issue here since C_e does not occur explicitly in (4.20).

In practice, static feedback laws are usually of limited interest because of their inability to eliminate steady-state tracking errors. As a countermeasure, integral action is frequently added in the controller, leading to PI control. It is therefore interesting to extend the presented approach to fixed-order output feedback. Since the dynamic output feedback synthesis approach presented in Section 4.3.1 doesn't allow to dictate the controller order, we employ a different trick here to convert the static feedback design problem into a fixed-order control problem. First, the plant is augmented with additional free states to be used as controller states. Then we again exploit the control over S to enforce the desired distributed controller structure.

Proposition 4.34. *Given a generalized plant P with realization given by (4.8) and an output feedback controller structure matrix $S \in \{0,1\}^{n_u \times n_e}$, then a dynamic output feedback controller $K \in \operatorname{sparse}(S)$ of fixed order n_{x_K} and guaranteeing a closed-loop $\mathbf{H}_2$-norm less than γ can be recovered from the solution of Problem 4.27 when assuming the state-feedback structure matrix*

$$\bar{S} = \begin{bmatrix} 0_{n_x - n_e} & 0 & 0 \\ 0 & S & S_C \\ 0 & S_B & S_A \end{bmatrix}, \tag{4.29}$$

where S_A, S_B, S_C are chosen to achieve the desired controller separation, and taking the state-output transformed and augmented plant

$$\bar{P}(s) = \left[\begin{array}{c|cc} \bar{A} & \bar{B}_w & \bar{B}_u \\ \hline \bar{C}_z & D_{zw} & D_{zu} \\ \bar{C}_e & 0 & 0 \end{array}\right] \tag{4.30}$$

with

$$\begin{aligned} \bar{A} &= \begin{bmatrix} Q^{-1}AQ & 0 \\ 0 & 0_{n_{x_K}} \end{bmatrix}, \\ \bar{B}_w &= \begin{bmatrix} Q^{-1}B_w \\ 0 \end{bmatrix}, \bar{B}_u = \begin{bmatrix} Q^{-1}B_u & 0 \\ 0 & I \end{bmatrix}, \\ \bar{C}_z &= \begin{bmatrix} C_z Q & 0 \end{bmatrix}, \bar{C}_e = \begin{bmatrix} C_e Q & 0 \\ 0 & I \end{bmatrix} = \begin{bmatrix} 0 & I_{n_e + n_{x_K}} \end{bmatrix}, \end{aligned} \tag{4.31}$$

where Q is the state-output transformation matrix per Lemma 4.31 with respect to C_e. The fixed-order output feedback controller K is recovered from the solution of Problem 4.27 by partitioning MY^{-1} as

$$MY^{-1} = \begin{bmatrix} 0_{n_x - n_e} & 0 & 0 \\ 0 & D_K & C_K \\ 0 & B_K & A_K \end{bmatrix} \tag{4.32}$$

and letting

$$K(s) = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array}\right]. \tag{4.33}$$

Proof. Let $\tilde{P}$ be $\bar{P}$ with $\bar{C}_e$ replaced by I_{n_x}. Note that $\mathcal{F}_l(\tilde{P}, MY^{-1}) = \mathcal{F}_l(P, K)$, which shows that K is indeed a fixed-order output feedback controller of order n_{x_K} guaranteeing a closed-loop $\mathbf{H}_2$-norm less than γ. □

Remark 4.35. Naturally, Proposition 4.34 can also be used to design fixed-order state-feedback controllers.

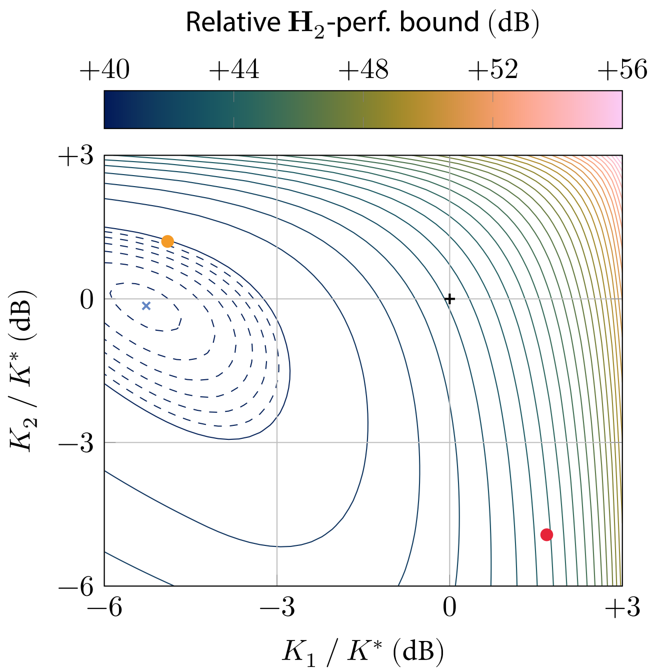

Figure 4.6: $\mathbf{H}_2$-performance guaranteed by the SI static decentralized feedback controller relative to the benchmark performance achieved by uniform optimal tuning.

4.3.3 Numerical Examples

We apply the newly developed output feedback design method of Proposition 4.32 to Examples 4.1 and 4.2. For lack of a more educated choice, we select $S = I$ for decentralized control, $\Gamma = I$, and use Algorithm 1 to compute R_T^*. To aid comparison, controller gains and performance levels are again expressed in dB relative to the optimal proportional feedback gain K^* for a single PLL system and the $\mathbf{H}_2$-norm achieved by the static feedback law diag(K^*, K^*), respectively. Bounds on the $\mathbf{H}_2$-norm are expressed relative to the respective actual $\mathbf{H}_2$-norm achieved by the controller yielded from the SI synthesis method.

Recall that Example 4.1 revealed that, given a distributed system of two daisy-chained PLLs as in Figure 4.1 and parameterized as in Table 3.1, the feedback gain of the first PLL system gain can be increased arbitrarily. The problem is therefore essentially ill-posed, as the control effort is not penalized in this formulation (cheap control) and the solution therefore unbounded. It is therefore not entirely unexpected that the optimal SI-based feedback gains for this case are found to be $K_1^{(\mathrm{SI})} = K^* + 68.6$ dB, $K_2^{(\mathrm{SI})} = K^* + 170.0$ dB. The $\mathbf{H}_2$-norm achieved by the SI controller is increased by 21.67 dB compared to the baseline performance. The guaranteed $\mathbf{H}_2$-norm is 89 % higher than the achieved performance, which shows significant conservatism in the performance bound. Manual inspection of the solution space around the SI optimum reveals that the convexified cost function of this ill-posed problem is extremely flat over a large range of controller gains, leading to an essentially arbitrary solution.

The story is very much different for the case of Example 4.2, where the controller gains are limited by an additional pole in G at 10 kHz. Figure 4.6 shows the SI convexified cost function over the same parameter range as Figure 4.2b. The two minima of the original cost function are marked for reference. The SI controller achieves an improved $\mathbf{H}_2$-norm relative to the baseline of -2.11 dB. The

N	SI Controller Perf. (dB)	Rel. Improvement (%)	Optimality Gap (%)	Perf. Bound Gap (dB)
2	-2.11	21.55	30.07	40.34
3	-11.79	74.26	38.85	41.81
4	-22.39	92.40	61.27	41.90
5	-33.18	97.81	103.28	41.87
6	-48.58	99.63	56.01	41.51
7	-57.35	99.86	168.71	40.36
8	-72.17	99.98	134.13	38.76

Table 4.1: Relative performance of the decentralized SI static output feedback controller for the system of Example 4.2 w.r.t. the uniform tuning benchmark vs. N, the number of PLL subsystems in the chain.

gap from guaranteed bound to actual $\mathbf{H}_2$-norm is however even more significant than in the former case, exceeding 40 dB. We note however, that a slightly different parameterized numerical example in the original paper [Sch+21] yielded much less conservative bounds, and it is very well possible that the conservatism of the convex restriction may be reduced by scaling the model differently, using different tolerances for matrix definiteness, or finding improved T and Γ. Variation of these hyperparameters also control the shape of the convexified cost function to the extent that it may be possible to further improve the solution. As in [Sch+21], the convexified cost function is less steep along the direction connecting the two local minima of the original cost function, but overall skewed to the inferior local minimum.

When considering daisy-chains of more than two PLLs, the complexity of the problem increases and gridding the parameter space becomes infeasible. Table 4.1 summarizes the performance of the SI controller for $N \geq 2$ up to $N = 8$. We note that the performance bound gap remains approximately constant, while the SI controller gains significant improvement over the baseline solution with increasing N. For all cases, the SI structured output feedback controller synthesis SDP problem was solved in negligible time using MATLAB, YALMIP, and MOSEK [The24a; Löf04; MOS24] on a standard personal computer. The optimality gap was determined using a local minimum certified solution found with systune [The24b], using nonsmooth optimization and significantly longer optimization times. Despite outliers due to finite precision induced numerical randomness, the optimality gap appears to increase monotonically with N, however not nearly as badly as the baseline tuning solution, which scales exponentially. It depends on the specific application whether it is feasible or worthwhile to scan over the solution space or use the SI-based single-step LMI synthesis method and accept the incurred conservatism.

Remark 4.36. One critical drawback of using systune in this case is the number of local minima in the cost function, as the best minimum can only be identified by starting the optimization from sufficiently many distinct parameter initializations. In [Sch+21], the number of local minima was shown to scale with $N!$, based on an argument on the combinatorial complexity of the original structured optimization problem when expressed using the Lyapunov equation with a linear cost function and quadratic equality constraints. Thus, for $N = 8$, one can expect to have over $40\,000$ local minima in the cost function, needing many more initial value iterates to fully explore.

4.4 Structured Robust Control by Sparsity Invariance

4.4.1 Full-Block S-Procedure LMI Conditions for LTV Uncertainties

Structured control problems often arise for large-scale distributed systems with correspondingly large number of model parameters. It is generally unrealistic to assume that the parameters are known exactly at all times, as component faults and degradation may slowly or suddenly change the model, deteriorating the performance of the synthesized feedback solution. This is especially problematic for $\mathbf{H}_2$-performance optimized systems because, as it was famously proved by Doyle for the equivalent LQG problem [Doy78], there are no meaningful stability margins that can be derived from the nominally optimizing solution. In other words, even minor alterations of the plant dynamics can have detrimental if not fatal (destabilizing) effects on the closed-loop. A ubiquitous goal in real-world distributed control is therefore not only satisfying the information constraints on the synthesized controller, but also securing a certain level of robustness against variation in critical plant parameters. We show in this section that it is possible to achieve such robustness, at least against linear parameter-varying uncertainty, using robust control techniques based on static multipliers and the full-block S-procedure, cf. [Sch00], in conjunction with the SI framework as introduced in Sections 4.3.1 and 4.3.2.

In Sections 4.3.1 and 4.3.2, the generalized distributed plant P was considered LTI. We now consider the composing subsystems to be uncertain LTV, more specifically as well-posed LFTs of an LTI part and an uncertain, time-varying, real parameter matrix each. For each subsystem $i = 1, ..., k$, let $\mathbf{\Delta}_i \subset \mathbf{R}^{n_{p_i} \times n_{q_i}}$ with $0 \in \mathbf{\Delta}_i$ be the sets of admissible parameter values, such that $P_i(\Delta_i) = \mathcal{F}_u(P_i, \Delta_i(t))$ where $\Delta_i(t)\colon [0, \infty) \to \mathbf{\Delta}_i$ is a possible realization of the LTV subsystem. There's no uncertainty associated with the network operator. Updating (4.2), we get

$$\begin{aligned} P(\Delta) &= \mathcal{F}_u(\text{daug}(P_1(\Delta_i), ..., P_k(\Delta_k)), \mathcal{N}) \\ &= \mathcal{F}_u(\mathcal{F}_u(\text{daug}(P_1, ..., P_k), \underbrace{\text{diag}(\Delta_i, ..., \Delta_k)}_{\Delta}), \mathcal{N}). \end{aligned} \tag{4.34}$$

With the LTV distributed system model in place and letting $\mathbf{\Delta} = \text{diag}(\mathbf{\Delta}_1, ..., \mathbf{\Delta}_k)$, we pose the robust version of Problem 4.10:

Problem 4.37. Given an uncertain distributed system $P(\Delta)$ as in (4.34) and a controller structure matrix $S \in \{0, 1\}^{k \times k}$, find a distributed controller $K \in \text{sparse}(S)$ that (by closing the control channel $e \to u$) internally stabilizes $P(\Delta)$ and minimizes the worst case performance on the channel $w \to z$ over all possible parameter trajectories $\Delta\colon [0, \infty) \to \mathbf{\Delta}$.

Using again the example of the $\mathbf{H}_2$-norm performance index, solving this extended problem requires to render the supremum condition in

$$\min_{K \in \text{sparse}(S)} \sup_{\Delta \in \mathbf{\Delta}} \|\mathcal{F}_l(P(\Delta), K)\|_2 \tag{4.35}$$

explicit. This is achieved by the full-block S-procedure [Sch00], which equivalently reformulates the stability and performance tests over the entirety of $\mathbf{\Delta}$ as a search for suitable multipliers from the set

$$\mathbf{M} := \left\{ M \in \mathbf{S}^{(n_p + n_q)} \mid \begin{bmatrix} \Delta \\ I \end{bmatrix}^\top M \begin{bmatrix} \Delta \\ I \end{bmatrix} \succ 0 \ \ \forall\, \Delta \in \mathbf{\Delta} \right\} \tag{4.36}$$

satisfying a to-be-defined LMI condition.

Continuing the SI approach, we derive the synthesis LMIs from the analysis problem. First, let P be such that $P(\Delta) = \mathcal{F}_u(P, \Delta)$ and admit the state-space representation

$$\begin{bmatrix} q \\ z \\ e \end{bmatrix} = \underbrace{\left[\begin{array}{c|ccc} A & B_p & B_w & B_u \\ \hline C_q & D_{qp} & D_{qw} & D_{qu} \\ C_z & D_{zp} & D_{zw} & D_{zu} \\ C_e & D_{ep} & D_{ew} & 0 \end{array}\right]}_{P(s)} \begin{bmatrix} p \\ w \\ u \end{bmatrix}. \tag{4.37}$$

Then, using the full-block S-procedure, the following well-known result on robust stability and performance analysis is derived:

Corollary 4.38 ([Sch00]). *Consider a generalized LTV system $P(\Delta) = \mathcal{F}_u(P, \Delta)$ as in (4.37), and a stabilizing output feedback controller $K \in \mathbf{RL}_\infty^{n_u \times n_e}$ admitting state-space realization (A_K, B_K, C_K, D_K). The interconnection of $\mathcal{F}_l(P, K)$ and Δ is well-posed and has at all times $\mathbf{H}_2$-norm less than γ if and only if there exist matrix variables $X \succ 0$, $Z \succ 0$ and multipliers $M, M_p \in \mathbf{M}$, partitioned conforming to $\begin{bmatrix} \Delta & I \end{bmatrix}^\top$ as*

$$M = \begin{bmatrix} W & U \\ U^\top & V \end{bmatrix}, \quad M_p = \begin{bmatrix} W_p & U_p \\ U_p^\top & V_p \end{bmatrix} \tag{4.38}$$

such that

$$[\star]^\top \left[\begin{array}{cc:cc:cc} 0 & X & 0 & 0 & 0 & 0 \\ X & 0 & 0 & 0 & 0 & 0 \\ \hdashline 0 & 0 & W & U & 0 & 0 \\ 0 & 0 & U^\top & V & 0 & 0 \\ \hdashline 0 & 0 & 0 & 0 & -I & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array}\right] \left[\begin{array}{ccc} I & 0 & 0 \\ \mathcal{A} & \mathcal{B}_p & \mathcal{B}_w \\ \hdashline 0 & I & 0 \\ \mathcal{C}_q & \mathcal{D}_{qp} & \mathcal{D}_{qw} \\ \hdashline 0 & 0 & I \\ \mathcal{C}_z & \mathcal{D}_{zp} & \mathcal{D}_{zw} \end{array}\right] \prec 0 \tag{4.39a}$$

$$[\star]^\top \left[\begin{array}{c:cc:c} -X & 0 & 0 & 0 \\ \hdashline 0 & W_p & U_p & 0 \\ 0 & U_p^\top & V_p & 0 \\ \hdashline 0 & 0 & 0 & Z^{-1} \end{array}\right] \left[\begin{array}{cc} I & 0 \\ \hdashline 0 & I \\ \mathcal{C}_q & \mathcal{D}_{qp} \\ \hdashline \mathcal{C}_z & \mathcal{D}_{zp} \end{array}\right] \prec 0 \tag{4.39b}$$

$$\mathrm{trace}(Z) < \gamma^2 \tag{4.39c}$$

$$\mathcal{D}_{zw} = 0, \tag{4.39d}$$

where

$$\left[\begin{array}{c|cc} \mathcal{A} & \mathcal{B}_p & \mathcal{B}_w \\ \hline \mathcal{C}_q & \mathcal{D}_{qp} & \mathcal{D}_{qw} \\ \mathcal{C}_z & \mathcal{D}_{zp} & \mathcal{D}_{zw} \end{array}\right] := \left[\begin{array}{cc|cc} A + B_u D_K C_e & B_u C_K & B_p + B_u D_K D_{ep} & B_w + B_u D_K D_{ew} \\ B_K C_e & A_K & B_K D_{ep} & B_K D_{ew} \\ \hline C_q + D_{qu} D_K C_e & D_{qu} C_K & D_{qp} + D_{qu} D_K D_{ep} & D_{qw} + D_{qu} D_K D_{ew} \\ C_z + D_{zu} D_K C_e & D_{zu} C_K & D_{zp} & D_{zw} + D_{zu} D_K D_{ew} \end{array}\right] = \mathcal{F}_l(P, K). \tag{4.40}$$

Corollary 4.38 is the general result for dynamic output feedback. For the synthesis case with variable controller matrices A_K, B_K, C_K, D_K, linearizing the LMIs (4.39) requires a change of variables incompatible with the SI framework (cf. Theorem 3.10). However, as in Section 4.3.2, the structured fixed-order output feedback synthesis problem can be derived from the static state-feedback problem with a reduced set of linearization variables and exploiting the SI property of Definition 4.13. In order to arrive at the desired linearization, we require the following two lemmas.

Lemma 4.39 (Linearization Lemma [SW17]). *Consider constant matrices A and S and matrices $B(\nu)$, $Q(\nu) = Q(\nu)^\top$ depending affinely on some parameter ν, and suppose that $R(\nu)$ (not necessarily affine in ν) can be decomposed as $TU(\nu)^{-1}T^\top$, $U(\nu)$ being affine. Then the non-linear matrix inequalities*

$$U(\nu) \succ 0, \quad \begin{bmatrix} A \\ B(\nu) \end{bmatrix}^\top \begin{bmatrix} Q(\nu) & S \\ S^\top & R(\nu) \end{bmatrix} \begin{bmatrix} A \\ B(\nu) \end{bmatrix} \prec 0 \tag{4.41}$$

are equivalent to the linear matrix inequality

$$\begin{bmatrix} A^\top Q(\nu)A + A^\top SB(\nu) + B(\nu)^\top S^\top A & B(\nu)^\top T \\ T^\top B(\nu) & -U(\nu) \end{bmatrix} \prec 0. \tag{4.42}$$

Lemma 4.40 (Dualization Lemma [SW17]). *Let $P \in \mathbf{S}^n$ be a non-singular matrix and let $\mathbf{U}$ and $\mathbf{V}$ be two complementary subspaces whose sum equals $\mathbf{R}^n$. Then*

$$x^\top Px < 0 \;\; \forall\, x \in \mathbf{U} \setminus \{0\} \quad \textit{and} \quad x^\top Px \geq 0 \;\; \forall\, x \in \mathbf{V} \tag{4.43}$$

is equivalent to

$$x^\top P^{-1}x > 0 \;\; \forall\, x \in \mathbf{U}^\perp \setminus \{0\} \quad \textit{and} \quad x^\top P^{-1}x \leq \;\; \forall\, x \in \mathbf{V}^\perp. \tag{4.44}$$

Applying Lemma 4.40 to (4.36) yields the set of dualized multipliers

$$\tilde{\mathbf{M}} := \left\{ \tilde{M} \in \mathbf{S}^{n_p+n_q} \mid \begin{bmatrix} I \\ -\Delta \end{bmatrix}^\top \tilde{M} \begin{bmatrix} I \\ -\Delta \end{bmatrix} \prec 0 \;\; \forall\, \Delta \in \mathbf{\Delta} \right\}. \tag{4.45}$$

Proposition 4.41 ([Sch+22]). *Given a generalized LTV system $P(\Delta) = \mathcal{F}_u(P, \Delta)$ as in (4.34) with $D_{ep} = D_{ew} = D_{eu} = 0$, well-posedness, uniform exponential stability and an upper bound on the closed-loop $\mathbf{H}_2$-norm less than γ are guaranteed by a static state-feedback controller $K \in$ sparse(S) for all trajectories of the parameter matrix $\Delta(t) \in \mathbf{\Delta}$, if there exist binary matrices T and R, real matrix variables M, $Y \succ 0$, $Z \succ 0$, and multipliers $\tilde{M}, \tilde{M}_p \in \tilde{\mathbf{M}}$ partitioned conforming to $\begin{bmatrix} I & -\Delta \end{bmatrix}^\top$ as*

$$\tilde{M} = \begin{bmatrix} \tilde{W} & \tilde{U} \\ \tilde{U}^\top & \tilde{V} \end{bmatrix}, \quad \tilde{M}_p = \begin{bmatrix} \tilde{W}_p & \tilde{U}_p \\ \tilde{U}_p^\top & \tilde{V}_p \end{bmatrix} \tag{4.46}$$

satisfying for all real and non-singular Γ*:*

$$\begin{gathered} T \leq S, \quad \mathrm{struct}(TR^{n-1}) \leq S, \\ M\Gamma \in \mathrm{sparse}(T), \quad Y\Gamma \in \mathrm{sparse}(R) \end{gathered} \tag{4.47a}$$

$$[\star]^\top \left[\begin{array}{cc|cc|c} 0 & I & 0 & 0 & 0 \\ I & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & \tilde{W} & \tilde{U} & 0 \\ 0 & 0 & \tilde{U}^\top & \tilde{V} & 0 \\ \hline 0 & 0 & 0 & 0 & -I \end{array}\right] \left[\begin{array}{cc} -(AY + B_u M)^\top & -(C_q Y + D_{qu} M)^\top \\ I & 0 \\ \hline -B_p^\top & -D_{qp}^\top \\ 0 & I \\ \hline -B_w^\top & -D_{qw}^\top \end{array}\right] \succ 0 \tag{4.47b}$$

$$\begin{bmatrix} Y & -(C_q Y + D_{qu} M)^\top & -(C_z Y + D_{zu} M)^\top \\ -(C_q Y + D_{qu} M) & D_{qp}\tilde{W}_p D_{qp}^\top - D_{qp}\tilde{U}_p - (D_{qp}\tilde{U}_p)^\top + \tilde{V}_p & D_{qp}\tilde{W}_p D_{zp}^\top - (D_{zp}\tilde{U}_p)^\top \\ -(C_z Y + D_{zu} M) & D_{zp}\tilde{W}_p D_{qp}^\top - D_{zp}\tilde{U}_p & D_{zp}\tilde{W}_p D_{zp}^\top + Z \end{bmatrix} \succ 0 \tag{4.47c}$$

$$\mathrm{trace}(Z) < \gamma^2. \tag{4.47d}$$

Proof. Naturally, Equation (4.47a) ensures by Theorem 4.19 that $K = MY^{-1} \in \mathrm{sparse}(S)$.

Write the analysis LMIs (4.39) for the static state-feedback case, i.e. $C_e = I, D_{ep} = D_{ew} = D_{eu} = 0$, and strip the redundant zero row and column:

$$[\star]^\top \left[\begin{array}{cc|cc|c} 0 & X & 0 & 0 & 0 \\ X & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & W & U & 0 \\ 0 & 0 & U^\top & V & 0 \\ \hline 0 & 0 & 0 & 0 & -I \end{array}\right] \left[\begin{array}{ccc} I & 0 & 0 \\ A + B_u K & B_p & B_w \\ \hline 0 & I & 0 \\ C_q + D_{qu} K & D_{qp} & D_{qw} \\ \hline 0 & 0 & I \end{array}\right] \prec 0 \tag{4.48a}$$

$$[\star]^\top \left[\begin{array}{c|cc|c} -X & 0 & 0 & 0 \\ \hline 0 & W_p & U_p & 0 \\ 0 & U_p^\top & V_p & 0 \\ \hline 0 & 0 & 0 & Z^{-1} \end{array}\right] \left[\begin{array}{cc} I & 0 \\ \hline 0 & I \\ C_q + D_{qu} K & D_{qp} \\ \hline C_z + D_{zu} K & D_{zp} \end{array}\right] \prec 0 \tag{4.48b}$$

$$\mathrm{trace}(Z) < \gamma^2. \tag{4.48c}$$

Permuting the rows and columns in (4.48a), we yield

$$[\star]^\top \left[\begin{array}{ccc|cc} 0 & 0 & 0 & X & 0 \\ 0 & W & 0 & 0 & U \\ 0 & 0 & -I & 0 & 0 \\ \hline X & 0 & 0 & 0 & 0 \\ 0 & U^\top & 0 & 0 & V \end{array}\right] \left[\begin{array}{ccc} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \\ \hline A + B_u K & B_p & B_w \\ C_q + D_{qu} K & D_{qp} & D_{qw} \end{array}\right] \prec 0. \tag{4.49}$$

Note that because (4.36) and $0 \in \mathbf{\Delta}$, it follows that $V \succ 0$ and $V_p \succ 0$. But then

$$\left[\begin{array}{ccc|cc} 0 & 0 & 0 & X & 0 \\ 0 & W & 0 & 0 & U \\ 0 & 0 & -I & 0 & 0 \\ \hline X & 0 & 0 & 0 & 0 \\ 0 & U^\top & 0 & 0 & V \end{array}\right] \prec 0 \quad \text{on} \quad \mathrm{im} \left(\begin{array}{ccc} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \\ \hline A + B_u K & B_p & B_w \\ C_q + D_{qu} K & D_{qp} & D_{qw} \end{array}\right) \tag{4.50a}$$

and

$$\left[\begin{array}{ccc|cc} 0 & 0 & 0 & X & 0 \\ 0 & W & 0 & 0 & U \\ 0 & 0 & -I & 0 & 0 \\ \hline X & 0 & 0 & 0 & 0 \\ 0 & U^\top & 0 & 0 & V \end{array}\right] \succeq 0 \quad \text{on} \quad \operatorname{im}\left(\begin{array}{cc} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \hline I & 0 \\ 0 & I \end{array}\right). \tag{4.50b}$$

Likewise for (4.48b)

$$\left[\begin{array}{cc|cc} -X & 0 & 0 & 0 \\ 0 & W_p & U_p & 0 \\ \hline 0 & U_p^\top & V_p & 0 \\ 0 & 0 & 0 & Z^{-1} \end{array}\right] \prec 0 \quad \text{on} \quad \operatorname{im}\left(\begin{array}{cc} I & 0 \\ 0 & I \\ \hline C_q + D_{qu}K & D_{qp} \\ C_z + D_{zu}K & D_{zp} \end{array}\right) \tag{4.51a}$$

and

$$\left[\begin{array}{cc|cc} -X & 0 & 0 & 0 \\ 0 & W_p & U_p & 0 \\ \hline 0 & U_p^\top & V_p & 0 \\ 0 & 0 & 0 & Z^{-1} \end{array}\right] \succeq 0 \quad \text{on} \quad \operatorname{im}\left(\begin{array}{cc} 0 & 0 \\ 0 & 0 \\ \hline I & 0 \\ 0 & I \end{array}\right). \tag{4.51b}$$

Since, in addition,

$$\operatorname{im}\begin{pmatrix} I \\ \bullet \end{pmatrix} + \operatorname{im}\begin{pmatrix} 0 \\ I \end{pmatrix} = \mathbf{R}^n, \tag{4.52}$$

the LMIs (4.48) may be readily dualized with Lemma 4.40. Utilizing

$$\operatorname{im}\begin{pmatrix} I \\ M \end{pmatrix}^\perp = \ker\begin{pmatrix} I & M^T \end{pmatrix} = \operatorname{im}\begin{pmatrix} -M^\top \\ I \end{pmatrix} \tag{4.53}$$

we find matrices spanning the complementary subspaces

$$\operatorname{im}\left(\begin{array}{ccc} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \\ \hline A + B_uK & B_p & B_w \\ C_q + D_{qu}K & D_{qp} & D_{qw} \end{array}\right)^\perp = \operatorname{im}\left(\begin{array}{cc} -(A + B_uK)^\top & -(C_q + D_{qu}K)^\top \\ -B_p^\top & -D_{qp}^\top \\ -B_w^\top & -D_{qw}^\top \\ \hline I & 0 \\ 0 & I \end{array}\right) \tag{4.54a}$$

$$\operatorname{im}\left(\begin{array}{cc} I & 0 \\ 0 & I \\ \hline C_q + D_{qu}K & D_{qp} \\ C_z + D_{zu}K & D_{zp} \end{array}\right)^\perp = \operatorname{im}\left(\begin{array}{cc} -(C_q + D_{qu}K)^\top & -(C_z + D_{zu}K)^\top \\ -D_{qp}^\top & -D_{zp}^\top \\ \hline I & 0 \\ 0 & I \end{array}\right). \tag{4.54b}$$

Reverting the row and column permutation, the dualized LMIs read

$$[\star]^\top \left[\begin{array}{cc|cc|c} 0 & Y & 0 & 0 & 0 \\ Y & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & \tilde{W} & \tilde{U} & 0 \\ 0 & 0 & \tilde{U}^\top & \tilde{V} & 0 \\ \hline 0 & 0 & 0 & 0 & -I \end{array}\right] \left[\begin{array}{cc} -(A+B_uK)^\top & -(C_q+D_{qu}K)^\top \\ I & 0 \\ \hline -B_p^\top & -D_{qp}^\top \\ 0 & I \\ \hline -B_w^\top & -D_{qw}^\top \end{array}\right] \succ 0 \tag{4.55a}$$

$$[\star]^\top \left[\begin{array}{c|cc|c} -Y & 0 & 0 & 0 \\ \hline 0 & \tilde{W}_p & \tilde{U}_p & 0 \\ 0 & \tilde{U}_p^\top & \tilde{V} & 0 \\ \hline 0 & 0 & 0 & Z \end{array}\right] \left[\begin{array}{cc} -(C_q+D_{qu}K)^\top & -(C_z+D_{zu}K)^\top \\ \hline -D_{qp}^\top & -D_{zp}^\top \\ I & 0 \\ \hline 0 & I \end{array}\right] \succ 0, \tag{4.55b}$$

where $Y = X^{-1}$ and

$$\begin{bmatrix} \tilde{W} & \tilde{U} \\ \tilde{U}^\top & \tilde{V} \end{bmatrix} = \tilde{M} := M^{-1} \quad \text{and} \quad \begin{bmatrix} \tilde{W}_p & \tilde{U}_p \\ \tilde{U}_p^\top & \tilde{V}_p \end{bmatrix} = \tilde{M}_p := M_p^{-1}. \tag{4.56}$$

Naturally $\tilde{M}, \tilde{M}_p \in \tilde{\mathbf{M}}$.

Before performing the linearizing change of variables $M = KY$, we need to apply Lemma 4.39 to (4.55b) to in fact yield a linear LMI. Permuting Y into the bottom-right corner, application of Lemma 4.39 yields

$$\begin{bmatrix} [\star]^\top \begin{bmatrix} \tilde{W}_p & \tilde{U}_p & 0 \\ \tilde{U}_p^\top & \tilde{V}_p & 0 \\ 0 & 0 & Z \end{bmatrix} \begin{bmatrix} -D_{qp}^\top & -D_{zp}^\top \\ I & 0 \\ 0 & I \end{bmatrix} & \begin{bmatrix} -(C_q+D_{qu}K)^\top & -(C_z+D_{zu}K)^\top \end{bmatrix}^\top \\ \begin{bmatrix} -(C_q+D_{qu}K)^\top & -(C_z+D_{zu}K)^\top \end{bmatrix} & Y^{-1} \end{bmatrix} \succ 0. \tag{4.57}$$

Permuting back and performing congruence transformation with $\operatorname{diag}(Y, I)$ yields

$$\begin{bmatrix} Y & -Y(C_q+D_{qu}K)^\top & -Y(C_z+D_{zu}K)^\top \\ -(C_q+D_{qu}K)Y & D_{qp}\tilde{W}_pD_{qp}^\top - D_{qp}\tilde{U}_p - \tilde{U}_p^\top D_{qp}^\top + \tilde{V}_p & D_{qp}\tilde{W}_pD_{zp}^\top - \tilde{U}_p^\top D_{zp}^\top \\ -(C_z+D_{zu}K)Y & D_{zp}\tilde{W}_pD_{qp}^\top - D_{zp}\tilde{U}_p & D_{zp}\tilde{W}_pD_{zp}^\top + Z \end{bmatrix} \succ 0. \tag{4.58}$$

Taking (4.55a) and (4.58) and substituting $M = KY$ yields (4.47b) and (4.47c). Hence, by Corollary 4.38, $P(\Delta)$ is well-posed and $K = MY^{-1}$ guarantees uniform exponential stability and $\|\mathcal{F}_l(P(\Delta), K)\|_2 < \gamma$, if (4.47d) is satisfied. □

Remark 4.42. The extension to fixed-order output feedback follows in complete analogy to Proposition 4.34.

Remark 4.43. The perhaps unconventional extra restriction $D_{ep} = 0$ is put into perspective by acknowledging the fact that the degrees of freedom in the construction of the uncertain LFR of $P(\Delta)$ can be exploited to ensure $D_{ep} = 0$ without conservatism or even additional strict-properness inducing dynamics. As an example, a system with unknown steady-state gain may be constructed with $D_{ep} \neq 0$ or *equivalently* $D_{qu} \neq 0$.

Even though the LMIs of Corollary 4.38 and Proposition 4.41 constitute a convex optimization problem, the set membership condition on the (dualized) multipliers generally translates to infinitely many constraints to be verified and is therefore not tractable. Tractability is commonly restored, generally at the cost of additional conservatism, by considering in place of $\mathbf{\Delta}$ a convex hull $\mathbf{\Delta}_c$ approximating $\mathbf{\Delta}$, defined through an arbitrary but finite number of corner vertices $\Delta_1, \ldots, \Delta_J \in \mathbf{\Delta}$ [Sch00]. Let

$$\mathbf{M}_c := \left\{ M \in \mathbf{S}^{n_p+n_q} \mid \begin{bmatrix} \Delta \\ I \end{bmatrix}^\top M \begin{bmatrix} \Delta \\ I \end{bmatrix} \succ 0 \;\; \forall\, \Delta \in \mathbf{\Delta}_c \right\}, \tag{4.59a}$$

$$\tilde{\mathbf{M}}_c := \left\{ \tilde{M} \in \mathbf{S}^{n_p+n_q} \mid \begin{bmatrix} I \\ -\Delta \end{bmatrix}^\top \tilde{M} \begin{bmatrix} I \\ -\Delta \end{bmatrix} \prec 0 \;\; \forall\, \Delta \in \mathbf{\Delta}_c \right\}. \tag{4.59b}$$

Then, membership in the strict inclusions

$$\mathbf{M}_c^\circ := \{ M \in \mathbf{M}_c \mid W \prec 0 \} \subseteq \mathbf{M}_c\,, \tag{4.60a}$$

$$\tilde{\mathbf{M}}_c^\circ := \left\{ \tilde{M} \in \tilde{\mathbf{M}}_c \mid \tilde{V} \succ 0 \right\} \subseteq \tilde{\mathbf{M}}_c \tag{4.60b}$$

$$\tag{4.60c}$$

is shown equivalently by the finite number of LMI conditions

$$W \prec 0 \quad \text{and} \quad \begin{bmatrix} \Delta_j \\ I \end{bmatrix}^\top M \begin{bmatrix} \Delta_j \\ I \end{bmatrix} \succ 0 \;\; \forall\, j = 1, \ldots, J, \tag{4.61a}$$

$$\tilde{V} \succ 0 \quad \text{and} \quad \begin{bmatrix} I \\ -\Delta_j \end{bmatrix}^\top \tilde{M} \begin{bmatrix} I \\ -\Delta_j \end{bmatrix} \succ 0 \;\; \forall\, j = 1, \ldots, J, \tag{4.61b}$$

respectively. Of course, the additional conservatism afforded with this tractable test is traced back to how closely $\mathbf{\Delta}_c$ approximates $\mathbf{\Delta}$ and to the additional convexity condition $W \prec 0$ ($\tilde{V} \succ 0$). The contribution of the latter is in fact null, if $D_{qp} = 0$, as then $W \prec 0$ ($\tilde{V} \succ 0$) is already implied and consequently $\mathbf{M}_c^\circ = \mathbf{M}_c$ ($\tilde{\mathbf{M}}_c^\circ = \tilde{\mathbf{M}}_c$) [Sch00]. Further, conservatism may be reduced by exploiting block-diagonal structure in $\mathbf{\Delta}$ and constraining the commuting sub-blocks of W ($\tilde{V}$) individually [Sch00].

4.4.2 Numerical Examples

In sight of practical challenges detailed in Section 5.2.3, we adapt the numerical example of Section 4.3.3 to a robust control setting by including an uncertainly damped resonance in the plant model of the PLL's controlled oscillator at 100 kHz. The transfer function of G_{vco} in the case study PLL (3.9) thus modifies to

$$\tilde{G}_{\text{vco}} = \frac{250}{s} \frac{s^2 + 2\omega_r s + \omega_r^2}{s^2 + 2\zeta\omega_r s + \omega_r^2}, \tag{4.62}$$

where $\omega_r = 2\pi \cdot 10^5$ and $\zeta \in [\underline{\zeta}, \bar{\zeta}]$ is the uncertain damping parameter. We only consider the well-posed problem variation based on Example 4.2. Because comparing controller parameters of various robust solutions to nominal ones has little meaning, a different representation of the numerical results is chosen that is better suited to the robustness aspect of the problem. This also alleviates the need to fix the number of decision variables to two. Hence, in the following, we consider a chain of $N = 3$ oscillators with independent uncertain damping parameters ζ_i, $i = 1, \ldots, N$, from a common range of possible realizations. We continue to consider only static feedback for each of the three PLLs, as

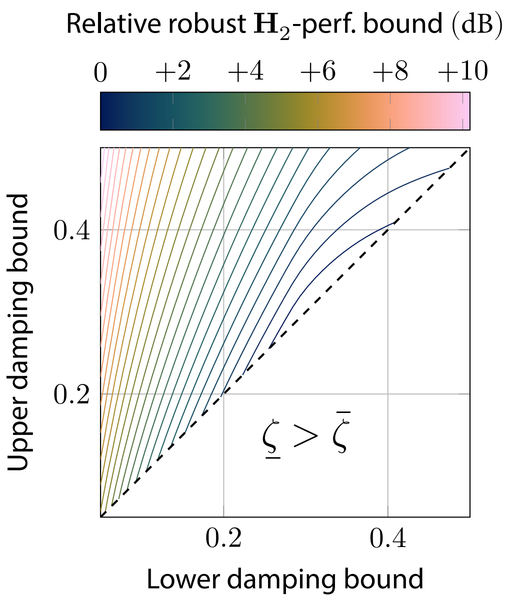

Figure 4.7: Guaranteed robust $\mathbf{H}_2$-performance bound depending on the range of the uncertain damping parameters ζ_i.

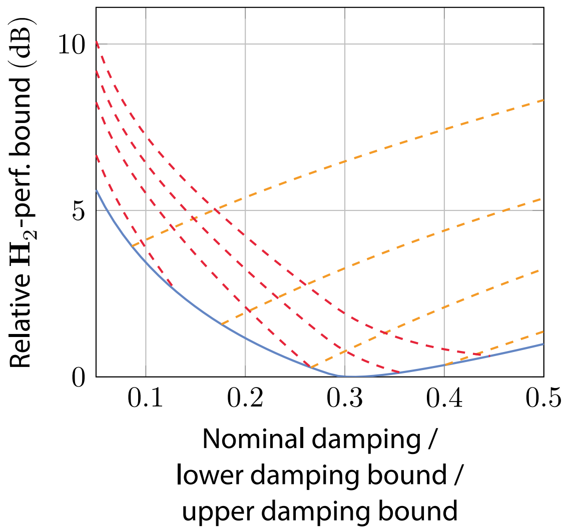

Figure 4.8: Normalized $\mathbf{H}_2$-performance bounds guaranteed by the SI nominal decentralized controller (——), and by the SI robust decentralized controllers for fixed lower bounds (- - -) and upper bounds (- - -) on the damping parameters ζ_i.

additional dynamics have secondary impact on the achieved performance. The final model is thus of 25^{th} order.

Figure 4.7 shows the effect of the additional robustness constraint on the guaranteed worst-case $\mathbf{H}_2$-norm, normalized to the lowest nominal $\mathbf{H}_2$-performance bound achieved when utilizing the SI approach via Proposition 4.32 over the realizable range of ζ given by $\underline{\zeta} = 0.05$, $\bar{\zeta} = 0.5$. The guaranteed performance deteriorates exponentially for decreasing lower bounds $\underline{\zeta}$, which is expected. For increasing upper bounds $\bar{\zeta}$ the guaranteed $\mathbf{H}_2$-norm also increases, however at a much slower rate, and is likely attributed to the increased conservatism due to the numerical challenges of Lyapunov function shaping for larger uncertainty sets. Notably, the minimum of the field does not occur at the extremal value $\underline{\zeta} = \bar{\zeta} = 0.5$, as one might expect, but rather close to $\underline{\zeta} = \bar{\zeta} = 0.3$. This is validated by Figure 4.8, showing the same minimum on the achievable $\mathbf{H}_2$-norm bound in the nominal damping case. Further, Figure 4.8 shows that the conservatism introduced by the additional robustness multiplier variables vanishes for vanishing uncertainty, seen by the convergence of the robust performance profiles (- - -, - - -) and the nominal profile (——).

4.5 Structured Robust Control via Integral Quadratic Constraints in the Frequency Domain

Models of large-scale distributed systems, such as timing distribution networks covering several kilometres in distance, are often affected by uncertainty that goes beyond the class of uncertain time-varying parameters considered primarily so far. In particular, for stabilized timing links, as discussed in Section 3.4.1, one generally wants to consider uncertain time delays, which cannot be covered adequately with the static multipliers and full-block S-procedure discussed in Section 4.4. While, for a controller optimized with the aforementioned approach, robust stability and performance can be certified post-synthesis using IQC analysis (utilizing dynamic multipliers), uncertainty classes characterized by dynamic multipliers (not taken into consideration directly in the synthesis step) lead to potentially significant suboptimality or insufficient robust stability margins. Generalizing multiple previous robust control approaches, there has been a significant effort in the field to utilize IQC theory beyond analysis for robust controller synthesis. Unfortunately, even though nominal, full-order, single- and multi-objective synthesis has been successfully rendered convex by suitable linearizing change of variables or projection, and even robust synthesis has been shown to be convex as long as the control channel is not affected by uncertainty [SV12], the additional decision variables introduced by multipliers for robustness have been shown to create non-resolvable bilinearities in the synthesis inequalities of the general case [Sch00; VS14].

A well-established approach to make robust IQC synthesis tractable relies on iterative procedures known previously from μ-synthesis, alternating between a nominal controller synthesis step and IQC analysis step of the loop closed with the fixed candidate controller. This approach is detailed in [VS14]. While in many cases practical, this approach cannot guarantee global (or even local) optimality of the solution, and the optimality gap is not easily quantifiable. Additionally, the order of the synthesized controller is a result of the algorithm and is usually of order equal or greater than that of the generalized plant. Lastly, the SI framework cannot be utilized in general to obtain a solution with a desired information structure. Ultimately, all these issues result from the core principle of formulating the controller optimization problem as convex SDP using LMIs. Another line of research, based on central works by Apkarian and Noll, attempts to circumvent the limitations of LMI-based methods by employing nonsmooth optimization techniques [AN06b; NA05; GA11] operating directly on the non-convex cost function. Structural and fixed-order constraints on the solution are thereby easily incorporated, and the subgradient descent method always provides at least a local certificate of optimality. Since the IQC framework is a highly capable tool to characterize uncertainty within the system by input-output criteria, very recent research has considered ways to combine non-smooth optimization with IQC-based robustness analysis [CS20; AN06a].

In the following, Section 4.5.1 introduces the required theory to express IQC-based robust stability and performance tests as an equivalent maximum singular value condition in the frequency domain [CS20], which can be solved in standard fashion using nonsmooth optimization [AN06b]. A complementary $\mathbf{H}_2$-performance bound in the frequency domain [Pag99] is discussed in Section 4.5.2, replacing state-space conditions involving Lyapunov variables from the LMI approach. Finally, Section 4.5.3 derives Lemma 4.53, which enables efficient implementation of the aforementioned frequency-domain bound on the $\mathbf{H}_2$-norm, which we then integrate into MATLAB's `systune` for numerical case studies.

4.5.1 Nonsmooth Optimization for Robust Structured Controller Synthesis via IQCs

We again start with the analysis case. As discussed in Section 2.3.3, the central challenge of applying Theorem 2.17 and the subsequent robust stability and performance tests of Corollaries 2.20 and 2.23 is in verifying the FDI (2.28) (respectively FDI (2.34)) for all frequencies ω on the imaginary axis. Instead of using the KYP Lemma 2.22 to express the FDI as an equivalent LMI condition, the work of Cavalcanti and Simões [CS20] aims to avoid the previously discussed challenges of the state-space LMI approach by equivalently expressing the IQC FDI as a maximum singular value test, a problem for which mature and efficient nonsmooth optimization algorithms exist. For this, they employ the following well-known transform, mapping positive real systems into bounded real systems and vice versa, thus linking passivity and small-gain theory.

Definition 4.44 (Bilinear Sector Transformation [DV75]).

$$\mathrm{sect}(X) := (I - X)(I + X)^{-1}. \tag{4.63}$$

The following lemma uses the operator $\mathrm{herm}(X) := X + X^{\mathsf{H}}$, i.e. the double Hermitian part of X.

Lemma 4.45 ([LSA94]). *For $X \in \mathbf{C}^{n\times n}$, $\bar{\sigma}(X) < 1$ if and only if $\mathrm{herm}(\mathrm{sect}(X)) \succ 0$. Vice versa, $\mathrm{herm}(X) \succ 0$ if and only if $\bar{\sigma}(\mathrm{sect}(X)) < 1$.*

In order to apply Lemma 4.45 to the IQC FDI, the left-hand side of (2.28) (respectively (2.34)) needs to be expressible in double Hermitian form. For this, the IQC stability and performance multipliers need to satisfy a special structure, which is introduced next.

Let $\tilde{\mathbf{\Pi}}$ be the set of dynamic IQC multipliers satisfying factorization of the form

$$\tilde{\mathbf{\Pi}} := \left\{ \tilde{\Pi}(s) = S(s)^{\mathsf{H}}\Phi(s)S(s) \mid \Phi(s) \in \mathbf{\Phi}, S(s) \in \mathbf{S} \right\}, \tag{4.64}$$

where

$$\mathbf{\Phi} := \left\{ \Phi(s) = \begin{bmatrix} \Sigma(s)^{\mathsf{H}}\Sigma(s) & \phi(s)^{\mathsf{H}} \\ \phi(s) & \mathrm{herm}(\Psi(s)) \end{bmatrix} \mid \Sigma(s) \in \mathbf{RL}_\infty^{n_q\times n_q}, \phi(s) \in \mathbf{RL}_\infty^{n_p\times n_q}, \Psi(s) \in \mathbf{RL}_\infty^{n_p\times n_p} \right\} \tag{4.65}$$

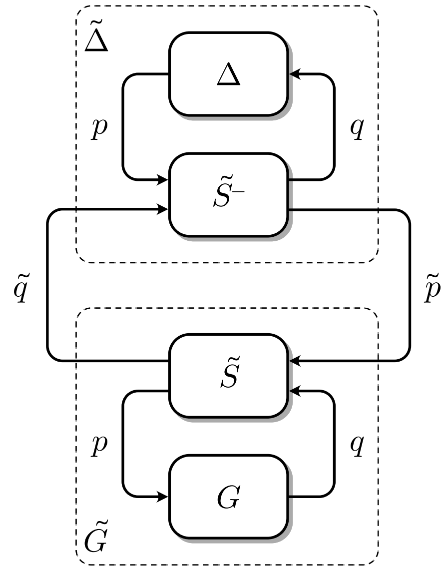

Figure 4.9: Transformed interconnection of an uncertain system.

and $S(s)$ is a transformation matrix from the set

$$\mathbf{S} := \left\{ S(s) = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \mid \right.$$
$$\left. S(s) \in \mathbf{RL}_\infty^{(n_q+n_p)\times(n_q+n_p)}, S_{11} \in \mathbf{RL}_\infty^{n_q\times n_q}, S_{22} \in \mathbf{RL}_\infty^{n_p\times n_p} \right\}. \quad (4.66)$$

The transformation matrix S maximizes the set $\tilde{\mathbf{\Pi}}$ and thus the uncertainty classes for which the approach of [CS20] is applicable.

Remark 4.46. Note that the class of dynamic IQC multipliers (4.64) is sufficiently general to cover most of the frequently encountered uncertainty classes [CS20].

Next, consider the standard uncertain system interconnection G – Δ Figure 2.3a and assume that the IQC defined by some $\tilde{\Pi} \in \tilde{\mathbf{\Pi}}$ is satisfied by Δ. Then observe that there exists the transformed interconnection $\tilde{G}$ – $\tilde{\Delta}$ shown in Figure 4.9 and given by $\tilde{G} := \mathcal{F}_u(G, \tilde{S})$ and $\tilde{\Delta} := \mathcal{F}_l(\Delta, \tilde{S}^{-})$, where

$$\tilde{S} := \begin{bmatrix} S_{12}S_{22}^{-1} & S_{11} - S_{12}S_{22}^{-1}S_{21} \\ S_{22}^{-1} & -S_{22}^{-1}S_{21} \end{bmatrix} \quad (4.67a)$$

and

$$\tilde{S}^{-} := \begin{bmatrix} -S_{11}^{-1}S_{12} & S_{11} \\ S_{22} - S_{21}S_{11}^{-1}S_{12} & S_{21}S_{11}^{-1} \end{bmatrix} \quad (4.67b)$$

are such that a) $\tilde{S}^{-}$ is the inverse of $\tilde{S}$ in the sense of the LFT operation, i.e. $\mathcal{F}_u(\mathcal{F}_u(G, \tilde{S}), \tilde{S}^{-}) = G$ and b) $\tilde{\Delta}$ satisfies the IQC defined by Φ of $\tilde{\Pi} = S^{\mathsf{H}}\Phi S$. This provides the prerequisite for the proof of the following specific case of Theorem 2.17 for dynamic IQC multipliers in $\tilde{\mathbf{\Pi}}$.

Theorem 4.47 ([CS20]). *Let $G \in \mathbf{RH}_\infty^{n_q\times n_p}$ be a stable transfer function and $\Delta\colon \mathbf{L}_{2_e}^{n_q} \to \mathbf{L}_{2_e}^{n_p}$ a causal and bounded operator. The feedback interconnection of G and Δ as in Figure 2.3a is stable if*

1. *for all $\alpha \in [0,1]$ the interconnection is well-posed for $\alpha\Delta$ in place of Δ,*
2. *for all $\alpha \in [0,1]$ the IQC defined by some $\tilde{\Pi} = \tilde{\Pi}^{\mathsf{H}} \in \tilde{\mathbf{\Pi}} \subset \mathbf{RL}_\infty^{(n_q+n_p)\times(n_q+n_p)}$ is satisfied for $\alpha\Delta$ in place of Δ, and*
3. *the maximum singular value condition*

$$\bar{\sigma}\left(\operatorname{sect}\left(-\begin{bmatrix}\phi(i\omega)\tilde{G}(i\omega)+\Psi(i\omega) & 0\\ \Sigma(i\omega)\tilde{G}(i\omega) & -\frac{1}{2}I\end{bmatrix}\right)\right) < 1 \;\; \forall\,\omega \in \mathbf{R} \tag{4.68}$$

is satisfied.

With the modified IQC stability Theorem 4.47, one immediately arrives at corresponding versions of the robust stability and robust performance tests of Corollaries 2.20 and 2.23. For robust performance analysis, consider the augmented multiplier obtained from the stability multiplier $\tilde{\Pi}$ and the performance multiplier $\tilde{\Pi}_p$ as

$$\tilde{\Pi}_a := \operatorname{daug}(\tilde{\Pi},\tilde{\Pi}_p) = \begin{bmatrix}\begin{bmatrix}\Sigma_1 & \\ & \Sigma_2\end{bmatrix}^{\mathsf{H}}\begin{bmatrix}\Sigma_1 & \\ & \Sigma_2\end{bmatrix} & \begin{bmatrix}\phi_1 & \\ & \phi_2\end{bmatrix}^{\mathsf{H}}\\ \begin{bmatrix}\phi_1 & \\ & \phi_2\end{bmatrix} & \begin{bmatrix}\Psi_1 & \\ & \Psi_2\end{bmatrix}+\begin{bmatrix}\Psi_1 & \\ & \Psi_2\end{bmatrix}^{\mathsf{H}}\end{bmatrix}, \tag{4.69}$$

and consider (2.34) in the permuted form

$$\begin{bmatrix}G_{qp}(i\omega) & G_{qw}(i\omega)\\ G_{zp}(i\omega) & G_{zw}(i\omega)\\ \hline I & 0\\ 0 & I\end{bmatrix}^{\mathsf{H}}\underbrace{\operatorname{daug}(\tilde{\Pi},\tilde{\Pi}_p)}_{\tilde{\Pi}_a}\begin{bmatrix}G_{qp}(i\omega) & G_{qw}(i\omega)\\ G_{zp}(i\omega) & G_{zw}(i\omega)\\ \hline I & 0\\ 0 & I\end{bmatrix} \prec 0 \;\; \forall\,\omega \in \mathbf{R}\cup\{\infty\}. \tag{4.70}$$

This again enables the use of the bilinear sector transform to express the IQC robust performance FDI as a maximum singular value constraint involving the components of $\tilde{\Pi}_a$. Equivalently, composite uncertainties $\Delta = \operatorname{diag}(\Delta_1, ..., \Delta_k)$ may be characterized by the composite multiplier $\tilde{\Pi} = \operatorname{daug}(\tilde{\Pi}_1, ..., \tilde{\Pi}_k)$. The transformation matrix S is constructed block-diagonally accordingly. If on the other hand a given uncertainty Δ satisfies the IQCs defined by multiple multipliers $\tilde{\Pi}_1, ..., \tilde{\Pi}_k$, the conic combination of all multipliers can only be utilized in Theorem 4.47 if they afford the same transformation matrix S, i.e. if for $i = 1, ..., k$ we have $\tilde{\Pi}_i = S^{\mathsf{H}}\Phi_i S$ and $\lambda_i \geq 0$ such that

$$\sum_{i=1}^{k}\lambda_i\tilde{\Pi}_i = S^{\mathsf{H}}\left(\sum_{i=1}^{k}\lambda_i\Phi_i\right)S. \tag{4.71}$$

Proceeding to synthesis is straight forward with this method, because, unlike the LMI approach, nonsmooth optimization in the frequency domain does not require the introduction of Lyapunov or other auxiliary matrix variables, nor linearization by change of variables, dualization or projection. We merely need to formulate a suitable objective function and pose hard constraints in a form suitable to the nonsmooth optimization algorithm. The following discussion is formulated in the setting of Problem 4.37 and interconnection Figure 2.2b.

The sought after structured dynamic output feedback control law is given as a fixed mapping $K(s,\kappa) \in \mathbf{RL}_\infty^{n_u\times n_e}$ of to-be-optimized controller parameters $\kappa \in \mathbf{R}^{n_\kappa}$ into the set of admissible controllers, e.g.

those of fixed-order n_{x_K} satisfying a certain information structure or, given a decentralized case with $n_u = n_e$, simply the set of block-diagonal PI controllers

$$K(s,\kappa) = \begin{bmatrix} K_P^{(1)} + K_I^{(1)} s^{-1} & & \\ & \ddots & \\ & & K_P^{(n_u)} + K_I^{(n_u)} s^{-1} \end{bmatrix} \tag{4.72}$$

with $\kappa = \left[K_P^{(1)}, K_I^{(1)}, ..., K_P^{(n_u)}, K_I^{(n_u)}\right]^\top$. Similarly, the augmented multiplier $\tilde{\Pi}_a(s,\chi)$ (4.69) shall be a mapping of the multiplier parameter vector $\chi \in \mathbf{R}^{n_\chi}$ into the family of multipliers $\tilde{\mathbf{\Pi}}_a$ for which the IQCs defined by $\tilde{\Pi}$ are satisfied by the uncertainty Δ and performance on the channel $w \to z$ with respect to $\tilde{\Pi}_p$ is guaranteed as required by Corollary 2.23.

We now state two variants of the synthesis problem based on Theorem 4.47. Both require constraints to be formulated in the form of functions

$$u(x) = \max_{y \in Y} v(x,y), \tag{4.73}$$

where $v\colon \mathbf{R}^n \times \mathbf{R}^m \to \mathbf{R}$ is continuous, $\nabla_x v(\cdot,\cdot)$ exists and is continuous, and $Y \in \mathbf{R}^m$ is compact [CS20].

We first restate the suboptimal formulation as it is stated in [CS20], i.e. with some fixed hard bound on the performance index characterized by the performance multiplier Π_p, and minimizing the maximum singular value in order to maximize confidence in the feasibility of the solution (provided it indeed satisfies (4.68)).

$$\begin{aligned}
\underset{\chi}{\text{minimize}} \quad & \max_\omega \bar{\sigma}\left(\text{sect}\left(-\begin{bmatrix} \phi(i\omega,\chi)\tilde{M} + \Psi(i\omega,\chi) & 0 \\ \Sigma(i\omega,\chi)\tilde{M} & -\frac{1}{2}I \end{bmatrix}\right)\right) && (4.74\text{a}) \\
\text{subject to} \quad & \mathcal{F}_l(P, K(\kappa)) \text{ stable}, \\
& h_i(\chi) \le 0, \quad i = 1, ..., n_h, && (4.74\text{b}) \\
& l_j(\kappa) \le 0, \quad j = 1, ..., n_l, && (4.74\text{c})
\end{aligned}$$

where $\tilde{M} = \mathcal{F}_u(\mathcal{F}_l(P, K(\kappa)), \tilde{S}(i\omega,\chi))$ and $h_i(\cdot), l_j(\cdot)$ are constraints of the kind of (4.73).

Optimization of the performance index via the optimization program (4.74) requires an outer bisection algorithm, decreasing a scalar bound γ on the performance index until stability can no longer be guaranteed by (4.68). This is the way it is done in [CS20]. Alternatively, one can always express the bound as a function of the multiplier parameter vector χ, $\gamma(\chi)$, and minimize directly. This leads to the alternate formulation

$$\begin{aligned}
\underset{\chi}{\text{minimize}} \quad & \gamma(\chi) && (4.75\text{a}) \\
\text{subject to} \quad & \mathcal{F}_l(P, K(\kappa)) \text{ stable}, \\
& \max_\omega \bar{\sigma}\left(\text{sect}\left(-\begin{bmatrix} \phi(i\omega,\chi)\tilde{M} + \Psi(i\omega,\chi) & 0 \\ \Sigma(i\omega,\chi)\tilde{M} & -\frac{1}{2}I \end{bmatrix}\right)\right) < 1 && (4.75\text{b}) \\
& h_i(\chi) \le 0, \quad i = 1, ..., n_h, && (4.75\text{c}) \\
& l_j(\kappa) \le 0, \quad j = 1, ..., n_l. && (4.75\text{d})
\end{aligned}$$

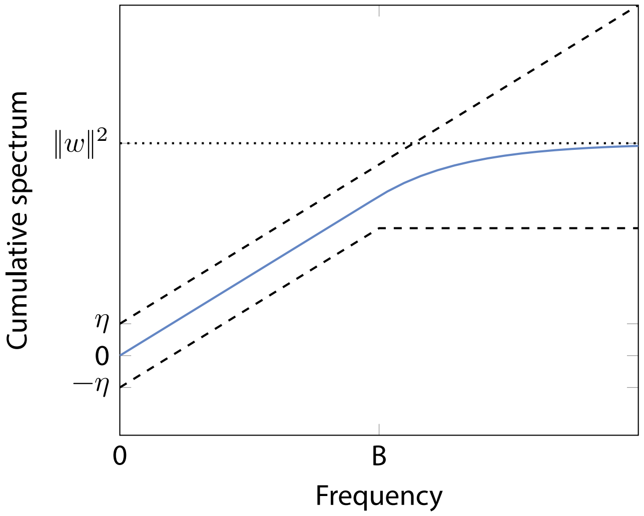

Figure 4.10: Cumulative spectrum of a signal with white spectrum up to accuracy η and bandwidth B, tapering off thereafter [Pag99].

Optimization problems of the kind (4.74) and (4.75) are well suited to the existing nonsmooth optimization algorithm method implemented in the MATLAB Control Toolbox [The24b], systune. For details on the algorithm and its implementation, we refer to [CS20; AN06b; GA11]. While the form (4.74) in conjunction with bisection optimization is readily supported by systune, the form (4.75) requires adaptation of systune in the form of additional "tuning goals" implementing the objective function $\gamma(\chi)$. We will return to this issue in Section 4.5.3.

4.5.2 Frequency-Domain Conditions for Robust H_2-Performance

Owing to the formulation of optimization program (4.75) in the frequency domain, the IQC robust $\mathbf{H}_2$-performance test of Lemma 2.25 is not well suited to this case, given that it again requires the computation of a Lyapunov matrix, bringing many downsides of the state-space approach back to the table. We therefore intend to provide with this section a corresponding $\mathbf{H}_2$-performance test via multipliers entirely in the frequency domain. The derivation of said multiplier along the lines of [Pag99; PF97; SEW23] however warrants a prior discussion on the formal generalization of robust $\mathbf{H}_2$-performance to nonlinear systems, because, unlike the robust control method via full-block multipliers discussed in Section 4.4, the IQC framework extends to uncertainties with fully non-linear dynamics. As we will see next, this has the consequence that the two common interpretations of the $\mathbf{H}_2$-norm in terms of (deterministic) impulse response energy and (stochastic) white-noise rejection are no longer equivalent as in the LTI (or LTV) case.

The first coherent account on generalization of the worst-case $\mathbf{H}_2$-norm under nonlinear, time-varying uncertainty in the sense of $\sup_{\Delta\in\mathbf{\Delta}}\|\mathcal{F}_u(P,\Delta)\|_2$ is given in [Sto93]. Clearly, the $\mathbf{H}_2$-norm of a strictly-proper, stable LTI system $G(s) \in \mathbf{H}_2$,

$$\|G\|_2^2 := \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(G(i\omega)^{\mathsf{H}}G(i\omega))\ \mathrm{d}\omega,$$

as defined in (2.20) of Definition 2.12, does not apply to nonlinear systems. The signal-based interpretations (2.22) and (2.23), based on the average (or expected) impulse response energy over all inputs (of random direction) and the expected output power when driven by unitary white-noise at the input, respectively, however do permit this extension, if only in semi-norms. The former is given by

$$\|T\|_{2,\mathrm{ir}}^2 := \sum_{i=1}^{n_w} \|Te_i\delta(t)\|_2^2 \tag{4.76}$$

and has already been bounded in Lemma 2.25. Here $\{e_1, ..., e_{n_w}\}$ is an orthonormal basis spanning the input space and $\delta(t)$ is the Dirac delta function. This generalization is lacking in the sense that it depends on the choice of basis and the fact that the impulse is assumed at time $t = 0$ [Sto93]. A mathematically rigorous generalization of the stochastic white-noise rejection interpretation in the frequency domain is given in [Pag99]. When generalizing the stochastic interpretation in continuous-time to a frequency-domain condition, one immediately encounters mathematical difficulties founded in the infinite bandwidth of continuous-time white noise. Paganini tackles this challenge by reversing the limiting process, first analysing the response of the system to signals with white spectrum up to accuracy η and bandwidth B and *then* letting those variables tend to zero and infinity, respectively. For this, we formally consider the set of deterministic signals

$$\mathbf{W}_{\eta,B}^m := \left\{ w \in \mathbf{L}_2 \mid \int_{-\beta}^{\beta} \hat{w}(i\omega)\hat{w}(i\omega)^{\mathsf{H}} \frac{\mathrm{d}\omega}{2\pi} \in \mathbf{G}_{\eta,B}^m \right\}, \tag{4.77}$$

where, using $\|X\|_{\max} := \max_{i,j}|X_{ij}|$,

$$\mathbf{G}_{\eta,B} := \left\{ g \in \mathcal{C}_{\mathsf{H}}^{m\times m}(\mathbf{R}_+) \mid \sup_{0\le\beta\le B} \left\| g(\beta) - \frac{\beta}{\pi} I \right\|_{\max} \le \eta, \right.$$
$$\left. \operatorname{trace}(g(\beta)) \le m\left(\frac{\beta}{\pi} + \eta\right) \quad \forall \beta \ge B \right\}, \tag{4.78}$$

that is, $\mathbf{W}_{\eta,B}$ is the set of all vector-valued, square-integrable signals with cumulative spectrum bounded as depicted in Figure 4.10. Then, along the lines of the Bartlett test, the following semi-norm is defined:

$$\|T\|_{\mathbf{W}_{\eta,B}} := \sup_{w\in\mathbf{W}_{\eta,B}} \|Tw\|. \tag{4.79}$$

We abbreviate the limit case as

$$\|T\|_{2,\mathrm{wn}} := \lim_{\substack{\eta\to 0\\ B\to\infty}} \|T\|_{\mathbf{W}_{\eta,B}}. \tag{4.80}$$

In order to motivate the new semi-norm $\|\cdot\|_{2,\mathrm{wn}}$ as a valid substitution for the $\mathbf{H}_2$-norm for nonlinear, time-varying systems, we cite the following result on the wide class of transfer functions with frequency response of bounded variation and bounded by a monotonically decreasing function.

Corollary 4.48 ([Pag99]). *For any $T \in \mathbf{RH}_\infty$, it holds that $\|T\|_2 \le \|T\|_{\mathbf{W}_{\eta,B}}$. If in addition $|T(i\omega)|^2 \in \mathrm{BV}_\downarrow(\mathbf{R})$, then $\|T\|_{2,\mathrm{wn}} = \|T\|_2$.*

Here, $\mathrm{BV}_\downarrow(\mathbf{R})$ is formally defined as

$$\mathrm{BV}_\downarrow{}^m(\mathbf{R}) := \Big\{ f \in \mathrm{BV}^{m\times m}(\mathbf{R}) \mid \exists\, g \in \mathbf{L}_1(\mathbf{R}_+),\, g \text{ monotonically decreasing and}$$
$$0 \le f(\omega) = f(-\omega) \le g(|\omega|)I \Big\}. \tag{4.81}$$

Having established equivalence of the $\|\cdot\|_2$ and $\|\cdot\|_{2,\mathrm{wn}}$ norms in the LTI case, an interesting relation to the induced $\mathbf{L}_2$-gain norm can be given for the nonlinear, time-varying case.

Lemma 4.49 ([Pag99]). *Given a nonlinear, time-varying system $P(\Delta)$ with LFR as in (2.12), the following inequality holds:*

$$\|\mathcal{F}_u(P,\Delta)\|_{\mathbf{W}_{\eta,B}} \leq \|P_{zp}\Delta(I-P_{qp}\Delta)^{-1}I\|_\infty \cdot \left\| \begin{bmatrix} P_{zp} \\ P_{zw} \end{bmatrix} \right\|_{\mathbf{W}_{\eta,B}} < \infty. \tag{4.82}$$

Remark 4.50. It is important to note that even for LTI uncertainty, Corollary 4.48 does not guarantee strict equality in the robust performance analysis case, i.e. in general

$$\sup_{\Delta\in\mathbf{\Delta}^{\mathrm{LTV}}} \|\mathcal{F}_u(P,\Delta)\|_2 \leq \sup_{\Delta\in\mathbf{\Delta}^{\mathrm{LTV}}} \|\mathcal{F}_u(P,\Delta)\|_{2,\mathrm{wn}}. \tag{4.83}$$

In particular, for multiple input, multiple output (MIMO) systems, a corresponding example showing conservatism of a factor $\sqrt{n_w}$ is provided in [ST00]. It is noted in [PF00] that this limitation also exists for the nonlinear extension of the deterministic impulse response interpretation, and is explained by the worst-case analysis summing over the worst-case responses in each output, which are attained by different realizations of the uncertainty, which in reality do not occur at the same time.

We now aim to characterize $\|\cdot\|_{2,\mathrm{wn}}$ as a performance index with a suitable multiplier and the performance IQC (2.32). We fundamentally rely on the multiplier proposed in [Pag99], but slightly adapt its definition, in order to not have to rely on frequency gridding to verify the robustness test. An analogous result is given in [JM99].

Corollary 4.51 ([SEW23]). *Given a (potentially nonlinear) system $T\colon \mathbf{L}_{2_e}^{n_w} \to \mathbf{L}_{2_e}^{n_z}$ and $Y = Y^{\mathsf{H}} \in \mathbf{RL}_\infty^{n_w\times n_w}$ of finite order satisfying*

$$\mathcal{I}\left(\Pi_p, Tw, w\right) \leq -\varepsilon\|w\|^2$$

with

$$\Pi_p = \begin{bmatrix} I & 0 \\ 0 & -Y \end{bmatrix}, \tag{4.84}$$

then

$$\|T\|_{2,\mathrm{wn}}^2 \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega))\ \mathrm{d}\omega. \tag{4.85}$$

Proof. Since $Y = Y^{\mathsf{H}}$ and of finite order, we have $Y(i\omega) \in \mathrm{BV}^{n_w}(\mathbf{R})$. Additionally, because $Y \in \mathbf{RL}_\infty$, it has no poles on the imaginary axis and there exists a monotonically decreasing function $g \in \mathbf{L}_1(\mathbf{R}_+)$ bounding $Y(i\omega)$, $\omega\in\mathbf{R}$, such that $Y(i\omega) \in \mathrm{BV}_{\downarrow}{}^{n_w}(\mathbf{R})$. Letting $z = Tw$, we have, by assumption,

$$\begin{aligned}
&\int_{-\infty}^{\infty} \begin{bmatrix} \hat{z}(i\omega) \\ \hat{w}(i\omega) \end{bmatrix}^{\mathsf{H}} \begin{bmatrix} I & 0 \\ 0 & -Y(i\omega) \end{bmatrix} \begin{bmatrix} \hat{z}(i\omega) \\ \hat{w}(i\omega) \end{bmatrix}\ \mathrm{d}\omega \\
&\qquad = \int_{-\infty}^{\infty} \hat{z}(i\omega)^{\mathsf{H}}\hat{z}(i\omega)\ \mathrm{d}\omega - \int_{-\infty}^{\infty} \hat{w}(i\omega)^{\mathsf{H}} Y(i\omega)\hat{w}(i\omega)\ \mathrm{d}\omega \\
&\qquad = 2\pi\|z\|^2 - \int_{-\infty}^{\infty} \hat{w}(i\omega)^{\mathsf{H}} Y(i\omega)\hat{w}(i\omega)\ \mathrm{d}\omega \\
&\qquad \leq -\varepsilon\|w\|^2,
\end{aligned} \tag{4.86}$$

which is equivalent to

$$\|z\|^2 \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{w}(i\omega)^{\mathsf{H}} Y(i\omega)\hat{w}(i\omega)\ \mathrm{d}\omega - \varepsilon\|w\|^2. \tag{4.87}$$

For the right-hand side, we utilize the following bound for $w \in \mathbf{W}^{n_w}_{\eta,B}$ derived in [Pag99]:

$$\begin{aligned}\sup_{w\in\mathbf{W}_{\eta,B}} \|z\|^2 = \|T\|_{\mathbf{W}_{\eta,B}} \leq \sup_{w\in\mathbf{W}_{\eta,B}} \frac{1}{2\pi}\int_{-\infty}^{\infty} & \hat{w}(i\omega)^{\mathsf{H}} Y(i\omega)\hat{w}(i\omega)\ \mathrm{d}\omega \\ \leq & \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega))\ \mathrm{d}\omega \\ & + \eta\left(2n_w \cdot g(B) + \sum_{k,l=1}^{n_w} (Y_{kl}(iB) + \mathrm{TV}(Y_{kl}))\right) \\ & + \frac{m}{\pi}\int_{B}^{\infty} g(\omega)\ \mathrm{d}\omega.\end{aligned} \tag{4.88}$$

In the limit, we immediately obtain

$$\lim_{\substack{\eta\to 0\\ B\to\infty}} \|T\|_{\mathbf{W}_{\eta,B}} = \|T\|_{2,\mathrm{wn}} \leq \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega))\ \mathrm{d}\omega, \tag{4.89}$$

which concludes the proof. □

To conclude this section, we summarize insight on the conservatism of this bound and the relation of $\|\cdot\|_{2,\mathrm{ir}}$ and $\|\cdot\|_{2,\mathrm{wn}}$ in addition to Remark 4.50.

Lemma 4.52 ([PF97]). *Consider an uncertain system $T(\Delta) = \mathcal{F}_u(T, \Delta)$ with $T \in \mathbf{RH}_\infty^{\bullet\times\bullet}$ as in (2.12) and $\Delta \in \mathbf{\Delta}$, the set of causal, bounded, nonlinear, time-varying operators on $\mathbf{L}_{2_e}$. Then the following relation holds*

$$\sup_{\Delta\in\mathbf{\Delta}} \|T(\Delta)\|_{2,\mathrm{ir}}^2 \leq \lim_{\substack{\eta\to 0\\ B\to\infty}} \sup_{\Delta\in\mathbf{\Delta}} \|T(\Delta)\|_{\mathbf{W}_{\eta,B}}^2. \tag{4.90}$$

The above Lemma 4.52 is the consequence of causality considerations in the two semi-norms [PF97]. The impulse response bound weakly exploits causality of Δ, which the white-noise rejection approach cannot, thus leading to consideration of (unrealistic) signals worse than the impulse. We however also need to acknowledge computational aspects. The robust performance tests Lemma 2.25 and Corollary 4.51 need to be rendered tractable with finite-dimensional approximations of stability and performance multipliers, in which case possibly

$$\mathrm{trace}(B_w^\top X B_w) \nleq \frac{1}{2\pi}\int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega))\ \mathrm{d}\omega. \tag{4.91}$$

In the LTI uncertainty case, it was indeed shown in [PF97] that substituting $Y(i\omega)$ with an arbitrary function evaluated over a tight frequency grid could essentially close the gap to the classical $\mathbf{H}_2$-norm, which was not achieved with LMI-based multiplier optimization. Neither approach can therefore be seen as inherently superior and the choice generally comes down to which is better suited to the method of optimization.

4.5.3 Efficient Implementation of a Robust H_2-Performance Bound

Corollary 4.51 provides a useful robust $\mathbf{H}_2$-performance test in the form of IQC multiplier wrapping a Hermitian transfer matrix integrable on the imaginary axis, whose frequency response is bounding the worst-case system response to white exogenous signals. For reasons discussed in the introduction to Section 4.5.1, we wish to solve the IQC synthesis problem involving this robust $\mathbf{H}_2$-performance bound using nonsmooth optimization methods, i.e. posing the problem as in (4.74) or (4.75). In either case, a finite-order parameterization is required for the transfer matrix $Y(s) = Y(s)^{\mathsf{H}}$ of the IQC $\mathbf{H}_2$-performance multiplier (4.84). A versatile parameterization for such Hermitian-valued transfer matrices is provided in [JM99]:

$$Y(s,\chi) = \Psi_Y(s,\chi) + \Psi_Y(s,\chi)^{\mathsf{H}} \tag{4.92}$$

with

$$\Psi_Y(s,\chi) = \sum_{i=1}^{N} \frac{X_i s + X_i a_i + Z_i b_i}{s^2 + 2a_i s + a_i^2 + b_i^2}. \tag{4.93}$$

Here, for $i = 1, \ldots, N$, $X_i, Z_i \in \mathbf{R}^{n_w \times n_w}$, and $a_i > 0$ and $b_i \in \mathbf{R}$ are scalars that determine the (negative) real and (absolute) imaginary parts of the location of each summand's pole in the complex plane, respectively. Note that for $b_i = 0$, the expression reduced to a single real pole at $-a_i$,

$$\frac{X_i s + X_i a_i}{s^2 + 2a_i s + a_i^2} = \frac{X_i(s+a)}{(s+a)^2} = \frac{X_i}{s+a}, \tag{4.94}$$

while $b \neq 0$ produces a complex-conjugate pole pair. As such, as long as all pairs a_i, b_i are distinct, $\Psi_Y(s,\chi)$ is the most general way to parameterize a transfer matrix in $\mathbf{RH}_2$. The parameters X_i, Z_i, a_i, b_i are collected in the parameter vector χ, for example as

$$\chi = [\operatorname{vec} X_1, \ldots, \operatorname{vec} X_N, \operatorname{vec} Z_1, \ldots, \operatorname{vec} Z_N, a_1, \ldots, a_N, b_1, \ldots, b_N]^\top. \tag{4.95}$$

The parameters are optimized under a nonsmooth optimization algorithm according to the hard and soft constraints in (4.75). This requires iterative computation of the current constraint and objective function values and their subgradients as the parameters are updated. With objective function (4.85), this involves the evaluation of an integral over the entire imaginary axis. To avoid computationally expensive numerical approximations, we provide here an elegant analytical solution, assuming the parameterization (4.92) is used.

Lemma 4.53 ([SEW23]). *Consider $Y(s,\chi) = Y(s,\chi)^\top \in \mathbf{RH}_\infty^{\bullet\times\bullet}$, parameterized as in (4.92). Then*

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} \operatorname{trace}(Y(i\omega,\chi))\, \mathrm{d}\omega = \sum_{i=1}^{N} \operatorname{trace}(X_i). \tag{4.96}$$

Proof. Following arguments in [JM99], we find for each summand in (4.93) that, initially for $b_i = 0$,

$$\begin{aligned}
&\lim_{B\to\infty} \int_{-B}^{B} \frac{X_i i\omega + X_i a_i + Z_i b_i}{-\omega^2 + 2a_i i\omega + a_i^2 + b_i^2}\, \mathrm{d}\omega \\
&= \lim_{B\to\infty} \int_{-B}^{B} \frac{X_i}{i\omega + a_i}\, \mathrm{d}\omega \\
&= \lim_{B\to\infty} \int_{-B}^{B} 2X_i \arctan\left(\frac{B}{a_i}\right)\, \mathrm{d}\omega \\
&=\pi X_i.
\end{aligned} \tag{4.97}$$

It then follows from analytic continuation of the function $\int_{-B}^{B} \frac{1}{i\omega+z}$ in the region $\mathrm{Re}(z) > 0$ that (4.97) holds for all $a_i > 0, b_i \in \mathbf{R}$. Hence,

$$\begin{aligned}
&\frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega, \chi))\ \mathrm{d}\omega \\
=&\frac{1}{\pi} \lim_{B\to\infty} \int_{-B}^{B} \mathrm{trace}(\Psi_Y(i\omega, \chi))\ \mathrm{d}\omega \\
=&\sum_{i=1}^{N} \mathrm{trace}(X_i).
\end{aligned} \tag{4.98}$$

□

An immediate consequence of Lemma 4.53 is the simple derivation of the gradient of the objective function (4.85) with respect to χ:

$$\nabla_\chi \left[\frac{1}{2\pi} \int_{-\infty}^{\infty} \mathrm{trace}(Y(i\omega, \chi))\ \mathrm{d}\omega\right] = \left[I_X(\chi_j)\right], \tag{4.99}$$

with $I_X(\chi_j) = 1$ if that entry of the parameter vector corresponds to a diagonal entry of an X_i, and zero otherwise.

On another practical note, observe that the robust-performance IQC FDI (2.34) in combination with the robust $\mathbf{H}_2$-performance multiplier (4.84) from Corollary 4.51 leads to

$$0 \prec G_{zw}(i\omega)^{\mathsf{H}} G_{zw}(i\omega) \prec Y(i\omega) \ \ \forall\, \omega \in \mathbf{R} \cup \{\infty\}, \tag{4.100}$$

which, with parameterization (4.92), implies strict positive realness of $\Psi_Y(s, \chi)$ (cf. Definition 2.3). As shown in [SEW23], the nonsmooth optimizer may fail to assert this implicit constraint for large ω, where, due to $Y(s, \chi)$ necessarily strictly proper, $Y(i\omega, \chi)$ converges to zero, leading to marginal infeasibility of the solution. In additional tests, adding an explicit hard constraint for $\Psi_Y(s, \chi)$ to be passive has shown to alleviate the problem, although at the cost of increased optimization time. In the scalar case, a computationally less expensive way to ensure that $\Psi_Y(s, \chi)$ is positive real is to constrain, rather than the entirety of $\Psi_Y(s, \chi)$, each summand of (4.93) to be positive real, realizing that clearly the sum of positive real functions again must be positive real.

Lemma 4.54. $\Psi_Y(s, \chi) \in \mathbf{RH}_\infty^{1\times 1}$ *as in (4.93) is strictly positive real if, for* $i = 1, \ldots, N$,

1. $a_i > 0$,
2. $b_i > 0$,
3. $X_i > 0$, *and*
4. $|Z_i| < X\frac{a}{b}$.

Proof. Consider a single summand of (4.93) in the scalar case, i.e. $X \in \mathbf{R}, Z \in \mathbf{R} \Rightarrow \Psi_Y(s, \chi) \in \mathbf{RH}_2^{1\times 1}$,

$$\frac{Xs + Xa + Zb}{s^2 + 2as + a^2 + b^2} =: \frac{p(s)}{q(s)}, \tag{4.101}$$

where, by definition of $\psi_Y(s,\chi)$, $a > 0$. Clearly, a transfer function is positive real if and only if its Nyquist plot is fully contained in the closed right-half plane, which is only possible if its relative degree is smaller or equal to one. The necessary and sufficient condition may be formulated in terms of the phase never exceeding ±90°. For $p(s)/q(s)$, which is of relative degree one and order one or two, the upper bound is implicitly assured (since $a > 0$), while the lower bound depends on the relative position of the zero and the complex conjugate pole pair. Factoring out X and letting $z = a + \frac{Zb}{X}$ and $c = a^2 + b^2$, such that

$$\frac{p(s)}{q(s)} = X\frac{s+z}{s^2+2as+c},$$

we infer $X > 0$, as otherwise, since $c > 0$, we have for $z > 0$, $\psi_Y(0,\chi) < 0$ or, for $z < 0$, that $\lim_{\omega\to\infty} \mathrm{Re}(\psi_Y(i\omega,\chi)) \to 0^-$. But then also necessarily $z > 0$. Furthermore,

$$\begin{aligned}
\angle\frac{p(i\omega)}{q(i\omega)} &= \angle p(i\omega) - \angle q(i\omega) \\
&= \tan^{-1}\left(\frac{\mathrm{Im}\{p(i\omega)\}}{\mathrm{Re}\{p(i\omega)\}}\right) - \tan^{-1}\left(\frac{\mathrm{Im}\{q(i\omega)\}}{\mathrm{Re}\{q(i\omega)\}}\right) \\
&= \tan^{-1}\left(\frac{\omega}{z}\right) - \tan^{-1}\left(\frac{2a\omega}{-\omega^2+c}\right) \\
&= \tan^{-1}\left(\frac{\frac{\omega}{z} - \frac{2a\omega}{-\omega^2+c}}{1+\frac{2a\omega^2}{-z\omega^2+zc}}\right) \\
&= \tan^{-1}\left(\frac{-\omega^3+\omega(c-2az)}{\omega^2(2a-z)+zc}\right).
\end{aligned} \tag{4.102}$$

Examining the denominator of the last argument, we find that −90° is not achieved for finite ω as long $2a > z$, i.e.

$$2a > z = a + \frac{Zb}{X} > 0 \quad \Leftrightarrow \quad Xa > Zb > -Xa. \tag{4.103}$$

Arbitrarily and without loss of generality fixing $b > 0$, we have

$$|Z| < X\frac{a}{b}, \tag{4.104}$$

which concludes the proof. □

Finally, we point out that a useful state-space form of $Y(s,\chi)$ (4.92) is easily derived from the modified Jordan canonical state-space form of a single summand of Ψ_Y given by

$$\frac{X_i s + X_i a + Z_i b_i}{s^2 + 2a_i s + a_i^2 + b_i^2} = \left[\begin{array}{c|c} I_{n_w} \otimes \begin{bmatrix} -a & b \\ -b & -a \end{bmatrix} & I_{n_w} \otimes \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ \hline \left[Z_i^{(1)}, X_i^{(1)}, \ldots, Z_i^{(n_w)}, X_i^{(n_w)}\right] & 0 \end{array}\right]. \tag{4.105}$$

4.5.4 Numerical Examples

In the following numerical example, we qualitatively and quantitatively investigate the nonsmooth optimization approach to the IQC framework in terms of computational efficiency, limiting aspects of using finite-order approximations for the IQC multipliers, and the benefits and possible downsides of the local certificate of optimality provided by the algorithm, particularly in comparison to the SI robust synthesis methods discussed in Section 4.4.

To illustrate how to use the method in practice, first consider the robust performance analysis of the following system as proposed in [PF97] and also discussed in [Sch+22]:

$$P(\delta) = \frac{1}{s^2 + (2+\delta)s + 1}. \tag{4.106}$$

Here, the uncertain damping $\delta \in [-1, 1]$ leads to an analytical worst-case $\mathbf{H}_2$-norm of

$$\sup_{\delta \in [-1,1]} \|P(\delta)\|_2 = 1/\sqrt{2}.$$

The corresponding uncertain LFR is easily constructed as $P(\delta) = \mathcal{F}_u(P, \delta)$ with

$$P = \left[\begin{array}{cc|cc} -2 & -1 & -1 & 1 \\ 1 & 0 & 0 & 0 \\ \hline 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{array}\right]. \tag{4.107}$$

Since the damping is unknown, but constant, the most general multiplier ensuring minimal conservatism is dynamic and of the form

$$\Pi(s) = \begin{bmatrix} D(s)^{\mathsf{H}} D(s) & W(s) - W(s)^{\mathsf{H}} \\ W(s)^{\mathsf{H}} - W(s) & -D(s)^{\mathsf{H}} D(s) \end{bmatrix} \tag{4.108}$$

with $D(s), W(s) \in \mathbf{RL}_\infty$. Once again, finite-order approximations for $D(s)$ and $W(s)$ are needed to make the problem amenable to numerical methods. We select the parameterization (4.93) with realization like (4.105). Selecting

$$\Pi_p = \begin{bmatrix} I & 0 \\ 0 & -Y(s) \end{bmatrix} \tag{4.109}$$

from Corollary 4.51 again with parameterization (4.92), we yield the augmented multiplier

$$\Pi_a = \operatorname{daug}(\Pi, \Pi_p) = \begin{bmatrix} \Sigma(s)^{\mathsf{H}} \Sigma(s) & \phi(s)^{\mathsf{H}} \\ \phi(s) & \operatorname{herm}(\Psi(s)) \end{bmatrix}, \tag{4.110}$$

with

$$\begin{aligned} \Sigma(s) &= \begin{bmatrix} D(s,\chi) & 0 \\ 0 & I \end{bmatrix}, \quad \phi(s) = \begin{bmatrix} W(s,\chi)^{\mathsf{H}} - W(s,\chi) & 0 \\ 0 & 0 \end{bmatrix}, \\ \Psi(s) &= \begin{bmatrix} -\frac{1}{2}(D(s,\chi)^{\mathsf{H}} D(s,\chi)) & 0 \\ 0 & -\Psi_Y(s,\chi) \end{bmatrix}, \end{aligned} \tag{4.111}$$

and naturally $Y(s,\chi) = \operatorname{herm}(\Psi_Y(s,\chi))$. Due to the Hermitian nature of the multipliers, no additional hard constraints are necessary. Nevertheless, we constrain the parameter vector χ such that the a, b, and X parameters in $\Psi_Y(s,\chi)$ are in fact non-negative using appropriate constraint functions $h_i(\chi) \lesssim 0$. Technically, the a parameters should be constrained to be non-zero, but this is omitted in this numerical setting and verified a-posteriori. Finally, we restrict $\Psi_Y(s,\chi)$ to be positive real, rendering the implicit condition explicit.

Running the nonsmooth optimization program for 100 randomly initialized multipliers, all runs converged to multipliers guaranteeing robust stability, yet at various local minima in terms of

$$\sup_{\delta \in [-1,1]} \|P(\delta)\|_{2,\mathrm{wn}}.$$

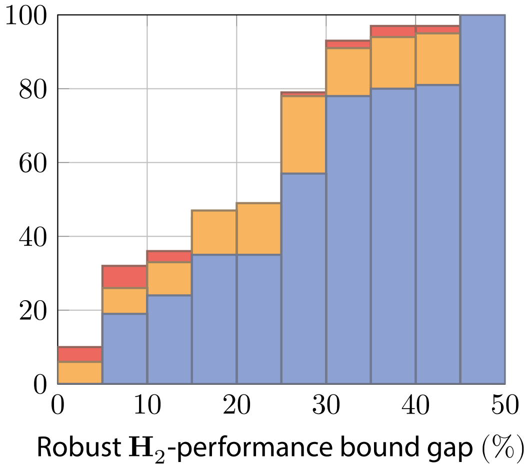

Figure 4.11: Relative robust performance gap guaranteed for $P(\delta)$ by a given number of out 100 runs. $D(s,\chi)$ and $W(s,\chi)$ were selected as sixth (——), fourth (——), and second order (——), respectively.

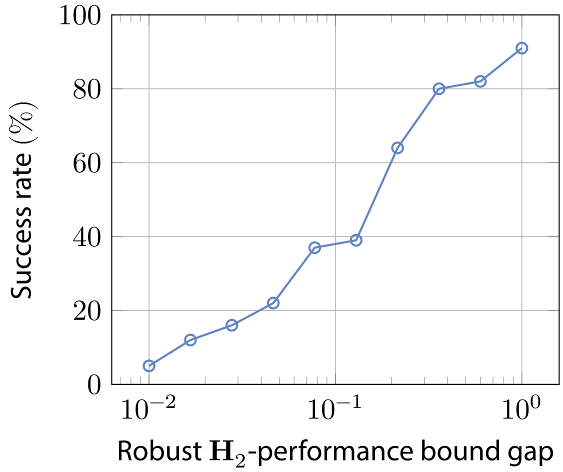

Figure 4.12: Percentage of 100 randomly initialized nonsmooth optimization runs successfully guaranteeing a given worst-case $\mathbf{H}_2$-performance bound for $P(\delta)$.

The distribution of results is displayed in Figure 4.11 for $D(s,\chi)$, $W(s,\chi)$ of sixth, fourth and second order, respectively. $\Psi_Y(s,\chi)$ was kept at second order throughout, as this is sufficient to bound the worst-case frequency response of the uncertain plant exactly, as will be derived shortly. The best result for sixth-order $D(s,\chi)$ and $W(s,\chi)$ virtually achieves the analytical worst-case norm of $1/\sqrt{2}$, being just 0.3 % off. While all runs had a performance gap of less than 50 %, almost half had a gap of less than 25 %. It is important to keep in mind that the probability of achieving a certain gap very much depends on the polytope from which the initial parameterization is sampled. This is shown by the fact that only very few runs achieved a gap less than 5 %, showing that the global minimum and local minima close to it have a small region of attraction. Designing the parameter polytope based on prior knowledge of the uncertainty and the resulting expected order of magnitude of the multiplier parameters will generally improve the expected gap.

For this simple example, the family of $\mathbf{H}_2$-performance multipliers that tightly bound the worst-case frequency response of $P(\Delta)^{\mathsf{H}}P(\Delta)$ up to a numerical gap proportional to ε is given by

$$\Psi_Y(s) = \left(\frac{1}{2}+\varepsilon\right)\frac{s+\frac{1}{2}+\frac{1}{\sqrt{3}}}{s^2+s+1}. \tag{4.112}$$

Fixing the performance multiplier in this fashion, the probability that optimization of random initial $D(s,\chi)$ and $W(s,\chi)$ converges to a feasible solution is plotted against the robust performance gap (controlled by ε) in Figure 4.12. A solution is considered feasible, if the IQC maximum singular value condition is not violated by more than a numerical tolerance of $1\cdot 10^{-5}$. Note that, due to the $\mathbf{H}_2$ setting, both $P(\delta)$ and Y are necessarily strictly proper and hence converge to 0 for $\omega\to\infty$. Consequently, the maximum singular value has a theoretical infimum of 1 when considering the entire imaginary axis. This causes slowing solver progress as the desired threshold of 1 is approached and causes the solver to terminate at possibly marginally infeasible solutions. It can be argued however, that the robust guarantee is not significantly weakened by this marginal infeasibility, just as, analogously, the feasibility of the LMI-based solution is also only guaranteed up to a numerical tolerance. It may therefore be numerically beneficial to limit the frequency range over which the singular value condition is enforced. In this example the range was limited to $100\,\mathrm{rad}$. However, it is difficult to quantitatively assess the impact of this hyperparameter. A more detailed discussion is hence left out.

Chapter 5

Optical Timing Synchronization at European XFEL and FLASH

5.1 Free-Electron Lasers

Free-electron lasers (FELs), along with radio telescope arrays and gravitational-wave observatories, belong to the most demanding applications for timing synchronization systems [KI09; Sch05; CS06; Kol+16]. This section introduces the core principles behind FEL light, its uses, and the consequential synchronization requirements found in such machines. We will also introduce the two FEL facilities at Deutsches Elektronen-Synchrotron (DESY), European XFEL and FLASH, where the pulsed laser-based synchronization (LbSync) system [Sch+15] that motivates this work is deployed, providing unparalleled synchronization performance.

5.1.1 Synchrotron Radiation and Pump-Probe Experimentation

Light with suitably controlled properties is essential to current and future experimental science across a wide breadth of fields, including structural biology, chemistry, material sciences, medicine, solid-state physics and many more [OF01; nat]. Various experimental techniques exist. For example, in material science, extremely energetic light flashes are used to study matter under extreme heat in order to find new materials and manufacturing processes [Zas+21]. In diffractive imaging and scattering crystallography, light-matter interaction causes a diffraction pattern on a detector, which can be used to infer structural properties of the analysed sample, such as a protein, (macro)molecule or atomic lattice [e.g. Vas+21]. Spectroscopy deducts composition, chemical bondings, and other physical properties of matter from absorbed and emitted photons at characteristic wavelengths [Kal+21]. For the inference to be accurate, the properties of the light used to probe the matter, such as wavelength, intensity, coherence, beam width and divergence, pulse duration, etc. need to be precisely controlled. More and more, not only the structure but also the dynamics of (sub)molecular particles, such as chemical reactions, are of interest. Because these processes are ultra-fast, extremely short light pulses are required as probe to temporally resolve the reaction dynamics. At the current frontier of research are the attosecond dynamics of electron motion in nanostructures and molecular orbitals, governing chemical composition and information transport in biological systems [KI09]. The generation of corresponding attosecond light pulses required for the experimental study of these electron dynamics has been recognized with the Nobel Prize in Physics in 2023 [LHu24].

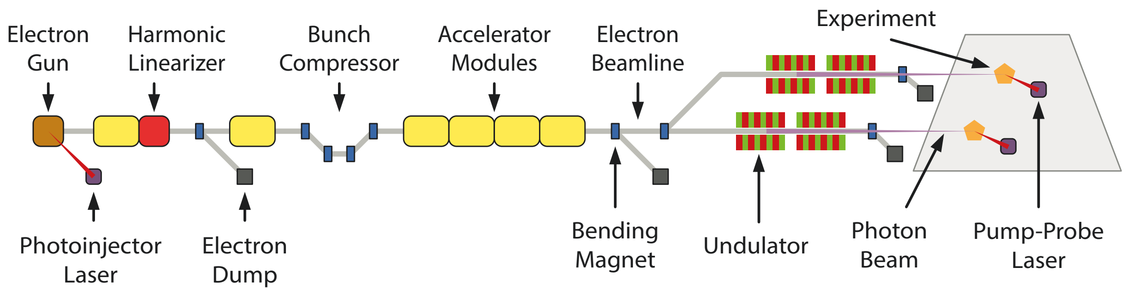

Figure 5.1: Simplified schematic depiction of an FEL based on a LINAC.

Due to the wide range of applications, and being a key-enabler to answer many of today's most prominent experimental research questions, the development of light sources producing shorter and ever more brilliant[1] photon pulses has gained significant attention and research progress. Past and present world-class light sources are universally based on synchrotron radiation, i.e. electromagnetic radiation emitted by electrons when deflected from their current path (e.g. by the Lorentz force exerted by a transverse magnetic field). This was originally observed in circular particle accelerators and just parasitically exploited at first, but eventually facilities dedicated to the generation of such synchrotron radiation were designed and built all over the world [Cou14]. An overview of the development generations of light sources is given in [Win98]. The current fourth-generation of light sources is spearheaded by FEL machines, which will be described next. Future (fifth-generation) light sources are expected to exploit plasma-based electron acceleration, not only further increasing the light's brilliance and achieving even shorter pulse durations, but most importantly reducing the scale and cost of the electron acceleration section by orders of magnitude [Lit+14; Hid+14; Ass+20].

The following paragraphs are loosely based on references [Sch+14; SSY00; Jae+20]. As the name already suggests, and opposed to "conventional" (i.e. quantum) lasers, where the laser light is emitted by electrons bound to atoms in a gain medium, FELs utilize electrons propagating freely in a vacuum. This removes the quantum-mechanical restriction of bound-electron lasers, where an electron may only absorb or emit energy in discrete quantities corresponding to the energy levels of the atom it is bound to, fixing the wavelength of emitted photons. With FELs on the other hand, the wavelength depends primarily on the electron energy, and can thus be varied continuously, extending into the hard X-ray regime for high enough energies. Schematically depicted in Figure 5.1, state-of-the-art FELs consist of three top-level pieces: 1. A *particle accelerator* providing packets of high-energy, i.e. ultrarelativistic electrons, 2. one or more *undulation sections*, generating the desired light from the electrons and supplying the 3. *experimental stations*, comprised of additional optics, the light-matter interaction point and a detector.

Although sometimes seen as a separate entity, we consider the particle source that generates bundles of free electrons (bunches) in a so-called *electron gun*, and subsequent injection into the main accelerating section as a part of the particle accelerator. Figure 5.2 shows a particular kind of electron gun, the *photoinjector*, where a high-energy laser pulse incident on a photocathode releases up to a few nC worth of free electrons, which are immediately accelerated by an electric field in order to mitigate space-charge effects pushing the cloud of electrons apart. The subsequent accelerator section may either be realized as a ring or linear accelerator (LINAC) structure, or a combination of both, however due inherent

[1] Quantity describing the total photon flux in a given conical opening angle per bandwidth. Also called "brightness" [Sch+14].

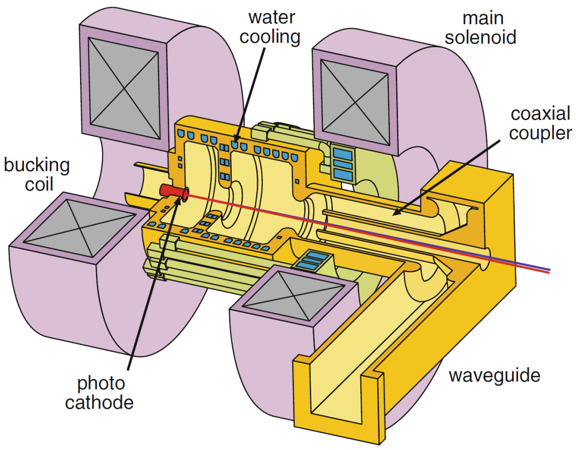

Figure 5.2: CAD drawing of a photoinjector showing the paths of laser and electron beams in red and blue, respectively. Drawing by E. Vogel (CC-BY 3.0) [SF15].

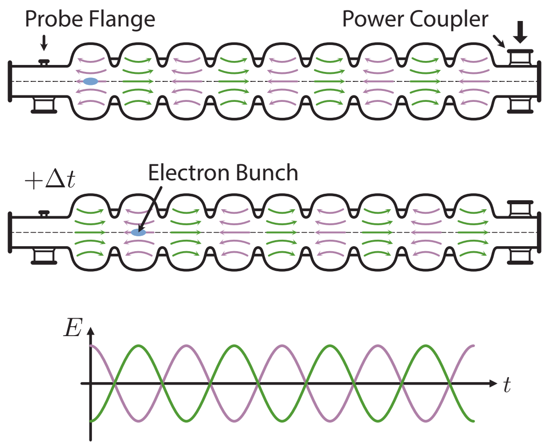

Figure 5.3: Acceleration of an electron bunch in a 1.3 GHz TESLA super-conducting RF cavity.

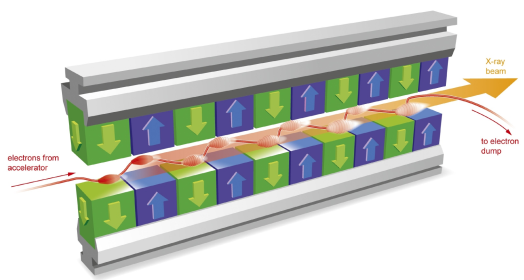

Figure 5.4: Schematic depiction of SASE in an undulator. ©European XFEL GmbH.

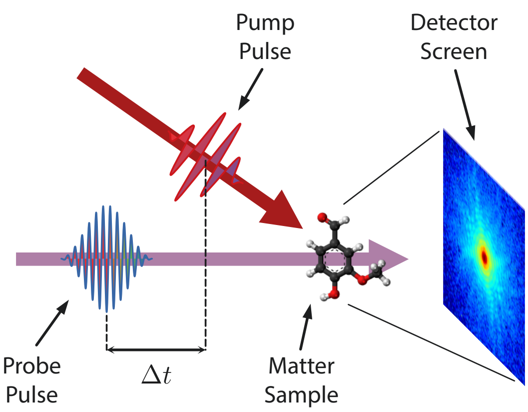

Figure 5.5: Schematic illustration of a pump-probe experiment.

equilibria conditions affecting electron bunches in a storage ring, their emittance and bunch lengths are fundamentally limited when compared to LINACs. The acceleration itself takes place in specially designed RF cavities, which support sustained oscillation of a free-space RF field with frequency such that the field in each cavity cell switches polarity exactly at the rate as the electrons move from cell to cell, causing the electrons to always experience an accelerating field. This is depicted in Figure 5.3. Due to their low mass, the electrons become ultrarelativistic and move essentially with the speed of light after passing only a few cavity cells. At this point, the electrons merely gain energy rather than effective speed from further acceleration, and the electron bunch can be focussed and steered using dipole, quadrupole, and sextupole magnets, similar to ray optics.

Having used the accelerator to bring the electrons up to the energy that is necessary for them to emit synchrotron radiation at the desired wavelength, the electrons undergo a self-amplified spontaneous emission (SASE) process in the undulation section. Here, the electrons are forced onto a sinusoidal trajectory by an *undulator*, an arrangement of alternating magnetic fields, shown in Figure 5.4. On their sinusoidal path, the electrons spontaneously emit synchrotron radiation. The SASE process is characterized by exponential increase in power due to an effect called *microbunching*, where electrons interact with the field emitted by other electrons, modulating the electron bunch into microbunches exactly half a wavelength apart. Subsequent spontaneous emission leads to quasi-coherent radiation, which scales in intensity with the number of electrons squared, rather than just the number of electrons, as is the case for incoherent emission. An alternative scheme to SASE is so-called *seeding*, where coherent emission is achieved not as a result of spontaneous emission, but by stimulated emission of the electron bunch by superposition with seed laser light. The advantage is reduced stochastic fluctuation in photon energies occurring with the SASE process because of relying on spontaneous emission. However, the generation of suitable seed laser light at the right wavelength is still a challenge. Currently, new schemes such as high-gain harmonic generation are explored [Lab+11].

The electrons are then separated from the beam path of the photon pulse using steering magnets, which only affect charged particles. The photon pulse continues to propagate to the experimental station, where additional optics may be used to focus and condition the beam depending on the experiment. For time resolved experiments, i.e. those studying the dynamics of chemical reactions and other processes, the experimental station also has to provide a "trigger" for the process being analysed, whose relative timing to the arrival of the FEL pulse determines the state of the process at which it is imaged. Such experiments are known as *pump-probe* experiments [Ati+24]. The pump excites a sample which then decays, or triggers a more complex reaction, and the probe captures the state of the process at a specified delay. This is depicted in Figure 5.5. While in this case the FEL pulse is the probe, the pump pulse is generated by a dedicated (non-trivial) laser system. To avoid confusion with other optical pumps, e.g. those used to excite the active media in optical amplifiers, we (imprecisely) call this laser system the "pump-probe laser".

5.1.2 Synchronization Requirements and Challenges

We now discuss the various precision synchronization needs and challenges in FEL machines, particularly those based on LINACs such as FLASH and European XFEL introduced later. As motivated in the beginning of this section, from a high-level perspective, the synchronization needs are derived from the experiment to be performed. In order to achieve the necessary spatial resolution, the wavelength of the synchrotron radiation needs to be sufficiently short, which is known as the diffraction limit. To resolve atomic structures, sub-nanometer wavelengths are required. Hard X-ray radiation with such short wavelength can only be generated from sufficiently energized electrons. To attain the required energy, the electrons need to traverse a sufficiently strong potential difference or sufficiently many subsequent potential differences. As the electric field strength is limited even in super-conducting RF cavities due to a lack of suitable amplifiers and electrical losses, scaling the number of RF cavities passed by the electrons is currently the only alternative, leading to very long LINAC sections in the order of many hundred meters. In addition, the RF field in each cavity needs to be precisely synchronized in phase to the arrival time of the electron bunch in order to ensure optimal energy transfer to the beam. Illustratively, in the extreme, a $180°$ phase shift would lead to acceleration in the opposite direction. Current L-band machines pose requirements in the order of $0.01\,\%$ and $0.01°$ RMS field stability for amplitude and phase, respectively [Sim06]. Consequently, LINACs for cutting edge FELs have dozens of RF stations, spanning multiple hundred metres, that require precision synchronization to generate and feed a phase stabilized RF field to the cavities (cf. introduction of the European XFEL below). Vice versa, the arrival time of the electrons also needs to be stabilized. Indeed, a phase error in the RF field of $0.01°$ at $1.3\,\text{GHz}$ is equivalent to a relative arrival time error of the electron bunch centroid of $21.4\,\text{fs}$. This is

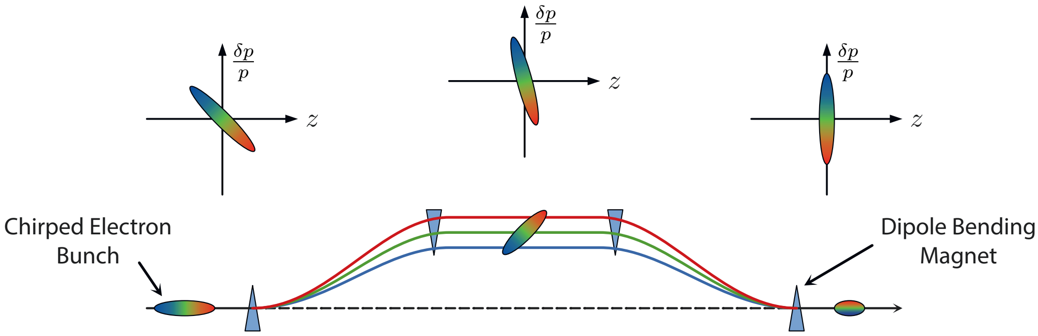

Figure 5.6: Schematic depiction of electron bunch compression in a magnetic chicane. Particles in the incoming chirped bunch have different energies and thus take different trajectories due to the varying deflection angle in the bending dipoles, leading to longitudinal compression.

achieved by synchronizing the laser pulses that eject the free electron bunches from the photocathode in the photoinjector and all RF stations to a common RF master timing reference.

Other than wavelength, the FEL pulses (and hence the electron bunches) also need to be sufficiently short in overall pulse duration, in order to temporally resolve ultra-fast dynamics hindered otherwise by motion blur effects. In practice, this adds even stronger synchronization requirements on the accelerating RF field: Because the optical power incident on the photocathode in the electron gun is limited, and because the initial bunch size is limited by space charge effects, the initial electron bunch is generally much longer than required later in the undulation section. A bunch compression scheme is needed to reduce the longitudinal extent of the electron cloud [Sch+14, chap. 8.6]. Since all electrons move with approximately the same speed, which is essentially the speed of light after passing only a few accelerating cavities, the only way to compress the bunch is to make electrons further ahead in the bunch take a longer path than those at the end. This is depicted in Figure 5.6. The electrons traverse a magnetic chicane, where early electrons are diverted to a longer path than trailing electrons. Because sufficiently fast modulation of a magnetic field to achieve this effect is very challenging, a viable alternative is to change the Lorentz force exerted on the electrons by modulating the electron energy across the longitudinal axis of the bunch. This is effectively achieved by introducing a marginal phase shift φ between electrons and accelerating RF field, such that the electrons are accelerated *off-crest*. The natural curvature of the sine-shaped RF field strength imprints an energy gradient onto the beam, called *chirp*, which then leads to the desired bunch compression in the magnetic chicane. A side effect is that the overall arrival times of the electron bunch at subsequent cavities and by extension of the FEL pulse at the experiment are affected by marginal error in the average bunch energy, which will offset the path delay in the magnetic chicane. The relevant equation for an incoming bunch with upstream arrival time jitter $J_{t,\mathrm{in}}$ is given by

$$\frac{J_{t,\mathrm{out}}^2}{c_0^2} = R_{56}\left(\frac{\sigma_A}{A}\right)^2 + \left(\frac{C-1}{C}\right)^2 \left(\frac{\sigma_\phi}{\omega_{\mathrm{RF}}}\right)^2 + \left(\frac{1}{C}\right)^2 \frac{J_{t,\mathrm{in}}^2}{c_0^2}, \tag{5.1}$$

where R_{56} is a characteristic dispersion coefficient of the magnetic chicane based on its geometry and magnetic field strengths, C is the resulting bunch compression factor, and A, ϕ and ω_{RF} are the amplitude, relative phase shift and the angular frequency of the driving RF field, respectively [Sch05].

This however can also be utilized as an actuation mechanism to stabilize the overall FEL pulse arrival time at the experiment [Din18; Czw+21], which is difficult to achieve feedforward from the onset due to

Figure 5.7: Overview of the DESY campus (right) and the European XFEL campus (left) at the other end of the 3.5 km LINAC.

the long duration of the photoinjector laser pulse and because of microscopic mechanical disturbances distorting the electron beam path of several hundred metres from photoinjector to experimental station by only a few micrometres [Grü+25]. The arrival time measurement for feedback can either be performed on the electron bunch after compression (sufficiently close to the undulation section) [Czw+21], or ideally directly on the FEL pulse at the experimental station [LG17; Grü+19]. However, due to lack of suitable parasitic (i.e. non-destructive) arrival time measurement schemes for the FEL photon pulse, electron bunch arrival time measurements are more common. In either case, the arrival time measurement again requires a highly stable time reference to compare the arrival time of photons or electrons to, adding yet another client in need of synchronization with respect to other parts of the machine.

Finally, on the experimental side, we have already established that in order to temporally resolve reactions via pump-probe experimentation, the arrival time of the FEL pulse (probe) with respect to the pump needs to be controlled to shorter time scales than that of the process itself. Having discussed arrival time stabilization of the electron bunch and by extension (approximately) the FEL pulse, the stabilization of the pump pulse is generally addressed by stabilizing the seed laser for the pump-probe laser system. This however is only sufficient if the subsequent optical components and the beam transport path to the interaction point are suitably stabilized as well, actively or passively.

Remark 5.1. Having discussed various needs to synchronization in FEL machines, we remark that these are only some examples among many possible applications that we deem essential with respect to the following exposition. Other examples include FEL seeding, experimental techniques other than pump-probe schemes such as e.g. terahertz streaking, and additional electron beam diagnostics and conditioning, e.g. laser heating.

5.1.3 DESY and the Free-Electron Lasers European XFEL and FLASH

We will now introduce the European XFEL and FLASH as two specific examples where a pulsed LbSync system (cf. Section 5.2) is employed to satisfy the previously outlined needs for synchronization in these facilities. As shown in Figure 5.7, both facilities are fully (FLASH) or partly (European XFEL) located on the Hamburg-Bahrenfeld campus of DESY. Founded in 1959, DESY today is one of the largest of the Helmholtz Association of German Research Centres, with particular emphasis on matter related research, in particular particle accelerators, photon science, particle, astroparticle, and high-energy

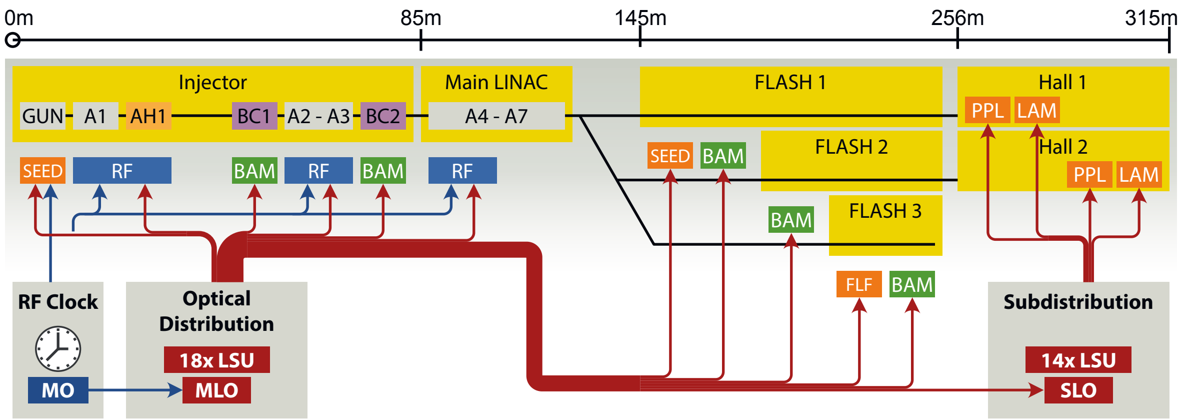

Figure 5.8: The DESY LbSync system at FLASH.

physics [see e.g. DES24]. It hosts various experimental science facilities and laboratories, ranging from accelerator research and development, dark matter research, electron research to third and fourth generation light sources, and has significantly contributed to the development of the related technologies, for example with the super-conducting 1.3 GHz TESLA cavity [Aun+00].

FLASH

Originally designed towards the specifications of the linear electron-positron collider TESLA[2], the TESLA cavity was initially put to use in the TESLA test facility (TTF) at DESY, which eventually evolved into FLASH, which was commissioned in 2004 as the world's first extreme ultraviolet and soft X-ray light source [Ack+07; Bey+23]. The layout of FLASH is depicted in Figure 5.8 along with the deployed LbSync system developed also at DESY. Relevant light-source parameters of FLASH in comparison to other facilities are listed in Table 5.1. As seen in detailed schematic Figure 5.9, its LINAC consists of a photoinjector [Sch+11; Wil+11], followed by three acceleration sections separated by two bunch compressor chicanes, spanning a total length of 256 m [Vog+11]. The train of electron bunches is then split into two photon beamlines with multiple experimental stations each, coming to a total length of about 315 m. The facility's synchronization is derived from a 1.3 GHz RF master oscillator (MO) (cf. Section 5.2.2) [Gas+20]. To overcome limitations of signal distribution via unstabilized coaxial cables, primarily significant drift caused by humidity and temperature changes, the MO is signal distributed optically to "far away" client stations using a LbSync system with stabilized optical fibre links (not depicted). As of today, clients include the photoinjector laser, optical reference modules (REFM-OPTs) for RF resynchronization at the RF stations (compensating coaxial RF cable drift), bunch arrival time monitors (BAMs) for electron arrival time measurement and longitudinal feedbacks, the laser heater system, a prototype seeding laser, and various experimental laser systems, including two modular pump-probe laser systems, the laser systems for the terahertz streaking experiment at FLASH and the plasma wakefield experiment FLASHForward. Recently, prototype laser arrival time monitors (LAMs) (cf. Section 5.2.3) have been installed at various interaction points in order to compensate drifts accumulated during pump-probe laser beam synthesis and transport to the interaction point [Sch+24].

FLASH is currently undergoing a multiple-phase upgrade with the FLASH 2020+ project [Bey+23]. The goal of the project is to increase the versatility of FLASH by equipping both SASE beamlines with

[2] Tera-Electronvolt Energy Superconducting Linear Accelerator, now superseded by the ILC project [Mic19]

Facility	Location	Commissioned	# Photon beamlines	Max. e^- energy (GeV)	Peak brilliance	Min. λ (nm)	Min. T_{pulse} (fs)	Max. $\frac{\text{pulses}}{s}$	LINAC RF Technology
FLASH 2020+	DESY Hamburg, Germany	2005	2	1.35	$>1 \cdot 10^{31}$	4	<10	$5 \cdot 10^3$	SCRF
European XFEL	DESY Hamburg, Germany	2017	3	17.5	$>5 \cdot 10^{33}$	0.05	<3	$27 \cdot 10^3$	SCRF
LCLS-II	SLAC Stanford, California	2023	2	15	$>2 \cdot 10^{33}$	0.05	<10	$1 \cdot 10^6$	SCRF
SwissFEL	PSI Aargau, Switzerland	2018	2	5.8	$>1 \cdot 10^{33}$	0.1	<20	100	warm
SACLA	Riken SPring-8 Center Harima, Japan	2011	3	8.5	$>1 \cdot 10^{33}$	0.08	<10	60	warm
PAL-XFEL	PAL Pohang, South Korea	2016	2	10	$>1.3 \cdot 10^{33}$	0.06	<20	60	warm
SHINE	SARI Shanghai, China	2026 (Planned)	8	3	$>1 \cdot 10^{33}$	0.05	<20	$1 \cdot 10^6$	warm

Table 5.1: Comparison of selected soft and hard X-ray free-electron lasers. Derived from [Hua+21; Eur].

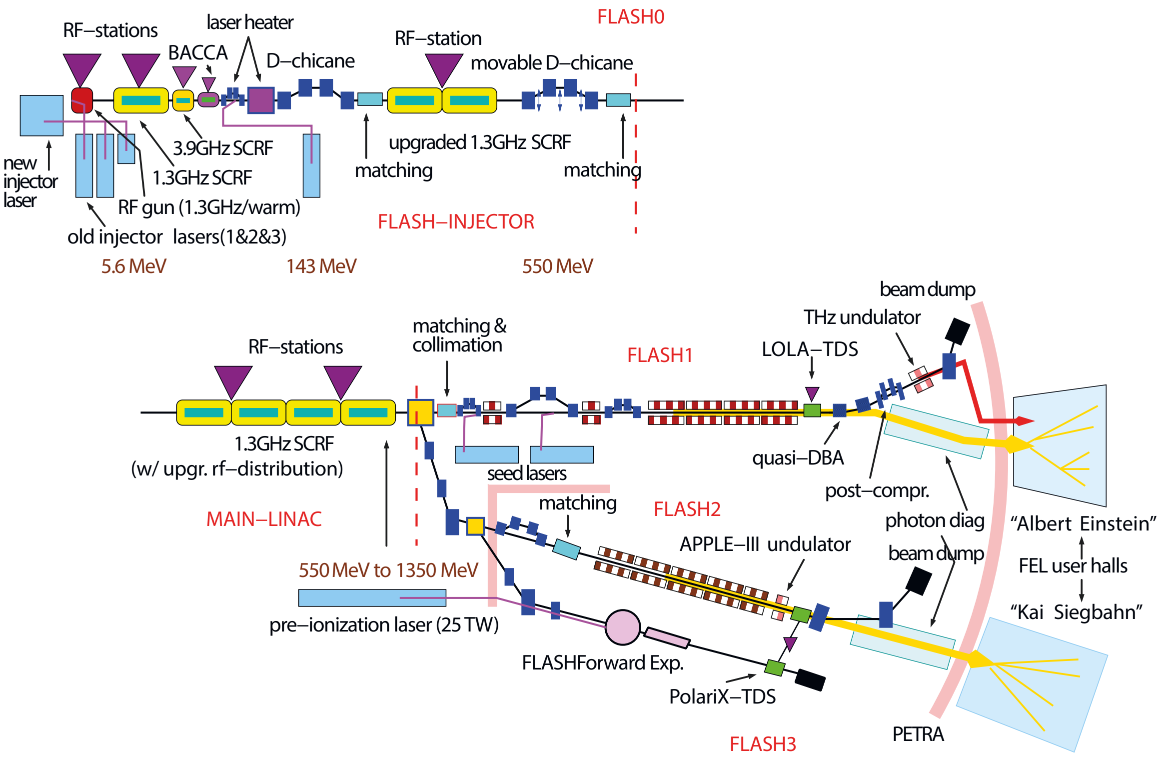

Figure 5.9: Schematic layout of FLASH after its 2020+ upgrade. Adapted from original drawing by M. Vogt et al. (CC-BY 4.0) [Vog+23].

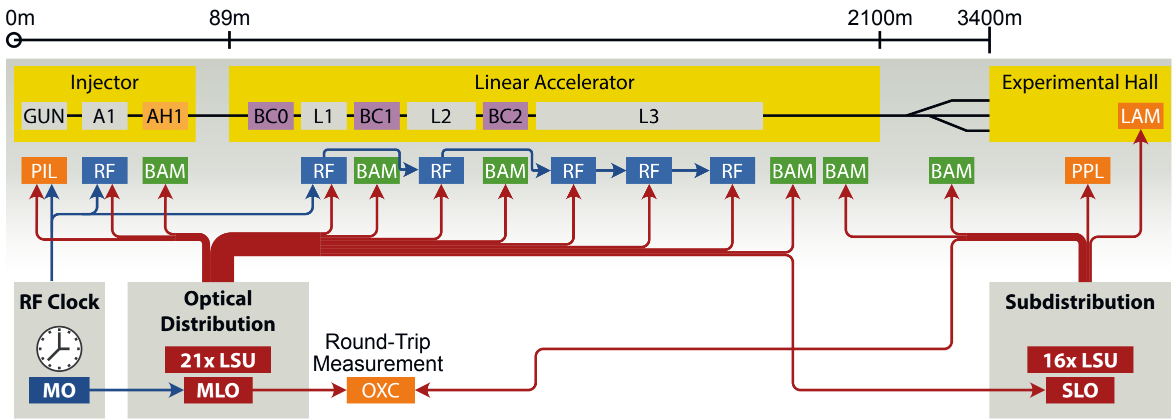

Figure 5.10: The DESY LbSync system at European XFEL.

tunable undulators, allowing for fully independent operation. The original FLASH beamline is also to be converted to fully seeded FEL operation, allowing for new experiments that require tighter control over the FEL pulse phase. The upgrade also entails a significant extension of the LbSync system with the addition of a subdistribution on the experimental side, which is needed to cover the significant increase in clients in need of synchronization with the 2020+ project and beyond.

European XFEL

Derived from a supplement to the TESLA technical design report (TDR) and based on experience gained from FLASH, the European XFEL was realized by 2017 and is currently the world's most powerful light source in terms of electron energy and peak brilliance (cf. Table 5.1) [Hua+21]. Although an international collaboration project, DESY holds a majority share of the European XFEL GmbH that built and operates the machine, and is in particular responsible for the operation of the LINAC [DES24; Eur24]. With an overall facility length of 3.4 km, it exceeds FLASH by approximately a factor 10, allowing for FEL pulses that extend into the hard X-ray regime and thus an entirely new range of experiments. The maximum electron energy of 17.5 GeV is realized with a 1.7 km long, superconducting LINAC, consisting of 768 TESLA cavities [Sch+23]. As seen in Figure 5.10, the layout of the European XFEL is indeed very similar to FLASH, with most of the differences just resulting from upscaling the facility. Nevertheless, the synchronization challenges are significantly hardened from a practical point of view when compared to FLASH, due to the sheer spatial extent of the facility. To this regard, European XFEL from the beginning had a subdistribution located at the experimental hall. The facility is also continuously being upgraded. Examples are the BAMs located in the undulation sections [Czw+24; SCS24], which were installed in 2023 and 2024 to yield additional insight into microseismic effects on the beamline tunnel caused by the tide of the North Sea, ocean waves in the North Atlantic, as well as earthquakes across the globe. All the above effects have already been observed in diagnostic data [Grü+25], but the total influence accumulated up to the interaction points are hard to predict – if at all possible – and are a hindrance for future experiments, requiring even shorter pulses and tighter arrival time synchronization. Future421251 synchronization related improvements and extensions to the machine include LAMs for tighter pump-probe arrival time stabilization at the interaction point (cf. Section 5.2.3), and the ASPECT project [Yan+23].

5.2 The DESY Pulsed Laser-Based Synchronization System

Having introduced the need for and context of the LbSync systems at European XFEL and FLASH, this section reviews technical aspects about implementation and deployment, with a primary focus on new developments and improvements since the previous works [Sch+15; Heu18; Lam17; Tit17; Lam+19]. We stress that the following discussion is mostly concerned with the methods and design choices utilized by the DESY LbSync system and is by no means universal. If deemed beneficial, comparisons and references to selected alternative implementations will be given.

5.2.1 High-Level Overview

Before the individual elements of the LbSync system are discussed, a summary of the central concepts and the high level topology of the system are reviewed. Recalling the structural layout of the LbSync systems at FLASH and European XFEL from Figures 5.8 and 5.10, its topology in the form of a synchronization tree is evident. The root is given by a drift-stable facility master reference, cf. Section 5.2.2, generated electronically and distributed to the particle accelerator's RF stations via (drifting) coaxial cables. The additional need to synchronize several laser systems in the accelerator and in the experimental hall to this reference, in combination with the greater fidelity of optical-to-optical timing detection methods, motivates the additional all-optical timing distribution. Since the optical distribution scales reasonably well, and because optical fibre links are generally easier to stabilize than coaxial cables, cf. Section 5.2.4, the optical timing distribution system is also used to provide a drift-stable optical reference to the RF stations, resynchronizing the disturbed RF reference received via coaxial cable [Lam+24], cf. Section 5.2.5. This approach requires good transfer of the RF master reference to the optical domain via a master mode-locked laser (MLL) in an optoelectric PLL, cf. Section 5.2.3. Close proximity of the two oscillators ensures minimal accumulation of out-of-loop disturbances. The same holds true for the photoinjector laser, which has relaxed jitter requirements due to long pulse lengths and jitter compression in the bunch compressors. Recently, additional drift compensation of the photoinjector laser against a stabilized optical fibre link from the optical master has been found to significantly benefit long-term FEL stability and maintaining machine timing [Mül+19].

5.2.2 Drift-Stable Facility RF Reference

At the root of the implemented synchronization network lies a 1.3 GHz MO, developed jointly by DESY, University of Warsaw and industry partners, providing exceptional phase noise performance and long-term stability. The oscillator is designed around a 1.3 GHz dielectric resonator oscillator (DRO), stabilized to GNSS via an intermediate oven-controlled crystal oscillator (OXCO), and a subsequent high-power amplifier (HPA). In the case of the 2014 version installed at European XFEL, the non-spurious integrated jitter accumulates to just 11 fs over 10 Hz to 10 MHz, cf. Figure 5.11. Spurious content (virtually entirely from the 50 Hz line) contributes another 23.82 fs in RSS. The latest generation, recently installed at FLASH has improved suppression of the 50 Hz power grid disturbance and reduces the phase noise in the critical region around 10 kHz with an improved DRO [Gas+20]. The total RMS jitter including spurious is just 12.83 fs. The next generation, to be installed at European XFEL in 2025, will further improve on the current design with better rejection of microphonics and environmental noise, e.g. from cooling fans, reducing the remaining spurious content in the spectrum. Additionally, the noise floor is to be lowered to approx. −178 dBc/Hz by tweaking the HPA.

We shortly discuss the relevance of the phase noise stability of the MO as common reference for the entire facility. The question is to which extent all other systems follow possible phase deviations from absolute

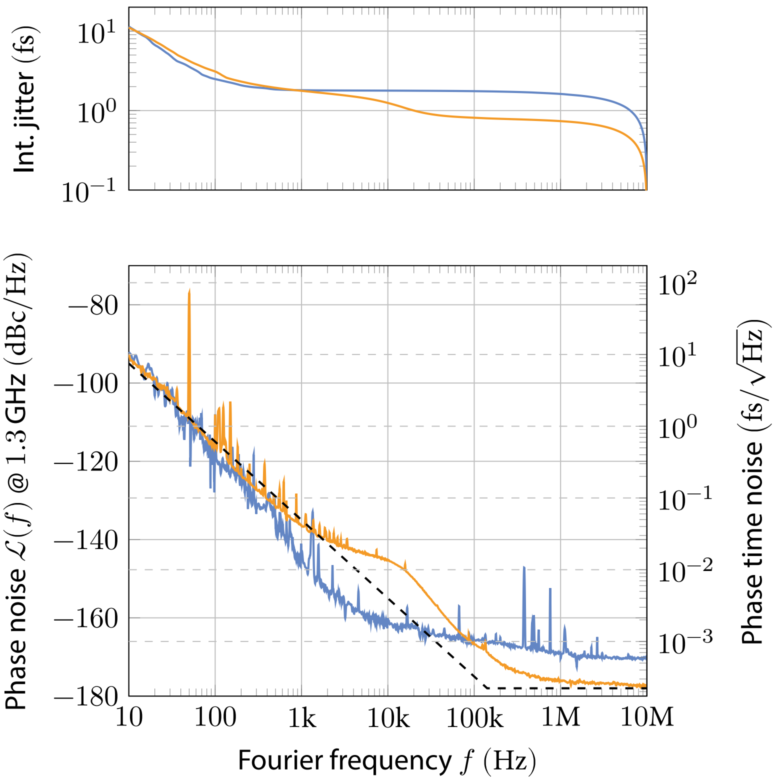

Figure 5.11: Phase noise spectra of the DESY 1.3 GHz MOs at FLASH (——) and European XFEL (——), courtesy of H. Pryschelski. (- - -) shows the prospective phase noise of future MO generations.

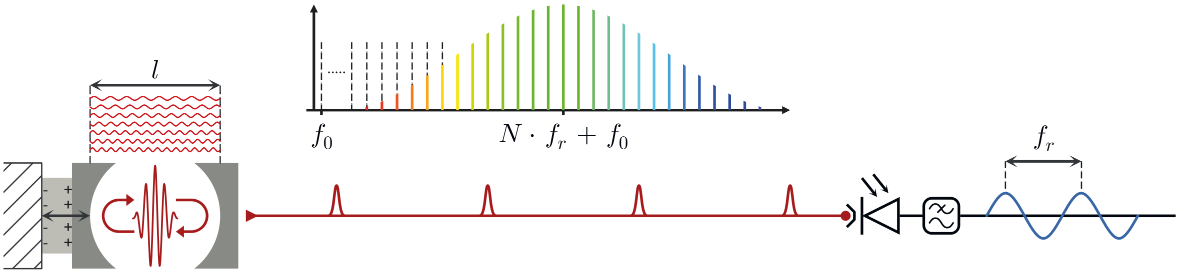

Figure 5.12: Illustration of optical frequency division with a femtosecond mode-locked laser. Several modes are superpositioned in the laser cavity, creating a train of femtosecond pulses. The cavity length and repetition rate are tuned with a piezo mounted mirror. Photodetection and filtering extracts the microwave frequency f_r, thus dividing the optical carrier.

time by the MO in unison, and at which points absolute phase noise is converted to relative phase noise. As has been remarked in Remark 3.15, the latter effect certainly appears on timing links stabilized with the auto-correlation principle, which is the case for the stabilized fibre links in the LbSync system. The former question is more difficult to answer accurately, as it depends on the locking bandwidths of all involved PLLs. Qualitatively, good phase noise at high offset frequencies and a low white floor are generally desirable for the RF field control in the accelerator because of the limited bandwidth of the resynchronization via the REFM-OPT. In general, this limited bandwidth is of minor concern, because the signal transport via the unstabilized coaxial RF cable picks up negligible high-frequency jitter and because of the low-pass characteristic of the RF cavities. High-gain RF field feedback regulation however will amplify high-frequency absolute phase noise in the MO signal, which limits overall performance of the accelerator's RF system [Lud+06]. Finally, absolute phase noise is of course relevant when comparing time resolved measurement data from two devices at different locations within the accelerator, as the absolute phase error of the reference integrated over the propagation delay difference to the devices shows up directly as measurement error.

Absolute timing stability of the MO therefore is indeed relevant for highest possible FEL performance and frontier experiments. This immediately raises the question of how the absolute phase noise of today's best RF sources can be improved further without sacrificing RF power or amplitude stability, which are also required for satisfactory machine operation. A detailed discussion is beyond the scope of this work, but we point the reader to the following recent references in the field of frequency metrology, investigating novel RF oscillators and generation of ultra-stable microwave signals from the optical sources [For+13; Xie+17; Nar+22; HJK21; Qui23].

5.2.3 Synchronization of Mode-Locked Laser Oscillators

Compared to microwave, optical sources and signals have several advantages, such as lower transmission loss, almost no sensitivity to EMI, and large information bandwidth. In time and frequency metrology, atomic energy transitions in the optical domain provide the most stable frequency standards known to mankind, achieving a frequency stability that is two orders of magnitude better than the current basis for the definition of the SI second, the microwave caesium-133 atomic clock standard [FB19; Dim+24; McG+19; Nak+20a]. However, microwave is still key in many practical applications, especially those driven by conventional electronics. Unsurprisingly then, the connection of the optical and microwave domains has been a key challenge of the past decades, and was eventually rendered practical using optical frequency combs, the development of which was in part recognized with the 2005 Nobel Prize in Physics [Hän06]. The principle of optical frequency division (OFD) performed by the frequency comb is illustrated in Figure 5.12. The most common implementation involves a (femtosecond) mode-locked laser (MLL), whose optical cavity sustains many hundred thousand phase-coherent longitudinal resonant modes with fixed wavelength spacing [FB19]. The optical frequency ν_n of the n^{th} mode is described by the comb equation

$$\nu_n = (n + n_0) \cdot f_r + f_0, \tag{5.2}$$

where n_0 is a large integer, $n = 0, ..., N$ is the mode number, f_r is the mode spectral distance, and f_0 is an offset frequency. In the time domain, the spectral comb, in accordance with the limit Fourier theory of Dirac combs, relates to a pulse train of extremely short pulses with hyperbolic secant waveform [KS16]. The repetition rate of the time domain pulses lies in the microwave domain (tens or hundreds of MHz are common), and is exactly equal to the spectral spacing of the optical modes, f_r, which in turn is given by the optical cavity length. Measuring the repetition rate with a suitable photodetector facilitates the desired phase-coherent optical frequency division by factors in the range 10^4 to 10^6 [Bar+05].

Utilization of the MLL as an optical frequency comb, academically meaning the use for metrological purposes, requires stabilization of the comb parameters f_r and f_0, where f_r, as already discussed, relates to the optical cavity length, and f_0, the so-called carrier-envelope offset, describes the rate of phase-slippage of the optical carrier wave with respect to the pulse envelope's centre of mass. For timing and synchronization purposes, stabilization of the repetition rate f_r is indispensable, as temperature and mechanical vibrations significantly impact the mechanical stability of the optical cavity and hence the long- and medium-term timing stability of the pulse train [Pas04]. Stabilization of the relative optical carrier phase f_0 is primarily necessary for optical spectroscopy applications, and for pulse shaping of few-cycle (attosecond) pulses [UHH02]. Stabilization methods for f_0 generally involve self-referencing and stabilization of a single mode in the optical spectrum to a high-fidelity optical reference cavity, such as an optical atomic clock. Details of these schemes are beyond the scope of this work, and the reader is directed to [FB19] for further information. Stabilization schemes for f_r involve feedback actuation of the optical cavity length either via temperature controlled thermal expansion, optical fibre stretchers or movable mirrors. Often combinations of two actuation mechanism are used to have both large dynamic range and fast response (high feedback bandwidth). The detection of the repetition rate f_r involves either comparison of the photodetected signal to a highly-stable RF reference and reversing the OFD gear-work to stabilize the optical mode spacing, or using optical self-referencing and optical mode referencing to jointly detect errors in f_r and f_0. Because the phase fluctuations of the original optical carrier get reduced by the same OFD factor over the entire range of offset frequencies when detecting f_r, this yields an extremely low phase noise microwave signal [FB19].

The MLLs used as master laser oscillators (MLOs) for FLASH and European XFEL generally show poor free-running frequency stability but exceptional phase noise at high offset frequencies. This is depicted in Figure 5.13. The poor frequency stability is not relevant, as the lasers are locked to the long-term stable MO. The true phase noise at high-offset frequencies is difficult to characterize truthfully, because, at a certain point, amplitude modulation (AM)-phase modulation (PM) conversion in the photodiode used to extract an RF signal for the phase noise analyser (cf. Section 5.2.3) dominates the true phase noise in the optical signal. The true phase noise can be estimated by balanced optical cross-correlation of two identical lasers (cf. Section 5.2.3), which is insensitive to intensity fluctuations, and attributing half of the phase noise to each laser. As shown in Figure 5.14, the estimated true phase noise of a MLL oscillator is significantly lower than what is measured using a photodiode. For an overview of

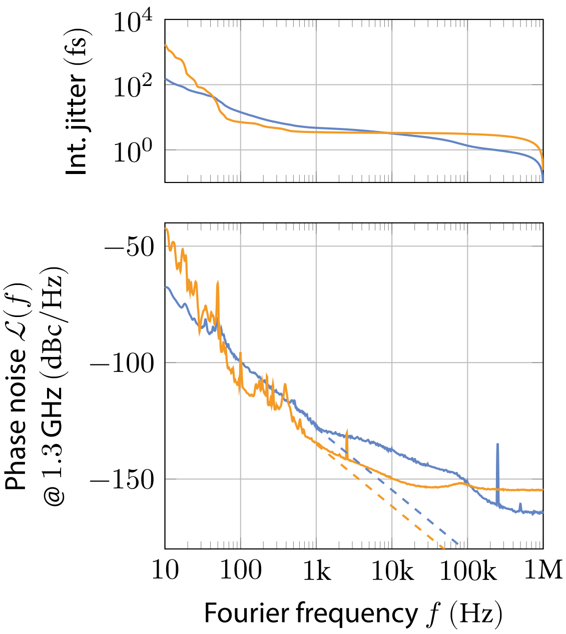

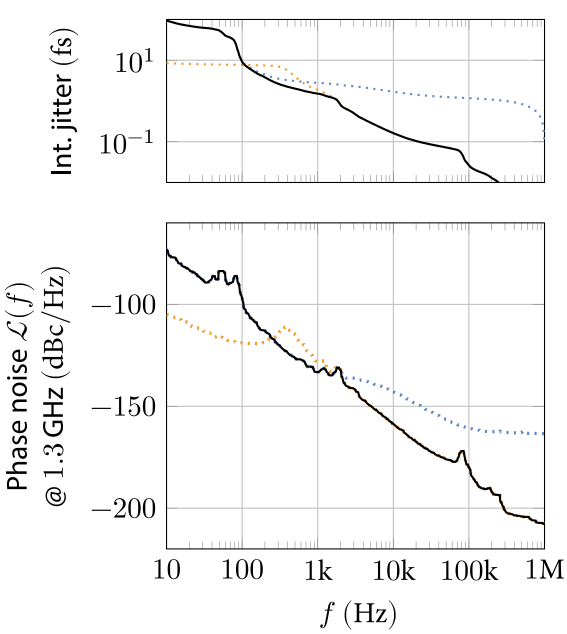

Figure 5.13: Phase noise spectra of the Menhir Photonics Menhir-1550 FLASH MLO (——) and the NKT Origami-15 European XFEL MLOs (——). The dashed lines estimate the true high-offset-frequency phase noise in absence of excess noise in the direct detection measurement setup.

Figure 5.14: Characterization of the phase noise of a Menhir Photonics Menhir-1550 MLL. (······) shows a direct detection based measurement, (······) the average residual phase noise of two coarsely locked identical lasers measured using a BOXC. The individual measurements are falsified by AM-PM in the photodiode and the phase locking, respectively. The true absolute phase noise is estimated from the combination of both measurements (——).

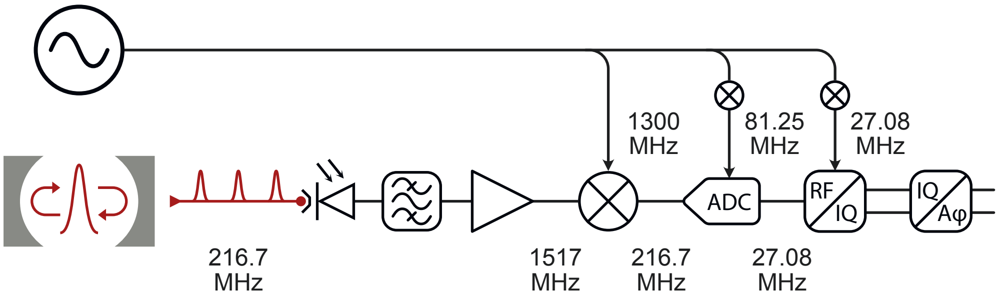

Figure 5.15: Harmonic heterodyne phase measurement of a MLL with an RF reference with a "golden" undersampling IQ-detector.

timing noise in MLL, see e.g. [KS16; Pas04]. For excess noise in photodetection of optical sources see e.g [Qui23].

Laser-to-RF Synchronization

Although in the future a fully optical master reference for the synchronization systems at FLASH and European XFEL is conceivable, the current generation is still based on the DESY RF MO as facility master, which is why the LbSync system uses a MLL in an optoelectric PLL with the MO as reference to transfer the MO phase to the optical domain. Stabilization of the carrier-envelope offset f_0 is not necessary in this application. High-sensitivity detection of the phase error of the MLO's optical pulse train with respect to the RF MO signal is crucial to ensure the required synchronization performance in the FEL facility. In the following, we summarize techniques used in the DESY LbSync system and notable alternatives from the literature.

Heterodyne Detection Scheme The heterodyne detection technique is known from standard RF applications, especially communication technology, and is described in detail in [e.g. Bra18]. As depicted in Figure 5.15, a suitable harmonic of f_r is extracted from the photodetected optical pulse train by band-pass filtering and subsequent RF amplification. The output signal of the photodiode, filter, amplifier (PFA) chain is then mixed with the reference to produce an intermediate frequency (IF) that is then digitized with an ADC for DSP IQ detection. When deriving the ADC clock from the same RF reference, the IF can be under sampled, i.e. the IF is purposely aliased into the first Nyquist zone, yielding a lower-frequency digital signal, while maintaining phase information[3]. In the subsequent digital processing, the signal is IQ detected against an LO derived also from the RF reference. Compared to direct IQ (baseband) sampling, this has the advantage of reduced harmonic distortion [SG22, chap. 5.2.1]. Selecting higher harmonics of the optical pulse train leads to higher sensitivity, but at the cost of ambiguity in the zero-phase point, the so-called *bucket problem*. Directly locking to a higher harmonic, the phase of the MLO will be stable with respect to the MO, but at a random discrete offset. This poses a significant problem for machine operation, as the timing sensitive accelerator components (re)synchronized to the MLO then operate in a different, suboptimal working point. A common solution is to first engage the PLL on the non-ambiguous fundamental frequency and then switch to detection based on a higher harmonic for greater sensitivity, or phase-unwrapping the higher harmonic signal based on the fundamental. The latter also increases the dynamic range of the detection via higher harmonic, but the finite bit-precision of the digital implementation needs to be

[3]Still, the sampled IF is required to be within the analog bandwidth of the ADC

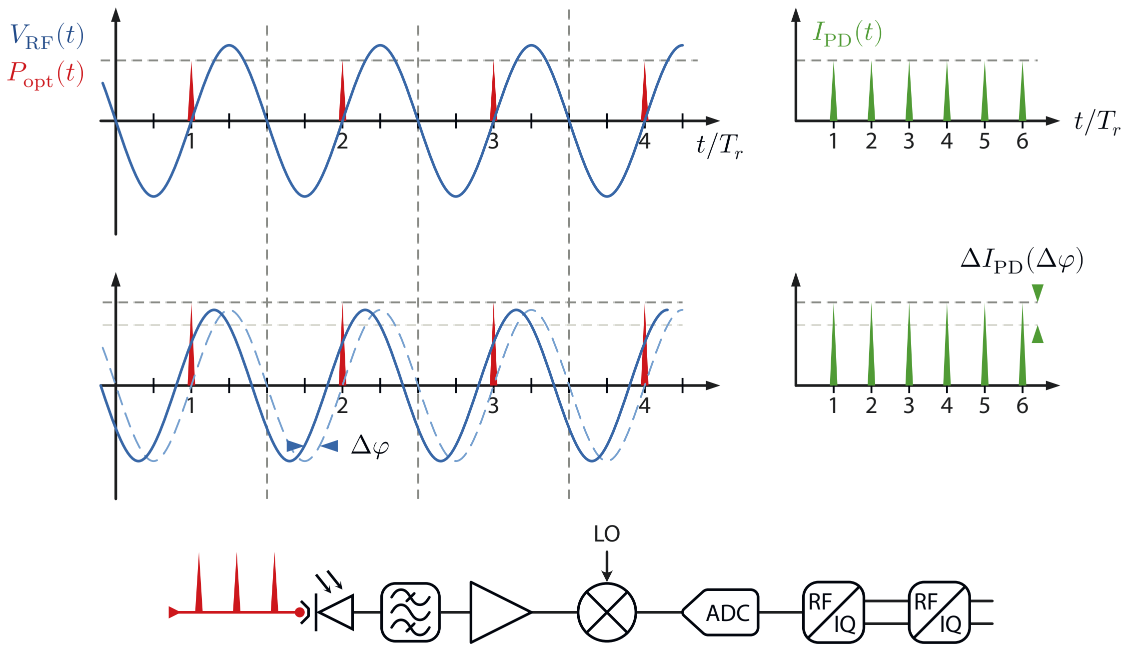

Figure 5.16: Mach-Zehnder modulator detection scheme. Laser pulse train is unmodulated if in-phase with the RF signal zero-crossings and otherwise amplitude-modulated proportionally to the phase error. This can be photodetected a IQ demodulated. Figure adapted from T. Lamb (CC-BY-SA 4.0) [Lam17].

taken into consideration [Heu18]. The heterodyne detection technique is limited in performance by the power of the reference and photodetected RF signals (determining the signal-to-noise ratio (SNR)), AM-PM conversion in the photodiode [Tit17], ADC jitter and aforementioned bit-resolution, as well as drift in the RF electronics (e.g. due to humidity) [Jan+15]. To address the electronic drift, a phase drift calibration method can be employed, see [Jan+15]. The detection noise floor of the 7$^{\mathrm{th}}$-harmonic ($1516.\overline{66}$ MHz) heterodyne detection implemented for the MLO in the DESY LbSync system is at approx. -133 dBc/Hz (40 as/$\sqrt{\mathrm{Hz}}$). More problematic however is the high flicker (i.e. $1/f$) noise rising out of the noise floor at $\sim$3 kHz, which is indistinguishable from true phase error between MLL and RF reference, and added to the MLL's output by the PLL controller.

DESY Mach-Zehnder Modulator Scheme Because of the poor long term drift stability of the heterodyne detection, and because of performance limitations due to the underlying direct conversion, a highly advanced Mach-Zehnder modulator (MZM)-based optical-microwave phase detection has been developed at DESY [Lam+11; Tit17; Lam17]. The principle of operation of the basic scheme is depicted in Figure 5.16. The laser pulses from the controlled laser oscillator are modulated in a quadrature-biased MZM by the oscillating voltage from the RF reference, leading to an amplitude modulation of the optical pulse train in case of a phase mismatch. While phase and bias voltage errors are indistinguishable in basic MZM schemes, advanced configurations can separate both effects, and can additionally be made insensitive to amplitude instabilities of the laser pulse train [Lam17]. These variants are based on dual-output MZMs in combination with a balanced readout or pulse interleaving and IQ-detection. Single-output MZMs can also be used, however these configurations face challenges from drift in the optical alignment and AM-sensitivity in the bias error detection [ESS18a; Lam17].

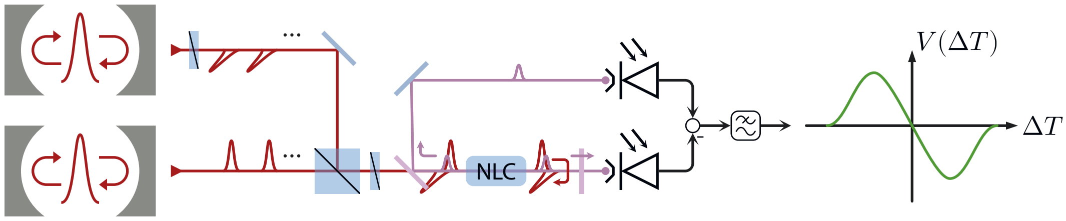

Figure 5.17: Balanced optical cross-correlation of two MLLs. Superposition of the orthogonally polarized pulses in the nonlinear crystal (NLC) causes second-harmonic generation (SHG), which is isolated using dichroic mirrors. Passing the NLC twice induces a group delay in the orthogonal pulses, leading to amplitude-insensitive timing detection.

Compared to direct conversion and heterodyne detection, the MZM-based phase detection suffers from significantly less AM-PM conversion, see [Tit17], where the AM-PM conversion coefficient of a simple MZM-based optical-microwave phase detector was shown to be roughly an order of magnitude smaller compared to direct conversion using a gallium arsenide photodiode. It also shows far superior long-term drift stability, showing out-of-loop drift deviation (averaged over 600 s) between two identical MZM detectors of less than 2 fs peak-peak over 7 days [SCS24]. The detection noise floor of the MZM phase detector used at European XFEL and FLASH to synchronize the MLO to the MO has not been thoroughly characterized to this date, but can be roughly estimated from the high-offset-frequency floor to be less than -155 dBc/Hz (3 as/$\sqrt{\text{Hz}}$). The aforementioned out-of-loop drift also certifies excellent flicker noise properties.

Other Optical-Microwave Schemes Parallel to the developments at DESY, several other optical-microwave detection schemes have been developed and published in the literature, driven by the need to overcome limitations of direct conversion based methods. Comparisons of selected schemes are found in [Ahn+22] and [Jae+20]. Many are based on actively or intrinsically biased Sangac loops. While some methods are also based on balanced phase-error dependent amplitude modulation via MZMs [KKP04; ESS18a; Sha+23], a general concern with the proposed schemes is the long-term stability due to either missing bias stabilization or sensitivity to environmental disturbances.

Laser-to-Laser Synchronization

A basic solution to synchronize two MLLs is to detect a suitable harmonic of the laser's repetition rates and perform standard RF synchronization, just as in the heterodyne detection scheme described in the previous section. However, this approach is severely limited in jitter performance and drift stability compared to fully-optical detection of the relative timing of the two femtosecond laser pulse trains, which can be achieved by *optical cross-correlation* [Sch+03]. In this scheme, the two laser pulse trains are orthogonally polarized and spatially overlapped in a special phase-matched crystal[4]. If the pulses are also temporally superpositioned in the crystal, its nonlinear properties cause generation of second-harmonic generation (SHG) or sum-frequency generation (SFG) light, which can be filtered and detected. The amount of light generated this way is maximized if subsequent pulses from both lasers continuously perfectly overlap in the crystal, thus implying synchronization of the two laser oscillators. In this simple configuration, the scheme is sensitive to intensity fluctuations in the laser pulse trains. A balanced configuration as depicted in Figure 5.17 is hence preferable. Here, the pulses pass the crystal

[4]Examples are periodically poled potassium titanyl phosphate (PPKTP) and beta barium borate (BBO).

twice and one polarization axis is group-delayed such that on each pass, the pulses partially overlap. The working point of the detector is chosen such that both passes of the crystal produce the same amount of SHG/SFG light. In this configuration, intensity fluctuations of either laser pulse train equally affect both forward and backward signals and thus cancel out in the difference signal. Relative timing changes however inversely affect forward and backwards signal, leading to a characteristic S-curve detector signal with reasonably linear response around the working point given by the centre zero-crossing. Perfect separation of amplitude and phase noise is however only achieved if the device is properly balanced, which is difficult to do perfectly in practice. Residual AM-PM conversion thus needs to be taken into consideration. Other factors limiting the precision of this detection method are optical pulse power, width and shape, photodiode shot noise, amplifier additive noise, ADC noise and drift, and wavelength compatibility of the two laser oscillators. The here presented balanced optical cross-correlator (BOXC) scheme is used in the DESY LbSync system extensively, in particular to synchronize to the subsidiary laser oscillator (SLO) to the MLO and the various pump-probe laser (PPL) oscillators to the SLO.

With the LAM, a new variant of this BOXC-based timing detector dedicated to the stabilization of the PPL arrival time at the experimental interaction point is currently under development at DESY [Sch+24]. The purpose of this device is to detect residual drift accumulated in the PPL system and beam path to the experimental interaction point and eliminate it via feedback to the synchronized PPL seed laser oscillator. Compared to the existing (two-colour) BOXC detectors, the LAM has to address multiple additional challenges, including tunable wavelength compatibility, power variation and burst operation. The former leads to a much more involved optical design. The latter means that, because FLASH and European XFEL and therefore the PPL delivery to the interaction point operate in burst mode with less than 1 % duty cycle, the SFG signal cannot be detected in baseband, as the SFG conversion efficiency is simply too low to generate sufficient average power. Pulse-to-pulse detection, integrating the power generated by each pair of laser pulses on forward and backward passes, respectively, is therefore necessary. Even though the initial goal is just drift compensation, the pulse-to-pulse detection also opens up possibilities for fast intra-burst arrival time stabilization.

Challenges in Synchronizing MLLs

We briefly discuss the intricacies of synchronizing MLLs, particularly factors limiting the bandwidth of synchronization, which were learned from experience in dealing with these systems over the time of this work. To the author's knowledge, virtually all commercially available, tunable MLLs use intra-cavity mirrors mounted on piezo crystals to actuate the cavity length and thus the laser's repetition rate with high bandwidth. Broader tuning ranges are achieved via an additional mirror mounted to a larger piezo with greater travel range or by using a cavity heating system, which is usually present anyway to stabilize the cavity temperature against environmental temperature variations. Fortunately, both the piezo expansion and the repetition rate depend affinely on the applied voltage and the cavity length, respectively, leading to overall linear system dynamics. Piezo crystals feature strong resonant frequencies [van28], which is for example exploited explicitly in quartz oscillators, but could lead to significant deterioration of the jitter performance, if unintentionally excited by the controller output. While this used to be a primary limiting factor for closed-loop synchronization bandwidth in early MLLs, special mounting of the piezo crystal [e.g. Nak+20b] has since improved the response and shifted all relevant resonances to 100 kHz and beyond. In other cases, higher-order controllers notching out the offending resonant frequencies need to be considered. The capacitive behaviour of the piezo crystal, often with primary serial capacitance of $\sim$100 nF, indicates low resistance and high currents at high frequencies, which could potentially damage the piezo when tuning the controller too aggressively.

Similarly, voltage steps to the piezo cause high mechanical stress and can cause cracks in the crystalline structure. In this case, the piezo can abruptly change its dynamic response depending on the applied bias voltage.

Another challenge arises not from resonances in the piezo's electromechanical response, but from mechanical resonances in the laser cavity or surrounding housing. In some lasers, these were observed as low as 400 Hz, while in more modern systems, values around 60 kHz to 80 kHz appear typical. Unfortunately, these resonances are not as easy to characterize as, say, attaching an impedance analyser to the tuning piezo's electrical contacts. Lower frequency resonances are generally excited not by the (low-power) tuning piezo, but rather by external sources such as motorized mechanics in proximity to the MLL. Here, active noise cancellation (ANC) techniques may be utilized to suppress the deterministic disturbance, preliminary research in this direction was conducted in [Ryb+17; Beh21]. Higher frequency resonances on the other hand have been observed to be excited by the piezo actuator itself and thus again would require controller-based notching or special engineering of the laser cavity. We conclude that synchronization of MLLs up to 100 kHz and beyond should be possible if the piezo and housing resonances are sufficiently suppressed, and the error is overall small enough so that the piezo is not overloaded.

5.2.4 Stabilized Optical Fibre Timing Links

The pulsed optical fibre timing links developed at DESY and deployed at FLASH and European XFEL are based on the auto-correlation principle described in Section 3.3.1 and subsequently described in more detail. Literature references are given in [Kim+08; Sch+15]. A reference laser pulse train is coupled into an optical fibre terminated at the synchronization client by a partially-reflective mirror, returning a fraction of the light back to the origin. The remaining optical power is transmitted for client use. Back at the origin, the reflected light pulses are directed to a BOXC, cf. Section 5.2.3, where the round-trip timing is detected against a fraction of source light separated before in-coupling to the optical fibre. Stabilizing the timing of the reflected laser pulses leads to a standing-wave-like fixing of the round-trip path length and thus compensation of environmental disturbances acting on the transmission media. For fast actuation, a few dozen meters of the optical fibre are wound around and glued to a tube made from piezoelectric material which strains the fibre when a voltage is applied to it.

While straight forward in theory, pulsed optical fibre timing links as described above are challenging to implement in practice due to several complications arising from optical effects occurring in the optical amplifiers required to overcome losses in the long optical fibres and in the optical fibre itself that deteriorate the laser pulse shape to the point where it becomes infeasible to be used for optical cross-correlation. These effects include chromatic dispersion, polarization mode dispersion, polarization axis cross-coupling and higher-order effects from the erbium-doped fibre amplifier (EDFA). For this reason, the pulsed optical fibre timing links at FLASH and European XFEL are made from polarization maintaining (PM) optical fibre and feature a special appendix scheme ensuring that these effects are mitigated. Preliminary details on the scheme are found in [Heu18] and references therein.

Another critical engineering problem of a LbSync system with many links is the distribution of the optical light from the laser oscillator to the individual links. Because this section of the system is inherently unstabilized, environmental disturbances need to be minimized. Indeed, when aiming for femtosecond precision, an optical path length change of merely a third of a micrometre, a percentile of the width of an average human hair, foils any hope of success. Assuming an average optical path length from laser to link BOXC of 3 m, this is a relative change of 10^{-7}. Besides mechanical vibrations and microphonic influences, timing drift is also caused by changes in the refractive index of air due to temperature, air pressure, humidity and CO_2 content. Linearizing the modified Edlén or Ciddor equation from [BD94; Cid96; SJZ00] at 1550 nm, 23 °C, 1 atm, 30 % relative humidity and 450 µmol/mol CO_2, one yields the approximate coefficients 3.06 fs/(m K), −0.88 fs/(m hPa), 0.04 fs/(m %) and −0.5 as/(m µmol mol). At European XFEL, this problem is addressed by stabilizing the the synchronization laser laboratory to <0.1 K RMS and <1 % relative humidity RMS. Because air pressure cannot be controlled with reasonable effort, the optical path lengths from the first beam split to the individual BOXCs should be kept as close as possible to mitigate non-common-mode drifts. The importance of controlling these environmental influences however depends strongly on the long-term stability requirements of the synchronization clients.

5.2.5 Laser-to-RF Re-Synchronization

Besides synchronization of lasers, as mentioned in Section 5.2.1, the LbSync system also provides a drift-stable optical timing reference to the LINAC RF stations of FLASH and European XFEL. Since the RF stations cannot directly utilize the optical signal, its phase information needs to be first transferred back to the microwave domain. Three options are shortly introduced and compared. Direct conversion by a PFA chain is generally not yet feasible because of the low RF power that can be extracted from the optical signal, mostly due to conversion efficiencies of the photodiode and excess phase noise added by high-gain RF amplifiers. In the future, modified uni-traveling-carrier photodiodes (MUTC-PDs) may offer sufficient conversion efficiencies to make this approach feasible, if the optical signal from the link has enough optical power and is sufficiently pure. The second approach involves a local VCO that is phase-locked to the optical reference using any of the optical-microwave phase detection schemes presented in Section 5.2.3. This however also is generally impractical due to the amount and cost of local oscillators needed, which would generally not achieve the excellent phase noise characteristics of the facility MO. The third option involves transmitting the MO signal to the RF station via uncompensated (i.e. drifting) coaxial cables, detecting the incurred drift via optical-microwave phase detection and correcting the phase of the RF signal using a voltage-controlled phase-shifter. This solution is currently implemented at FLASH and European XFEL with the DESY developed optical reference module (REFM-OPT) [Lam17]. One original key motivation for this solution was that the accelerators could still operate using the drifting RF reference, if the optical reference was interrupted due to system failure. To the author's knowledge, this fall-back mode of operation was however never used or needed at European XFEL for prolonged periods of time.

The REFM-OPT features a dual-output MZM-based phase detection scheme with intrinsic insensitivity to amplitude fluctuations on the optical reference. It is engineered to mitigate mechanical vibrations entering the optical compartment. Even though the phase-shifter also marginally deteriorates the phase noise qualities of the RF signal, the excellent high-offset-frequency phase noise of the MO is mostly maintained. The feedback loop commonly achieves a locking bandwidth of approximately 3 kHz to 5 kHz [Lam+24], which is sufficient to eliminate the majority of the residual phase error incurred on the coaxial cables. The output power of the REFM-OPT is additionally feedback stabilized via an amplifier and a voltage controlled attenuator at the RF input. It should be noted that all RF components are located before the final splitter from which one output is fed to the opto-electric phase detector. As such, all components are in-loop and their drift is compensated, with the exception for any differential output drift of the final splitter, which has been thoroughly characterized and mitigated by in-rack temperature stabilization. Detailed stability performance metrics are reported in [Lam+24].

At European XFEL, seven REFM-OPTs are in use. The devices installed in the L3 section of the accelerator supply multiple RF stations each and are daisy-chained to save MO outputs and cabling. This can be afforded because the phase noise requirements in the L3 RF stations are much more relaxed compared to the stations situated before the bunch compressors.

5.2.6 MicroTCA.4: Digitization and Control Electronics Platform

The DESY LbSync system is to a large extent build on an electronic hardware platform following the MicroTCA.4 standard [Sch+13]. The various advanced mezzanine cards (AMCs) and Rear Transition Modules (RTM) cards fulfil various tasks including implementing custom control and signal processing algorithms via FPGAs, analog down-conversion of RF signals, digitization of analog RF and DC signals, outputting analog control and biasing voltages, driving of piezo actuators via high-voltage amplifiers, moving stepper motors, interfacing vendor specific hardware via low-level protocols, and high level control of the algorithms and hardware platform from a Unix computer system. The latter also enables integration of the system into the accelerators control system, which in the case of FLASH and European XFEL is DOOCS [Wil+17]. For more information on the used AMC and RTM cards and their purposes, we refer the reader to references [Fel+15; Heu18].

The FPGA firmware developed for laser synchronization features multichannel digital IQ demodulation, two 6$^{\text{th}}$-order single input, single output (SISO) state-space controllers with 24-bit fixed-point arithmetic, and a DMA area for PCIe data readout of up to 16 channels with buffers fitting 2^{15} samples. In standard setups with 1.3 GHz master frequency, the ADCs are clocked with $f_{\text{ADC}} = \frac{1.3\,\text{GHz}}{16} = 81.25\,\text{MHz}$ and the signal processing algorithms including the controller run with $\frac{f_{\text{ADC}}}{48} \approx 1.7\,\text{MHz}$ update rate. Closed-loop bandwidths of up to $\sim$160 kHz are feasible. For data readout, the signals are downsampled so that the buffers capture approximately one machine cycle of 100 ms, leading to a rate of $\frac{f_{\text{ADC}}}{48\cdot 16\cdot 5} \approx 338.54\,\text{kHz}$.

The link stabilization firmware supports four links per set of MicroTCA.4 cards and has since [Heu18] been replaced with a pure VHDL implementation based on the DESY Firmware Framework [But+15] instead of Rapid-X. The overall functionality of the firmware however was not changed. It now features differential input stages with IIR filtering and downsampling and 18-bit PI controllers with increased resolution on the integrator state compared to the previous state-space controller implementation, resulting in a smaller overall steady-state error due to numerical underflow. For the future, a 3$^{\text{rd}}$-order high-resolution state-space controller may be implemented to suppress the first resonance of the fibre-strecher piezo at $\sim$18 kHz. The ADCs are clocked at 5 MHz, then decimated to 1 MHz which is also the update rate for the digital controller implementation. This still suffices for closed-loop bandwidths in the order of 100 kHz. Similar to the laser synchronization firmware, the data is downsampled to $\sim$326.8 kHz for PCIe readout.

An important aspect of digital control implementations is the digital processing latency, which is usually the dominant contributor to the overall loop delay. Using the DAC on the PZT4 piezo driver RTM, the overall processing latency was observed to be jittering in the range 10 µs to 12 µs for a laser-lock setup clocked at 81.25 MHz and 13 µs to 15 µs for a link-lock setup. The variability is explained by the random timing of the change in controller output in relation to the SPI DAC communication cycle. The worst-case latency for the laser-lock setup corresponds to almost 45° phase loss at 10 kHz, significantly reducing stability margins and limiting the achievable closed-loop bandwidth. Utilizing the front-panel DAC on the SIS8300-L2 AMC, the latency is reduced to worst-case 2 µs, corresponding to just 7° at 10 kHz. Even though the link stabilization application would also benefit from a reduced loop latency, e.g. utilizing the AD84 RTM DACs, the gain margin increase would be minimal.

5.2.7 Modelling of Complex Distributed Systems

In the past, performance improvements in large experimental science facilities such as European XFEL were primarily achieved on the subsystem level. On a macro scale, coarse performance indices were defined for each subsystem which would ensure satisfactory overall performance of the machine, however the technological challenges were limited to individual engineering problems solved by small teams of experts. As further improvement of the subsystems becomes increasingly challenging as the performance approaches physical limits of thermal noise, heat dissipation, etc., the joint optimization of multiple subsystems becomes an increasingly important tool to achieve better machine performance [Qia22]. This has created a need for suitable models of large complex cyber-physical systems, which are challenging to create and maintain because they generally are cross-domain, encompass many model parameters and require a thorough understanding of all involved subsystems. For example, while a team of experts can fairly easily model and optimize one particle accelerator RF station, a global model of the European XFEL would have to maintain parameters for 24 active RF stations and also take into consideration coupled dynamics via the beam consecutively passing through all 784 cavities[5]. Depending on the purpose of the model, additional subsystems and coupled dynamics may need to be taken into consideration, e.g. magnets, lasers, diagnostic devices, geophysics, etc.

One of goal of this work is the creation of a unifying model for the entire European XFEL LbSync system and its validation using data taken from live system operation. Though ultimately requiring multiple iterations, the modelling process can be summarized as follows:

1. *Literature research.* Based on the central references [Heu18; Sch+15; Löh09; Din18] a literature survey was conducted, collecting references relevant to the synchronization at European XFEL and FLASH in general or dealing with particular subsystems or components of the LbSync system.

2. *Knowledge graph.* Based on references discussing the high-level topology of the LbSync system, a (multi-layer) block diagram of the system was created, the blocks of which contained references to relevant publications.

3. *Abstraction.* Across all gathered dynamic subsystem models, we identified common properties such as structure and quantities of interest, eliminating aspects of the models not relevant to the timing stability. This for example allows laser-to-RF PLLs to be modelled just like laser-to-laser or pure RF PLLs, thus reducing overhead.

4. *Identification of coupled dynamics.* From the abstracted subsystem models, we inferred a joint model for the quantities of interest relevant for subsystem coupling, i.e. signals passed from one subsystem to another.

5. *Time and frequency domain modelling.* Due to the block diagram structure, the model could directly be replicated in MATLAB Simulink [The24a], yielding a time domain model of the system. However, for control system design in particular, a frequency domain representation in terms of transfer functions (or equivalent state-space models) is more useful. This was developed using MATLAB's control system and robust control system toolboxes [The24b; The24c] based on the subsystem's dynamic equations.

6. *Device characterization.* The coefficients of the dynamic subsystem models needed to be identified from the actual systems installed at European XFEL. Though some parameters (e.g. amplifier gains) can be assumed to be virtually identical across all devices from a single production batch

[5] Including the injector accelerator module and third-harmonic section.

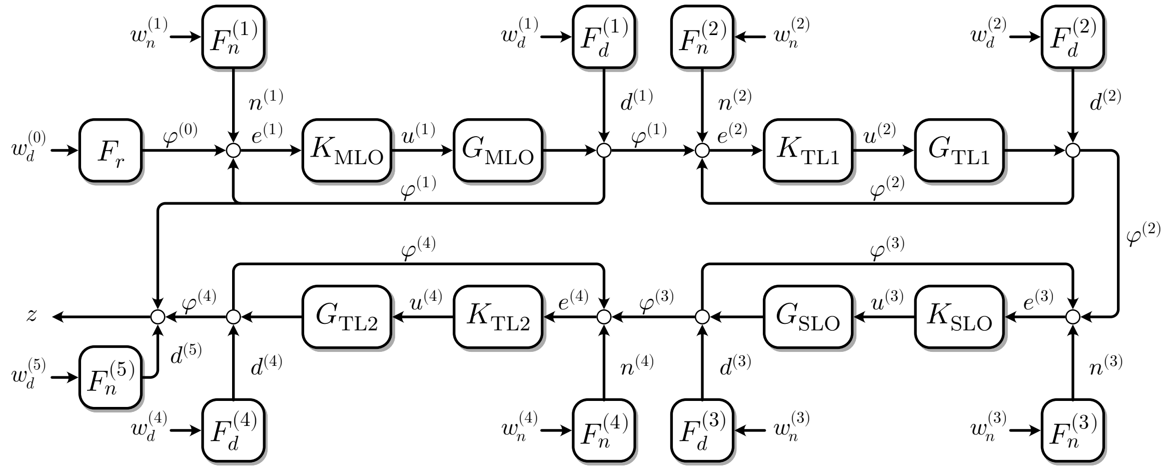

Figure 5.18: Block diagram of the round-trip experiment at European XFEL. (0) through (4) abbreviate the MO, MLO, SLO link (TL1), SLO and return link (TL2) subsystems, respectively. The Round-Trip BOXC is abbreviated by (5).

or model revision, others, like phase noise profiles for laser oscillators, are highly dependent on environmental conditions and device health, deteriorating with time. Because of this, a stringent measurement protocol was developed, giving context to each measurement and ensuring that only measurements from a coherent system state are used to build the overall model.

7. *Measurement post-processing.* Where possible, necessary post-processing of measurement data was done using MATLAB scripts, ensuring that the processing is reliable and reproducible. For control system design, finite-order approximations of certain measurements are needed. Often these are best fitted by hand, as expert knowledge dictates which features of a measurement are real or artificial, i.e. need to be incorporated or omitted in the finite-order approximation.

8. *Model parameter tree.* We organized the properties that were obtained from individual measurements into a parameter tree replicating the topology of the LbSync system. The model coefficients are then derived from the parameter tree using dedicated model building code. This also allows merging multiple parameter trees from several measurement campaigns, using only the latest identified values.

Due to the significant workload attached to the identification of all required system parameters, and because of restricted access to the facility when in operation, only a subset of the subsystems included in the LbSync system at European XFEL was identified so far. Most importantly, the current version of the frequency-domain model covers the MO, MLO, SLO, the MLL oscillator at the soft X-ray port (SXP) instrument, the interconnecting pulsed optical fibre timing links as well as a return link used for performance evaluation (cf. Section 5.2.8). This allows to model an entire chain from MO to PPL, which is the primary concern for experimental scientists at the moment. In a next step, the model is meant to be augmented with a longitudinal timing stability model of the accelerator to accurately predict the arrival time jitter at the experiment.

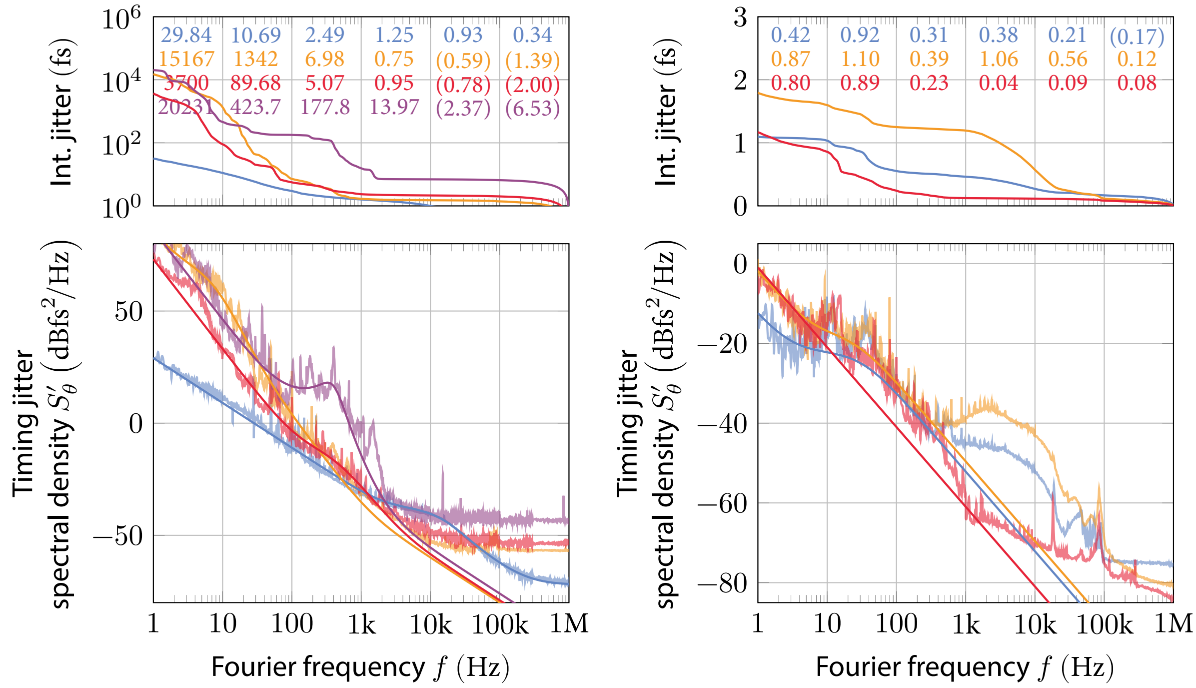

(a) The RF MO (——), MLO (——), SLO (——), and SXP instrument MLL oscillator (——).

(b) The SLO link (——), round-trip experiment link (——), and the SXP instrument link (——).

Figure 5.19: Low-order phase and timing noise models for relevant LbSync subsystems at European XFEL. The shaded lines indicate the absolute noise measurements used for fitting. Coloured numbers indicate per-decade integrated jitter values. Values in parentheses are limited by the measurement setup. Integrated jitter computed without spurs.

5.2.8 Model Validation and Round-Trip Performance Evaluation

In order to quantify the synchronization performance and long-term drift stability of the LbSync system, a dedicated round-trip experiment was set up and evaluated at European XFEL. As depicted in Figure 5.10, the reference signal supplied by the MLO via a long optical fibre timing link to the synchronization subdistribution laboratory in the experimental hall, approximately 3.5 km away from the MO, is returned to the MLO in the main synchronization laboratory via the SLO and another long timing link. By detecting the timing between the returned signal and the MLO using a BOXC, the synchronization performance over a total close to 8 km of optical fibre and an intermediate "repeater" oscillator can be upper bounded[6]. We will also use this data, along with data collected from the individual subsystems, to validate the LbSync system model described in Section 5.2.7 and further for control system design in Section 5.3. The simplified and control systems oriented block diagram of the experiment is shown in Figure 5.18.

[6]We claim that this measurement upper bounds the performance, because the experiment design suffers distinctively from the delay line effect (cf. Section 3.3.1), especially when the MLO's excellent absolute phase noise at high offset frequencies is artificially deteriorated by tightly locking it to the MO. For the synchronization chain topology in the accelerator, the delay line effect is however less dominant.

	MLO	SLO	SXP MLL
Centre Repetition Rate (MHz)	216.7	216.7	54.17
Repetition Rate Tuning Gain (Hz/V)	9.85	4.95	0.42
Driver Gain (V/V)	25.18	25.18	41.25
Driver Bandwidth (kHz)	150	150	150
Loop Latency (µs)	2	2	9
Proportional Feedback Gain	$7.26 \cdot 10^{-6}$	$181 \cdot 10^{-6}$	$441 \cdot 10^{-6}$
Integral Feedback Gain	$23.9 \cdot 10^{-3}$	$606 \cdot 10^{-3}$	$112 \cdot 10^{-3}$

Table 5.2: Parameters of selected MLL oscillators from the European XFEL LbSync system.

	SLO	SLO Return	SXP
Nominal Optical Round-Trip Time (µs)	37.24	39.20	1.76
Fibre Stretcher Optical Path Delay (fs/V)	36.42	45.58	39.77
Driver Gain (V/V)	83	83	83
Driver Bandwidth (kHz)	1.0	1.0	1.0
Piezo (Anti-)Resonance Frequency (kHz)	18.54 ± 0.24	18.54 ± 0.24	18.54 ± 0.24
Piezo (Anti-)Resonance Damping	$\sim 2 \cdot 10^{-3}$	$\sim 2 \cdot 10^{-3}$	$\sim 2 \cdot 10^{-3}$
Loop Latency (µs)	14	14	14
Proportional Feedback Gain	$26.4 \cdot 10^{-6}$	$40.3 \cdot 10^{-6}$	$80.0 \cdot 10^{-6}$
Integral Feedback Gain	2.43	0.18	4.00

Table 5.3: Parameters of selected pulsed optical fibre timing links from the European XFEL LbSync system.

The model coefficients are derived from several individually measured system parameters, which are summarized in Tables 5.2 and 5.3. Measured absolute phase noise power spectral densities are depicted in Figure 5.19 together with the fitted responses of the low-order colouring filters used later for synthesis.

Remark 5.2. The small feedback gains might be understood as an indication of improper model scaling. However, this is not the case, as the closed-loop residual errors indeed track in the few or even sub-femtoseconds. The order of magnitude on the gains hence rather indicates "over-actuation" (in the sense of magnitude rather than degrees-of-freedom) of the system. However, one also needs to take into consideration that the digital controller implementation saturates at ± 1 and still needs to be able to catch the PLL during locking and counteract strong disturbance events such as mechanical shocks without saturating. These effects are not taken into consideration for "nominal steady-state" operation, for which the phase noise colouring filters are designed. Without having to account for anomaly events and PLL catching, the piezo driver could be very-well removed from the system.

Sound understanding of the principles and limitations of the employed measurement techniques is crucial to obtain proper system models. In the following, we summarize key limitations and assumptions made for post-processing of measurements in order to arrive at the final model and achieve the desired consistency with the validation data.

1. The absolute phase noise measurements of the laser oscillators at Fourier frequencies beyond a couple of kHz are dominated by the photodiode used for direct conversion (cf. Section 5.2.3) and hence extrapolated beyond 10 kHz, from which point on the MLL timing stability is assumed to be limited by amplified spontaneous emission (ASE) with −20 dB/dec slope [Pas04; KS16].

2. The timing noise of the long pulsed optical fibre timing links at Fourier frequencies beyond $\sim$1 kHz is dominated by the phase stability of the source MLL via the delay line effect (cf. Section 3.3.1) and hence extrapolated as $1/f^2$.

3. Since the exact optical path lengths of the pulsed optical fibre timing links are unknown, they were estimated based on available documentation and construction schematics and possibly fine-tuned by aligning the elimination notches in the error signal spectra.

4. All parameters can only be measured or identified up to some certainty. While certainly useful, a full characterization of the respective confidence intervals however is beyond the scope of this work.

We base our validation on the correct prediction of the timing jitter spectral density at the round-trip cross-correlator. As shown in Figure 5.21, there is good qualitative and quantitative agreement between model output (——) and the timing jitter spectral density (——) measured at the BOXC detecting the round-trip timing error. Except towards very low Fourier frequencies, the slope virtually matches exactly, the characteristic elimination ripples due to the delay line effect can be clearly made out in the measurement and are well aligned to the simulation. Quantitatively, the discrepancy between simulation and measurement does not exceed a few dB for the most part, which may in part be attributed to calibration error. This error is linked to the accuracy of the motorized optical delay line used as absolute timing change reference, which we estimate to be better than ± 15 %. The simulation increasingly underestimates the timing jitter spectral density towards lower Fourier frequencies. Cross-referencing the results show in Figure 5.21 with the individual subsystem results in Figure 5.20, we reason that the discrepancy towards very low Fourier frequencies is due to environmental influences not included in the model, primarily temperature and air pressure fluctuations affecting the out-of-loop free-space optical distribution in the laser laboratory, as well as possibly electronic drift in the balanced

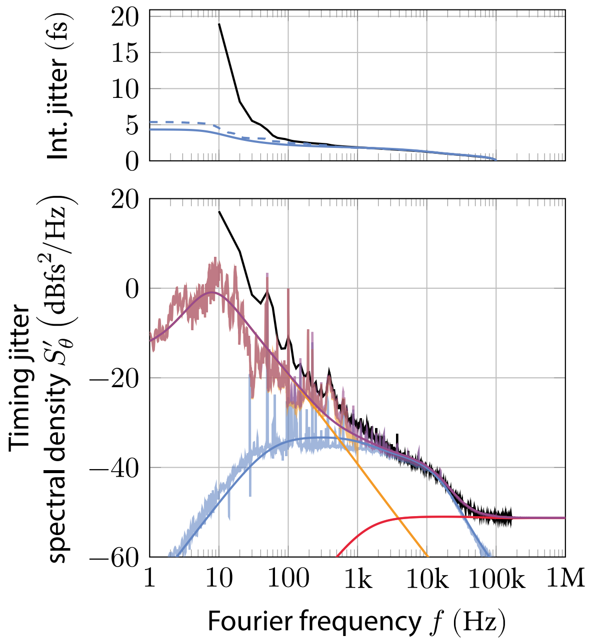

(a) MLO

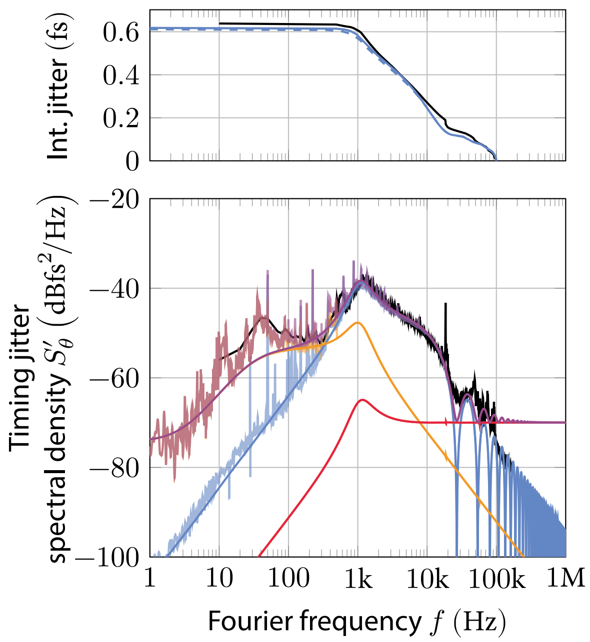

(b) SLO link

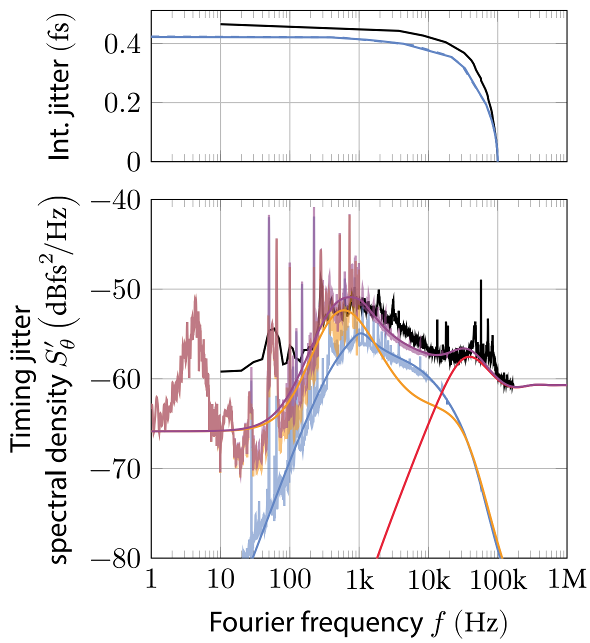

(c) SLO

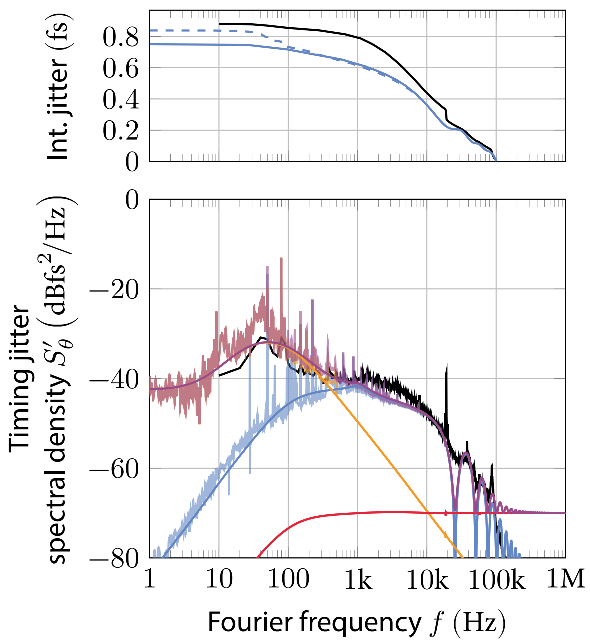

(d) Return link

Figure 5.20: PSDs of the in-loop residual error signals for the four subsystems comprising the LbSync round-trip experiment at European XFEL. The contributions from the reference signal (——), output noise (——) and the measurement noise (——) add up to the total simulated residual error (——) in RSS. An in-loop measurement (——) is used as validation data. The solid lines are computed based on low-order colouring filters, whereas the shaded / dashed lines indicate the measured noise spectra used for fitting, transformed by the closed-loop transfer function.

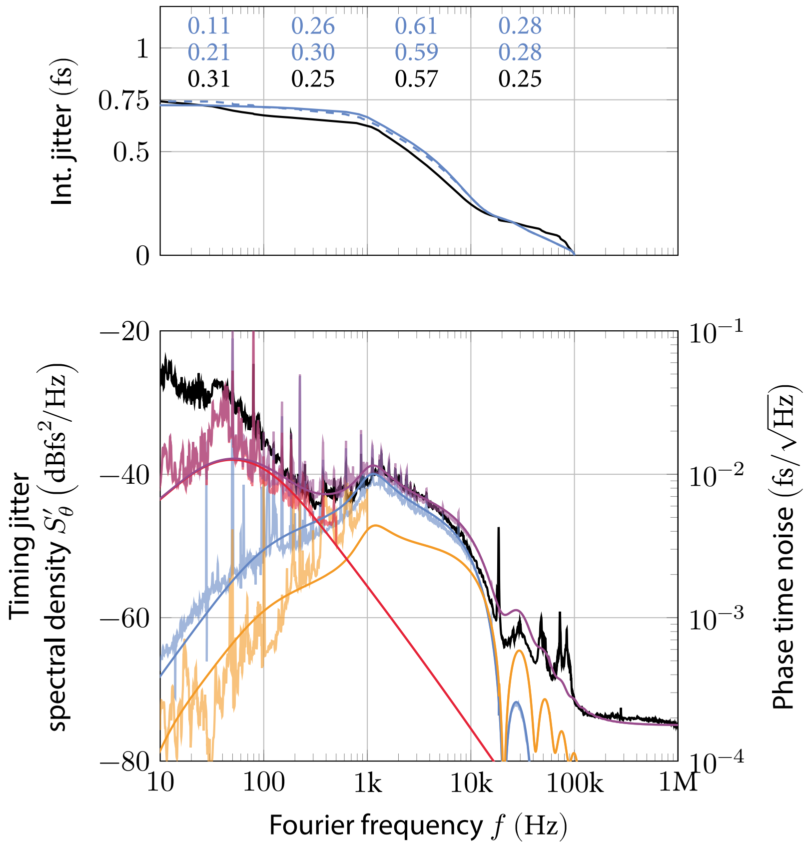

Figure 5.21: Comparison of the timing jitter spectral density measured with the round-trip experiment BOXC (——) and the LbSync system model output (——). The dominant contributions come from the MO (——), the MLO (——) and the return link (——). The shaded lines indicate cleaned-up input noise PSDs, transformed by the closed-loop model transfer function.

photodetectors. Non-stationary behaviour of the relatively weakly suppressed noise on the return link (——) can also not be ruled out, because the measurements were conducted successively rather than simultaneously. AM-PM effects are expected to play a secondary role, due to good balancing of the BOXCs. Using low-order noise-colouring filters, which later are required for synthesis, rather than computing the residual round-trip jitter directly from the measured absolute phase noise spectra yields a virtual identical result, which is also visible from the dashed integrated jitter curve.

Based on the simulation, the MO (——) is made out as a main contributing factor to the round-trip timing jitter, adding ~ 0.6 fs in-between 1 kHz to 10 kHz. Here, the delay line effect puts significant demands on the upstream reference with respect to absolute phase noise, if the overall performance is to be further improved. The second-largest contribution is found in the 10 Hz to 100 Hz range, where the weakly locked return link (——) and climate instabilities add up to ~ 0.3 fs jitter in RSS. Finally, the decade 10 kHz to 100 kHz is dominated by the detection noise floor of the SLO BOXC (not plotted), due to the high SLO PLL bandwidth, and the delay line ripples of the MLO. Despite many visible spurs in the timing jitter PSD, the integrated jitter curves show that these add no significant RMS timing jitter. In total, the overall performance of the Bayesian optimization tuned system achieves ~ 750 as 10 Hz to 1 MHz integrated jitter over ~ 8 km of signal transport. To the author's knowledge, no comparable sub-femtosecond synchronization performance has so far been demonstrated at any other facility in operation. While the round-trip experiment provides a valuable stand-alone characterization of the LbSync system, it is not a perfect timing stability indicator for user operation in the facility for two main reasons. One is the aforementioned delay line effect, which, due to the round-trip configuration, is much more pronounced in this experiment than in any point-to-point configuration within the European XFEL. Indeed, the delay line effect scales at most with the relative distance of any two nodes to their common ancestor in the timing synchronization tree. We say "at most" because an intermediate oscillator like the SLO causes jitter to be uncorrelated beyond the locking bandwidth, thus the phase noise just adds up, rather than being amplified. The other reason is that the figure of merit for the experiments is actually the relative arrival time jitter of the PPL pulse and the FEL pulse at the interaction point. The former is additionally impaired by the PPL system, including high-power optical amplifiers, compressors and distribution. Deployment of LAMs (cf. Section 5.2.3) at multiple experiments is planned for 2025 at both European XFEL and FLASH, allowing for slow drift compensation. A fast feedback is also conceivable in the future. The FEL pulse jitter is essentially dependent on the longitudinal stability of the electron bunches in the LINAC, which depends on the RF feedbacks as well as other longitudinal stability or energy feedbacks in the accelerator. While the relative arrival time jitter at the interaction point can be measured with dedicated photon arrival time monitor (PAM) setups, the required devices are still under development [LG17; Czw+21] and parasitic use during other experiments is still a major challenge. Nevertheless, in 2022, during a dedicated campaign at European XFEL, PAM data was already used for online black-box optimization of the LbSync feedback gains [Lüb22].

5.3 Controller Design Study for the LbSync System at European XFEL

Having validated the derived system model for the LbSync system at European XFEL, we now turn to a full controller design study using the robust distributed controller synthesis methods developed in Chapter 4, demonstrating their applicability to large-scale systems. Because the SI methods only solve convex approximations of the original problem and the nonsmooth optimization methods provide a local rather than global certificate of optimality, we require suitable benchmarks to evaluate the performance gap of the synthesized suboptimal solutions. These will be introduced in Sections 5.3.1 and 5.3.2. Obtaining a global certificate of optimality is computationally infeasible for the discussed eight- and ten-dimensional parameter spaces, which is why we use a reference black-box tuning algorithm based on Bayesian optimization (BO) as reasonable estimate of the global optimum and a nominal, fully centralized controller, which admits a convex solution, as lower performance bound.

Two study cases are considered. The first is identical to the round-trip experiment discussed in Section 5.2.8, allowing for real-world black-box tuning and validation of the synthesized controllers (at least in a nominal setting). The second study case considers the configuration primarily relevant for facility operation, namely the optimal synchronization of the PPL laser to the MO and by extension, the FEL pulses at the experiment. The topology is similar to the round-trip experiment shown in Figure 5.18, but has another MLL PLL after the second link and the performance variable z is computed as the difference of the output phase of that PPL oscillator to the MO. Here, no real-world validation data is available, as the required PAM experiment is only available during dedicated facility studies. Comparison to PAM data is planned in the near future. For reasons that will become clear in the subsequent exposition, the primary figure of merit, besides any (robust) upper performance bounds derived directly or indirectly from the optimization variables of the respective synthesis problems, will be the nominal timing jitter between the MO and the respective endpoint integrated in the range 1 Hz to 1 MHz, closely approximating the closed-loop $\mathbf{H}_2$-norm of the system.

5.3.1 Benchmark: Black-Box Tuning via Safe Bayesian Optimization

In recent years, BO has emerged as one of the most prominent and promising black-block tuning algorithms for computationally expensive optimization problems [Wan+23], especially in accelerator physics [Rou+24]. BO revolves around the minimization of an unknown cost function $f \colon \mathbf{X} \rightarrow \mathbf{R}$ on a (generally compact) d-dimensional set of inputs $\mathbf{X} \subset \mathbf{R}^d$, called tuning parameters. A Gaussian process model is employed to model the knowledge and uncertainty about the cost function based on experimental evaluations of the cost function at discrete points, a preliminary mean function and a covariance function (or *kernel function*). With each new iteration, a new sample of the cost function is obtained and posterior distributions for yet uncertain cost function points can be calculated. Any hyperparameters for the mean and kernel function can also be updated based on maximum likelihood or maximum a-posteriori methods. The algorithm then proceeds in an exploration and exploitation fashion, iteratively updating the Gaussian process model and evaluating the so-called acquisition function given the latest estimate of the cost function in order to determine the next point to be experimentally evaluated. Noise in the experimental observations is accounted for by assuming additive white Gaussian noise (AWGN) with suitable variance. Typical termination conditions are either the maximum number of iterations or a convergence threshold in the exploitation phase.

For safety critical applications, which we broadly define as tuning problems where certain subsets of the tuning parameter space cause undesirable or "catastrophic" events such as damage to system components due to oscillatory or unstable loop dynamics, special variations of the BO algorithm were proposed in order to satisfy additional constraints which are to be satisfied with certain confidence or based on assumptions regarding the underlying cost function, during the tuning procedure [Lüb+23]. In the case of the LbSync system, this corresponds to a maximum in-loop RMS error signal observed at any subsystem, which is considered to be a precursor to loosing phase-lock and potentially damaging the piezo actuators in the laser cavities. As oscillatory behaviour at any subsystem is also visible downstream, it is generally sufficient to enforce a safety threshold on the RMS round-trip measurement.

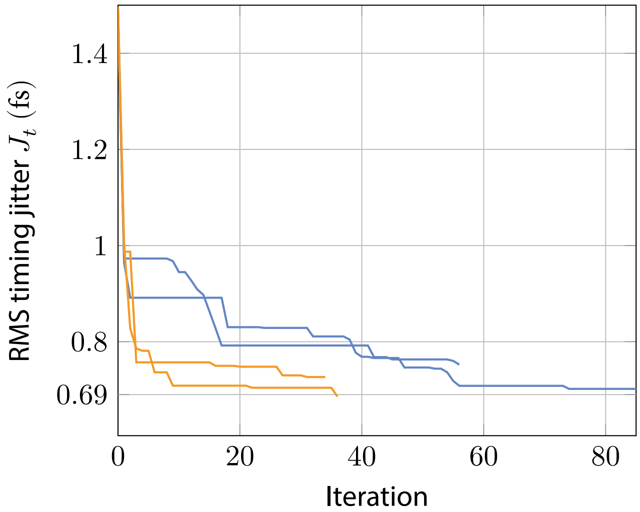

Figure 5.22: Tuning progress of the LbSync round-trip experiment optimization using single-task (Section ——) and multi-task (——) BO. The RMS timing jitter is computed over a 100 ms window.

The results presented next were kindly provided by Jannis Lübsen, who developed and implemented the MoSaOpt and SaMSBO tuning algorithms [Lüb+23; LHE24] and performed the tuning on the LbSync round-trip experiment at European XFEL. The support model used for the SaMSBO algorithm was provided by the author and is a preliminary version of the model validated in Section 5.2.8 and used for synthesis in the subsequent sections. As shown in Figure 5.22, the multi-task algorithm successfully exploited the support model to achieve faster convergence and achieved overall slightly better performance compared to a single-task algorithm (UCB–GP [Sri+12]). The final jitter value of 0.69 fs does not contradict the result from Section 5.2.8, as the RMS jitter value used for BO tuning is calculated over data obtained from the MicroTCA digitization system, operating at a smaller sampling rate and analog bandwidth compared to the baseband analyser that measured the trace in Figure 5.21. We note that one of the single-task optimization runs violated the safety-threshold, which indicates that the algorithm's hyperparameters were not suitably matched to the given unknown cost function, a limitation that multi-task BO (SaMSBO) generally is less sensitive to. In order to see whether the obtained feedback gains could potentially be further improved, we used them to hot-start the systune nonsmooth optimization algorithm operating on a 1^{st}-order Padé approximation of the system model. The result coupled the MLO even more loosely to the MO and showed higher feedback gains for virtually all other subsystems. In simulation, the closed-loop round-trip integrated timing jitter over 1 Hz to 1 MHz could thus be reduced from 750 as to 490 as, an improvement of $\sim$35 %. Regrettably, the result could not be verified on the actual system, as the round-trip experiment had meanwhile been decommissioned. The additional performance improvement was primarily achieved by relaxing the MLO lock even more than was permitted by the hyperparameters of the safe BO algorithm, and given the good agreement of measurement and simulation for the BO tuned system, we are confident that the 490 as seen in simulation are realistically achievable.

5.3.2 Benchmark: Fully Connected Controller Architecture

In contrast to the previous section, the benchmarks presented next are based on a fully connected controller architecture, intended to provide a frame of reference for the performance penalty due to

Model	Synthesis / Tuning method	Additional Constraints	Guaranteed Performance (fs)	Simulated Performance (fs)
Round-Trip Jitter	(Multi-task) safe BO	Decentralized PI feedback	-	0.75
	ε-pert. + ARE	-	-	0.25 (0.29)
	LMI	Damping >0.1 Decay >0.1	46.9	1.74 (1.75)
	NS	16$^{\text{th}}$-order controller	-	0.27 (0.31)
PPL chain	ε-pert. + ARE	-	-	1.55 (1.57)
	LMI	Damping >0.1 Decay >0.1	43	3.02 (3.03)
	NS	16$^{\text{th}}$-order controller	-	1.63 (1.65)

Table 5.4: Benchmark performance results for the LbSync system at European XFEL. Results based on the Padé-approximated plant are given in parentheses.

distributed and partially connected controller architectures. Not imposing any fixed-order constraint on the controller, the synthesis problem is convex but still singular due to a missing penalty on the controller effort and invariant zeros on the imaginary axis in the channel $w \to e$, cf. Section 3.2.2 and Definition 3.6. We thus compare solutions based on an ε-perturbed model and the corresponding ARE solution (cf. Section 3.2.2 and Theorem 3.8) as well as suboptimal controllers obtained by LMI methods (cf. Section 3.2.2 and Theorem 3.10). To ensure finite closed-loop $\mathbf{H}_2$-norm, we impose a first-order bandlimit at 1 MHz on any applicable feedthroughs (primarily the MO phase noise colouring filter) and roll-off any uncontrollable modes in the colouring filters at 1 Hz, so that the $\mathbf{H}_2$-norm closely approximates the integrated jitter over 1 Hz to 1 MHz (cf. Section 3.1.2). To ensure well-posedness of the LMI problem, we explored several pole-region constraints as well as weighting the solution size against the achieved performance bound. Lacking an informed and numerically sensible choice for each involved constraint, we chose not to enforce strictness of the LMIs by means of a definitness tolerance. Rather we ensured feasibility of the solution by sanity checking the closed-loop pole locations.

The full-order synthesis methods produce controllers with possibly many more states than are necessary to attain (close to) optimal performance. Having eight and nine internal delays respectively, the round-trip experiment model and the PPL chain model additionally grow quickly in size with increasing order of the chosen Padé-approximation, as do the synthesized controllers. An insufficient Padé-order poses the risk of an unstable closed-loop, due to insufficient stability margins. On the other hand, a too large model is sensitive to numerical instabilities that can also lead to unstable closed-loop dynamics. Additionally, we have observed that with controllers synthesized even from low-order Padé-approximations, the closed-loop frequency response becomes numerically ill-conditioned, making the numerical evaluation of the closed-loop $\mathbf{H}_2$-performance unreliable. Therefore, we also take into consideration fixed-order solutions obtained from nonsmooth optimization, as a numerically more well-behaved alternative. The results are summarized in Table 5.4. While achieving stability for the given Padé-approximated model, no synthesized full-order controller achieved stability with the time-delay

Model	Synthesis / Tuning method	Additional Constraints	Guaranteed Performance (fs)	Simulated Performance (fs)	Optimality Gap (dB)
Round-Trip Jitter	LMI + SI	Static output feedback	160	2.4	19.65
	LMI + SI	1^{st}-order dynamic output feedback	177	2.4	19.65
	NS	Static output feedback	-	0.57	7.16
	NS	PI feedback	-	0.49	5.85
	NS	2^{nd}-order dynamic output feedback	-	0.47	5.48
PPL chain	LMI + SI	Static output feedback	246	10.3	16.45
	LMI + SI	1^{st}-order dynamic output feedback	246	10.3	16.45
	NS	Static output feedback	-	4.5	9.26
	NS	PI feedback	-	2.4	3.80
	NS	2^{nd}-order dynamic output feedback	-	2.1	2.64

Table 5.5: Decentralized controller synthesis performance results for the LbSync system at European XFEL. Optimality gap w.r.t. best benchmark result.

model when evaluated via time-domain simulation. This is perhaps not surprising, given the lack of guaranteed stability margins from these synthesis methods [Doy78]. Nevertheless, we argue that the benchmarks are still valid, because, by perturbation theory, there exists a "close" controller that achieves stability and virtually the same level of performance for the real system by marginally shifting the closed-loop poles to the stability region, even if we cannot easily compute it. We also note that the benchmarks are solely meant as lower-bounds and that the computed solutions are not fit for real-world implementation, due to very fast dynamics and large high-frequency gain in the controller responses. Attempts to eliminate such fast dynamics by appropriate pole-region constraints in the synthesis LMIs produced overly conservative results, likely due to the large dynamic range of the model coefficients needed to cover six decades in the frequency domain posing a challenge for SDP solvers with finite numerical precision.

5.3.3 Nominal Decentralized Controller Synthesis

We now evaluate the SI structured output feedback synthesis method from Section 4.3 and corresponding solutions from the nonsmooth optimization-based synthesis framework using the previously benchmarked case studies given by the round-trip experiment and PPL synchronization model. For now, we continue to assume that all system parameters are certain and employ finite-order approximations for all time delays. Recall that the latter simplification is required for the SI method to be utilizable. A first-order Padé approximation was used throughout.

Table 5.5 summarizes the synthesis results. While no hard robust stability certificate could be obtained from IQC analysis with dynamic multipliers for constant & bounded time-delays, based on time-domain simulations, all associated closed-loop TDSs were deemed stable for the nominal delay value. Failure to certify stability by multiplier theory does not invalidate this assessment, because the robust stability test is conservative due to finite-order approximations for the dynamic multipliers, the chosen dynamic basis and the fact that the test has to simultaneously be satisfied for all internal delays. Additionally, the test implicitly asserts stability not only for the nominal delay, but all delay values below the given upper bound, which was set to the nominal delay value. Since the closed-loop systems all show severely underdamped modes from the piezo resonances of the optical fibre stretcher, it seems likely that the solution space of this conservative stability test vanishes.

An immediate observation from Table 5.5 is the poor performance achieved by controllers synthesized using the SI framework. Enforcing the high degree of separability on the Lyapunov matrix needed for decentralized control clearly introduces impractical levels of conservatism, even for the static output feedback scenario. Additionally, the state-space realization obtained from the state-output transformation given in Lemma 4.31 is not numerically stable. Indeed, truncating very small coefficients from the transformed system matrices already leads to order-of-magnitude changes in the reported performance bound. We reason that the impact of these numerical properties is made even more severe by the singular problem setting, which by itself is enough to render the controller recovery numerically unstable (due to unboundedness of the central solution, cf. Section 3.2.2), if not constrained to (marginally) suboptimal solutions by means of LMI regions. The additional LMI constraints however increase the conservatism even more, and we have found that, even with relaxed pole-region constraints, the solver often was unable to find feasible solutions at all. We also observed that the SI framework with dynamic output feedback was unable to exploit the additional freedom for the controller to reduce the guaranteed bound. In fact, the controllers input and output matrices evaluate to virtually zero, resulting again in static output feedback. It is not clear whether this is perhaps a direct consequence of the structural constraints on the optimization variables. A direct comparison between 0^{th}-order feedback and 1^{st}-order feedback is not worthwhile because of the numerical sensitivity of the optimization problems. It hence cannot be said which setting will potentially perform better.

In contrast, the nonsmooth optimization based synthesis is highly stable. Virtually all random initial conditions converged to the same respective solutions, hinting that they are indeed globally optimal. The optimization time per iteration was in the order of seconds and thus negligible. Considering the PPL synchronization case study, while switching from static output feedback to PI feedback improved the achieved performance by $46\,\%$, 2^{nd}-order dynamic output feedback only yielded another $14\,\%$ improvement w.r.t. PI feedback. Taking the achieved performance of 2.1 fs as an approximate value for the lower bound that is asymptotically approached with increasing controller order, and comparing to the benchmark results from Table 5.4, we estimate the performance gap due to decentralized feedback to be approximately 0.55 fs ($35.5\,\%$). A higher degree of controller interconnection and the extension from PI to PID feedback are therefore considered worthwhile changes for implementation.

5.3.4 Robust Decentralized Controller Synthesis

The decentralized controllers synthesized above for the nominal plant were heuristically checked for closed-loop stability given the nominal time delay values. For applications such as the LbSync system, which are sensitive to closed-loop instability both due to component and operational downtime costs, hard stability guarantees are welcome, if not required. We therefore synthesize and evaluate robust controllers using the IQC framework and nonsmooth optimization in the frequency domain,

Model	Controller	Delay grouping	Diag. perf. multiplier	# Decision variables	Success rate
Round-Trip Jitter	PI	-	-	274	43 %
		x	-	294	39 %
		-	x	58	58 %
		x	x	78	50 %
	2nd-order	-	-	309	40 %
		x	-	329	35 %
		-	x	93	43 %
		x	x	113	40 %
PPL chain	PI	-	-	394	32 %
		x	-	414	8 %
		-	x	64	46 %
		x	x	84	23 %
	2nd-order	-	-	429	28 %
		x	-	449	14 %
		-	x	99	32 %
		x	x	119	30 %

Table 5.6: Robust decentralized controller synthesis performance results for the LbSync system at European XFEL.

cf. Section 4.5. The loop delays from the digital feedback control implementation were identified as the biggest opponent to closed-loop stability, and the propagation delay on the optical fibre links as significant impact factor for overall performance due to the delay line effect (cf. Section 3.3.1). We thus constructed suitable, loop-shifted LFT forms for both the round-trip experiment model and the PPL chain model where the delays are represented by $\mathcal{S}_\tau$ (cf. Section 2.3.4). We then solved the decentralized controller synthesis problem using the $\mathbf{H}_2$-performance multiplier from Corollary 4.51 and conic combinations of time-delay multipliers (2.41) and (2.43) with static scalings as in Lemma 2.29 for $\nu = 0$. Dynamic scalings are difficult to implement based on MATLAB's systune because of the inability to conjugate transpose tunable system objects in the Control System Toolbox.

The synthesis results are summarized in Table 5.6 and Figure 5.23. In contrast to the nominal synthesis, virtually each of the 250 randomly initialized iterations for each configuration resulted in a different local minimum, with a large percentage not converging to a feasible solution at all. The complexity of the optimization problem is dominated by the size of the performance multiplier, which scales quadratically with the number of exogenous system inputs and linearly with the multiplier order, i.e., for the PPL chain system, the performance multiplier as in (4.92) alone consists of 732 decision variables. To reduce the optimization times to reasonable values, we considered the slightly less general set of multipliers obtained by setting $Z_i = b_i = 0 \; \forall \, i$ in (4.93), halving the performance multiplier's number of decision variables. We also considered the case where additionally $X_{ij} = 0 \; \forall \, i \neq j$, reducing the number further to 36. While the achieved nominal performance values are comparable to the nominal synthesis, the robust performance bounds are two orders of magnitude larger. When performing LMI based robust $\mathbf{H}_2$-performance analysis as in [VSK16; Vee+21], similar robust performance bounds are reported, despite using delay multipliers with dynamic scalings. As shown in Figure 5.23, there is no significant correlation between the robust performance bound and the nominal performance.

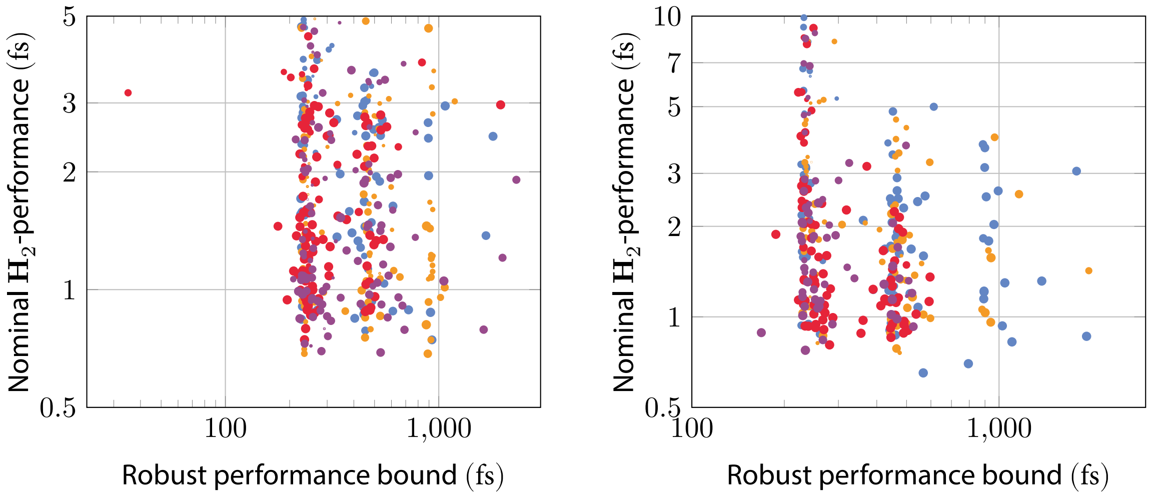

(a) Round-Trip experiment, decentralized PI feedback.

(b) Round-Trip experiment, decentralized 2nd-order feedback.

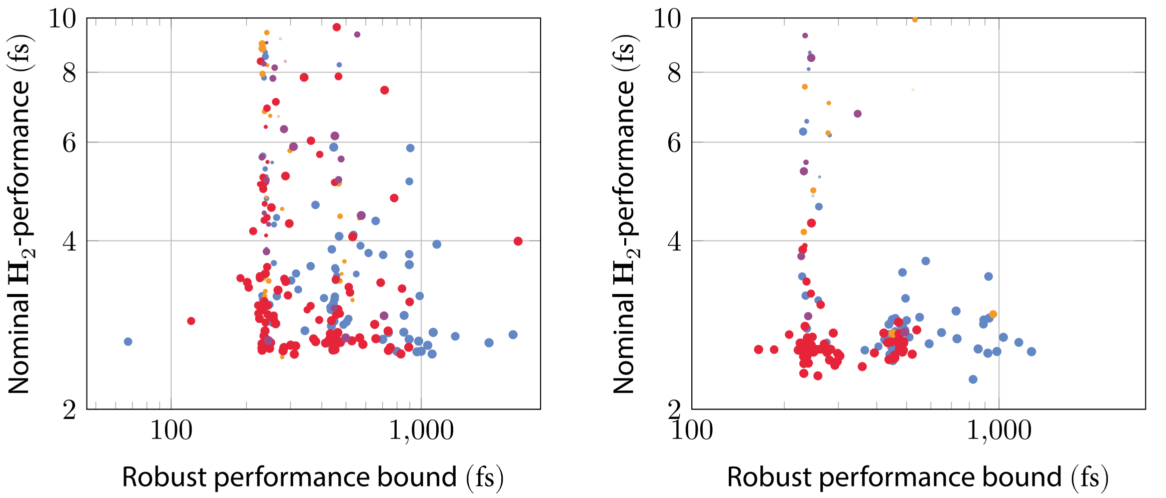

(c) PPL chain, decentralized PI feedback.

(d) PPL chain, decentralized 2nd-order feedback.

Figure 5.23: Distribution of nonsmooth robust synthesis results. ● no additional constraints, ● grouped delay multipliers, ● block diagonal performance multiplier, ● grouped delay multipliers and block diagonal performance multiplier. Mark size inversely proportional to maximum singular value objective.

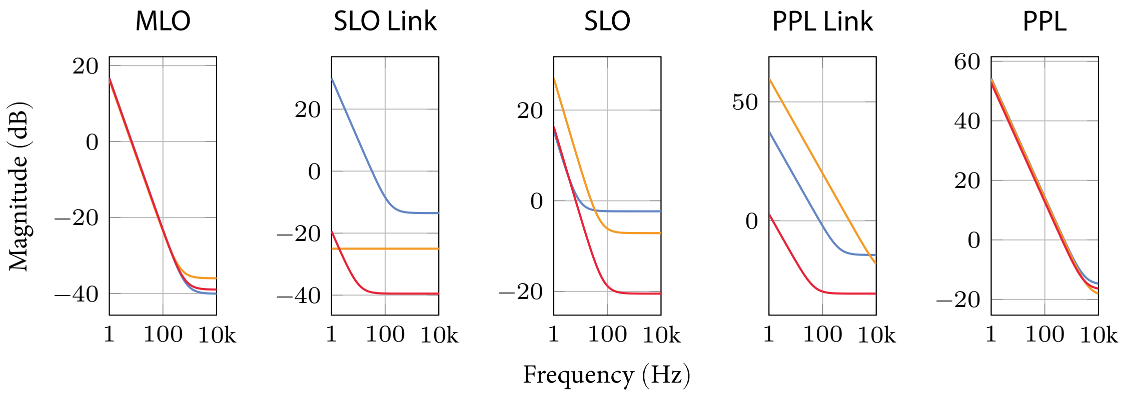

Figure 5.24: Comparison of the best nominal decentralized PI controller (——), the robust controller guaranteeing the best robust performance bound (——) and the robust controller providing the best nominal performance (——) for the PPL jitter in the LbSync system at European XFEL. Shown left to right are the local feedback controllers for the MLO, the SLO link, the SLO, the PPL link, and the SXP PPL oscillator.

To assess if the robust bound is excessively conservative due to individually bounding each delay by an independent multiplier, we considered another synthesis scenario in which the delays are grouped by size and covered by a common multiplier for repeated uncertainties, see [VSK16]. Instead of nine delay multipliers with two decision variables each, we then have three multipliers with a total of 38 decision variables. Figure 5.23 shows that, statistically, this additional freedom cannot be exploited by the optimizer in terms of guaranteed performance, rather the probability of converging to feasible solution decreased because of the additional multiplier variables. Comparing the results from Hermitian (●) and real diagonal (●) performance multipliers, fewer decision variables benefit the consistency of the achieved results, as shown by the increased clustering in the lower left corner.

It's further worthwhile discussing the visible outliers with robust guaranteed performance of less than ∼200 fs. The perhaps most straightforward explanation is "luck", in that only a small subset in the sampling space for the initial random multiplier parameterizations leads to a region of better local minima. Investigating the pole locations a_1, a_2, a_3 of the optimized performance multiplier for the two outliers in Figure 5.23c and the next best results in terms of guaranteed robust performance, we did indeed find that the combination of pole locations were somewhat isolated for the outliers compared to the rest, but not clearly enough in any one direction to allow for a decisive conclusion on how these parameters ought to be initialized.

Figure 5.24 shows the best nominal decentralized PI feedback controller and the robust decentralized PI feedback controllers achieving the best robust performance guarantee and best nominal performance in comparison for the PPL chain case study. As one expects, the robust controllers have overall lower feedback gain, although not persistently across all frequencies. Together with the discussion on outliers above, this raises the question on how to determine a suitable set of initial multiplier parameter values to sample from, and whether warm-starting from the best nominal controller satisfying the desired controller structure can significantly improve the quality and reliable computation of the robust synthesis results. An in-depth answer to the question is beyond the scope of this controller design study and left as future work.

5.3.5 Comparison of Synthesis Methods

Concluding this controller design study, we found that, among the nominal and robust structured controller synthesis procedures presented and developed in this work, the fundamental differences of the methods played out to different advantages and disadvantages depending on the properties of the synthesis problem in terms of system order, input-output dimensionality, multiplier size, number and size of (internal) delays, structural constraints, and problem regularity. A qualitative summary is given in Table 5.7. It should be noted that, bar a few inherent numerical challenges due to e.g. the state-output form and the transformation needed to obtain it, the perceived insufficiencies of the obtained synthesis result, such as marginal stability or overly conservative performance guarantees, are a consequence of the difficulty of the synthesis problem rather than issues with the synthesis methods themselves. Compared to other controller synthesis or tuning techniques, such as (safe) BO or other methods from the realm of machine learning, the methods discussed in this work aim to provide hard mathematical stability and performance guarantees whereas many other approaches only provide probabilistic or heuristic margins. The challenge thus is in designing synthesis techniques that find suboptimal solutions to the (in the case of structured control NP-hard) optimization problem that are tractable, sacrifice as little performance as possible and are certifiably robust. To this end, we feel that for highly structured control problems, nonsmooth methods operating in the frequency domain are superior to time-domain LMI methods because of the high conservatism incurred by convexification of the synthesis problem for models of the size of the LbSync system, and because of their ability to take into account dynamic multipliers. Nevertheless, because fully randomized parameter initializations do not scale indefinitely, large nonsmooth optimization problems need to be suitably initialized, a task that approximate solutions obtained via LMIs and evolutionary searches could help with [WF05; FW06].

5.3.6 Performance Headroom and Limiting Factors

This section discusses the insight gained from the previous controller design study into the current performance headroom for the LbSync system at European XFEL, bottlenecks and recommended changes to the system to be implemented. Figure 5.25a shows the simulated jitter between the PPL oscillator and the MO, similar to Figure 5.21 for the round-trip experiment. In stark contrast however, this configuration does not benefit from decoupling the MLO from the MO, thus increasing the contribution of the MZM based opto-electric PLL to the total RMS timing jitter. Compared to BOXC-based detection, the lower SNR of the MZM detection limits the loop gain and thus the achievable total jitter. The higher absolute phase noise of the MO at Fourier frequencies beyond 1 kHz is also transferred to the MO, which is necessary to minimize the relative phase error between PPL oscillator and MO, but also limits the performance of the optical fibre timing links, as the delay line effect converts this absolute phase noise to relative phase noise as observed in the time-of-flight detector. The higher the closed-loop bandwidth of the stabilized timing links, the more artificial jitter is added to the output, even though the in-loop RMS error is reduced. This emphasizes the need for a global perspective on the LbSync system and joint optimization of the subsystems. It is expected that jointly considering the LbSync system and the longitudinal dynamics of the electrons in the accelerator itself will bring further coupling effects to light.

Recently, the loop delay of the MicroTCA.4 laser PLL setup was investigated, shown to be a limiting factor and reduced by a factor of 6, yielding additional gain margin and allowing for tighter reference locking, yielding up to 50 % in-loop performance, cf. Section 5.2.6. Additional reduction of the loop latency will show diminishing returns, as the waterbed effect will be increasingly dominant. Therefore,

Technique / Synthesis Method	Advantages	Disadvantages
Finite-order approximations	• Straightforward in practice	• Increases system size • Approximation / no hard guarantees
ϵ-Perturbation	• Arbitrarily close gap to infimal performance • Straightforward in practice	• No stability margins (critical for approximated systems) • Numerically ill-behaved for small ε (proper scaling required)
Algebraic Riccati equations	• Efficiently solved	• Full-order controllers • Only for regular(ized) problems
Linear matrix inequalities	• Efficient solvers available • Control over solutions by Lyapunov-shaping	• Conservative • Problem scales quadratically with system order • Structure constraints difficult to incorporate
Nonsmooth optimization	• Local certificate of optimality • Trivial inclusion of structure constraints	• Intractable for moderately many decision variables (multipliers)
Sparsity invariance	• Closest convex restriction for structure constraints	• Conservative
State-output transformation	• Fixed-order output feedback • Compatibility with static multipliers	• Numerically ill-behaved • No support for dynamic multipliers
Integral quadratic constraints	• Supports many classes of uncertainties	• In some cases conservative

Table 5.7: Qualitative comparison of different techniques and methods for handling of time-delay systems, singular problems, structured and robust control synthesis based on the presented LbSync controller design study.

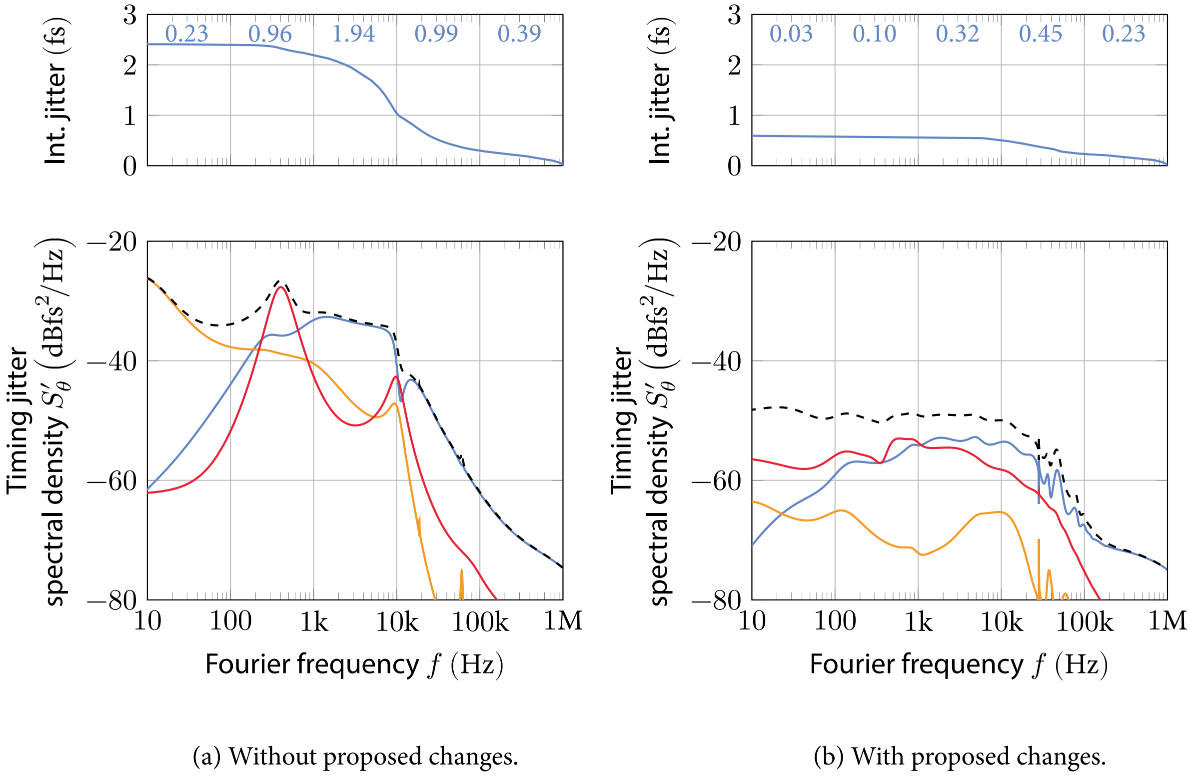

Figure 5.25: Timing jitter of the PPL oscillator at the SXP experiment w.r.t. the MO at European XFEL (- - -). The three most significant contributions come from the MO (——), the MLO (——) and the PPL (——).

the SNR of the MZM and BOXC phase detectors should be increased first. This is a non-trivial task, as the MZM optical-to-RF phase detector used for the MLO at European XFEL is already supplied with excess of 30 dBm RF power. The BOXC SNR is limited by the optical power at the end of a long timing link, which in turn is controlled by the maximum gain of the EDFA amplifiers and needs to be traded-off against optical deterioration of the pulse train. Some minor improvements may be achieved here by operating the EDFA pump diodes in feedback, stabilizing the overall relative intensity noise (RIN) of the optical pulse train, reducing AM-PM noise in imperfectly balanced BOXCs. Another option worth investigating would be higher SHG and SFG conversion efficiencies, which could be obtained with in-fibre implementations [Chr+23]. For the optical fibre timing links, a reduced loop latency is secondary, as the closed-loop bandwidth is limited on long links by the delay line effect and on all links by the piezo fibre stretcher's resonance around 18 kHz. Assuming the delay line effect was mitigated by a more phase pure MO reference and the piezo fibre stretcher reengineered, or its resonance cancelled by a higher-order feedback controller, the expected impact on the total jitter would still be marginal, because of the low noise affecting the link at frequencies beyond the current locking bandwidth.

Lastly, we discuss the need for and effect of higher order feedback controllers and a more highly interconnected information structure. As simulations in Section 3.2.3 have shown, the discussed PLLs benefit from a lead-compensator like boost at higher frequencies. Higher-orders provide little benefits except for plants affected by highly structured disturbances. Notching of piezo or cavity resonances is generally not necessary with modern, well-engineered MLLs. As already discussed, the optical fibre timing links might eventually benefit from notching the limiting 18 kHz resonance, but mechanical engineering solutions to this problem also appear not yet exhausted. With respect to controller interconnection, clear benefits of additional interconnections have been demonstrated. For implementation, the question remains how latency on the FPGA interconnections will deteriorate this benefit, especially for connections of MicroTCA.4 crates located at opposite ends of the facility. However, just considering interconnections of feedback systems located in the same lab, a significant performance improvement can be observed.

In the following list, we summarize recommended changes to the LbSync system and expected cumulative performance improvements for the synchronization of the PPL oscillator, ordered by how quickly those changes could be implemented in the real system. The impact of each measure should not be viewed in isolation, as for example higher-order controllers can only unfold their full potential, if the closed-loop bandwidth is not bottlenecked by insufficient phase or gain margin.

1. The currently used PI feedback controllers should be augmented with a lead compensator. In order to be implementable with the current hardware, the pole should be not faster than 100 kHz. This reduces the total integrated timing jitter by 12.4 % from 2.42 fs to 2.12 fs.

2. At present, the PPL PLL has not yet received the upgrade reducing the loop latency from 12 µs to 2 µs. If this was implemented, and the bandwidth of the piezo driver, which is currently artificially limited to 10 kHz, increased to 100 kHz, the jitter could be lowered another 7.1 % to 1.97 fs.

3. After the 2025 / 2026 shutdown, the next generation of the DESY MO is expected to deliver absolute phase noise performance close to what is hypothesized in Figure 5.11. With this important upgrade, the jitter is lowered a significant 48.7 % to 1.01 fs, close to attosecond performance.

4. Assuming the MZM laser-to-RF phase detection could be improved to lower the noise floor from currently -51.3 dBfs/$\sqrt{\text{Hz}}$ to -60 dBfs/$\sqrt{\text{Hz}}$, another 8.2 % could be scraped of the timing jitter, lowering it to 0.93 fs.

5. Next we consider the removal of the $18\,\mathrm{kHz}$ piezo resonance in the piezo fibre stretchers by efficient notching or mechanical redesign. At the same time, the bandwidth of the piezo driver can be then opened up to $\sim 100\,\mathrm{kHz}$. This has very little benefit for the overall timing jitter, the simulation shows a reduction of merely $1.4\,\%$ to $0.93\,\mathrm{fs}$.

6. Finally, we propose additional controller interconnection. Interconnection of each MLL PLL controller with the controller of the subsequent optical fibre timing link showed vanishing performance gain in simulation. On the other hand, interconnecting all the MLL PLL controllers with suitable interconnection delays gains yet another $34.8\,\%$, achieving a final relative PPL jitter of $600\,\mathrm{as}$. Compared to the original value of $2.42\,\mathrm{fs}$, this accumulates to a total reduction of $75.4\,\%$.

As shown in Figure 5.25b, the relative timing jitter of the PPL oscillator is again (or still) dominated by the already improved MO, followed by the PPL oscillator. As a next step, an RF reference derived from a photonic source should be considered, cf. Section 5.2.2, which could lower the absolute phase noise of the master reference in the critical region of $400\,\mathrm{Hz}$ to $400\,\mathrm{kHz}$ to levels close to $-60\,\mathrm{dBfs}/\sqrt{\mathrm{Hz}}$. Achieving sub $100\,\mathrm{as}$ jitter will however require improvements also on the other subsystems and pose a significant challenge for the next decade.

Chapter 6

Summary, Conclusion and Outlook

6.1 Summary and Conclusion

The analysis, design and implementation of large-scale synchronization systems is highly multidisciplinary. They combine aspects of time and frequency metrology, control theory, electrical engineering and RF technology, signal and probability theory and, in the case of extremely precise laser-based implementations, aspects of optics and photonics. On top, depending on the application, additional domain knowledge, for example from the field of particle accelerator physics and technology, is required. In this thesis, theoretical considerations were decoupled as much as possible from the implementation details, despite providing a hopefully well-rounded introduction to the practical realm of RF systems, optical timing metrology and FEL facilities as one possible application.

In the preliminaries, foundations for the mathematical abstraction of the optimal design challenge were laid out in the form of essentials from functional analysis and LTI system theory. Accounting for the stochastic nature of the timing disturbances any synchronization system seeks to eliminate, a primer on signal and probability theory was provided, with a specific discussion on the stochastic properties of phase instabilities as found in all oscillators, regardless of their technology. Furthermore, foreseeing the importance of time delays in large-scale synchronization networks, and the challenge to account for their infinite-order dynamics, we introduced the IQC framework for robust stability and performance analysis, allowing to eliminate the explicit occurrence of the problematic time-delay terms in the system model. Facilitating the control theoretic analysis, two elementary feedback loops were identified and modelled, the PLL and the auto-correlation timing link, with which any distributed synchronization system can be adequately described from an informational point of view. Additionally, the notion of optimality was formalized in the form of the infinite-horizon integral square error (ISE) and the disturbance rejection problem, phrased in terms of the $\mathbf{H}_2$ closed-loop system norm.

Both optimal design problems were identified as singular, due to a physically motivated lack of control action cost, causing the optimal feedback law to be unbounded. Additionally, for PLLs based on VCOs with integral-type dynamics, the resulting invariant zeros cause the optimal controller to produce instable closed-loop pole-zero cancellations on the imaginary axis. As such, controlled suboptimal solutions were discussed as the only way to obtain implementable feedback solutions. Here, the advantages of LMI synthesis methods have been demonstrated, offering higher control over the solution when compared to standard ARE methods applied to regularized approximations of the singular problem. Nevertheless, because of colouring filters modelling the noise acting on the system, both methods produce high-order controllers, whereas it was shown that low-order controllers obtained by e.g. nonsmooth optimization perform virtually equally well in terms of the closed-loop $\mathbf{H}_2$-norm and are additionally easily constrained so that they are implementable. In the case of the auto-correlation timing link, time delays cause additional complication of the formalism. Here, classic finite-order approximations were compared to hard guarantees obtained using the IQC framework. Again, the close to ill-posed SDP associated with the analysis and synthesis LMIs for the singular problems posed a challenge for numerical solvers – either the solution was virtually unbounded, or additional constraints to enforce boundedness incurred excessive conservatism in the solution to the point of infeasibility. However, with careful numerical conditioning of the model and constraints, implementable results were obtained that outperform previous heuristic tunings.

Novel results for the efficient synthesis of robust (sub)optimal controllers with structural constraints were provided (Chapter 4), accounting for the distributed nature of any large-scale synchronization system. The need for a joint consideration of all involved local feedback loops was demonstrated by means of a simple example. A suitable modelling framework was developed for the tree-like structure of synchronization networks, and the ISE notion generalized to weighted paths to arbitrarily many synchronization clients. The non-convexity of the structured control problem was handled two-fold. First, with the SI framework, the broadest class of convex restrictions (i.e. those convex approximations providing solutions that are guaranteed feasible for the original problem), has been characterized and subsequently extended to dynamic fixed-order output feedback by exploiting an explicit state-space representation. In practice, this state-output form however is frequently numerically ill-behaved, both in terms of computation and in numerical sensitivity of the solutions obtained from it, increasing the requirements on numerical conditioning of the generalized plant model. By dualization, the method was further shown to be compatible with the full-block S-procedure, allowing for robust synthesis under structural LTV uncertainty. Importantly, it was shown that the additional incorporation of multipliers incurs no additional conservatism as the uncertainty approaches zero. Lacking support for the dynamic multipliers needed to effectively bound (uncertain) time delays, an alternative approach operating in the frequency domain was developed. Here, based on recent literature results, the central FDI condition of the IQC framework was translated into a maximum singular value condition that (despite being nonsmooth) can be optimized directly in the frequency domain using subgradient bundle methods. This allows the richly studied IQC framework to be applied efficiently to structured control problems, without need for convexification, as long as the multiplier admits a weakly restrictive structure. Instead, local certificates of optimality are constructed. For this method, a suitable multiplier formulating a $\mathbf{H}_2$-performance bound was developed, along with an efficient implementation for the industry standard MATLAB systune function, based on an analytic solution to the frequency domain integral over said multiplier. A challenge for this method remains the number of optimization variables, especially when dynamic multipliers are involved, leading to many local minima. Exploring the landscape of the cost function thus requires possibly many randomly initialized optimization runs. Numerical examples are provided for each of the novel methods.

The theory developed in this work was subsequently transferred to practical application (Chapter 5). After providing fundamentals on state-of-the-art FELs and relevant aspects of longitudinal stability motivating the need for a (sub)femtosecond precision synchronization system extending several kilometres, the fundamental principles of optical synchronization and latest developments on the DESY LbSync system at European XFEL and FLASH were presented, yielding the parameters for a large subsystem of the network, validated by in-facility measurements. A short comparison to black-box optimization methods in the form of Bayesian optimization was drawn. In a subsequent detailed controller design study, we compared nominal and robust structured output feedback controllers synthesized with the newly developed methods to a series of benchmark results providing lower bounds on the achievable performance. Based on the experience gained from this simulation study, qualitative advantages and disadvantages of the different methods have been summarized. Quantitatively, due to the intractability of gridding such high-dimensional search spaces, the overall conservatism of the obtained solutions cannot be accurately assessed. While large gaps between the guaranteed robust performance and the achieved nominal performance were observed, this may either be attributed to conservatism of the employed methods and reduced-order multipliers, or an unknown worst-case realization of the internal time delays. From the point of view of the LbSync system application, the delay line effect and the absolute stability of the master timing reference were identified as pressing bottlenecks for system performance, but also the addition of lead compensator structures to the local controllers showed significant headroom for improvement. Simulations further showed that increasing the controller interconnectivity has significant potential for further improvement of residual synchronization jitter. In conclusion, this in-depth controller design study for the LbSync system, for the first time taking into account the coupled dynamics of PLLs and stabilized timing links, has shown significant room for improvement up to 75 % nominal performance while at the same time guaranteeing robust stability, and thus provides an authoritative basis for future developments and improvements not only for the LbSync systems at European XFEL and FLASH, but synchronization systems in general.

6.2 Outlook

6.2.1 Robust Structured Output Feedback

The challenging combination of several design criteria such as robustness against (nonlinear) time-varying uncertainties, distributed, decentralized or other forms of structured control in conjunction with output feedback remains an active field of research. In the case of structured control, the problem is known to be NP-hard and will remain so, but the search for suboptimal solutions with low levels of conservatism that are also tractable to compute, appears not yet exhausted. In the following, selected ideas for future work are presented that the author deems most worthwhile pursuing.

With respect to the sparsity invariance framework, research into numerically stable computation of the state-output form would be welcome, making this approach more of a plug & play synthesis method. This likely needs to go hand-in-hand with more stable numerical algorithms for SDP problems, particularly exploiting arbitrary precision computation. It would also be interesting to look for connections between the set of output feedback controllers obtainable from the augmented state-output system given the requirements on separability of the Lyapunov matrix. Informed choices for the optimal matrices T and Γ (cf. Corollary 4.21) also remain an open problem. More generally, formulating numerically well-behaved LMI problems from singular $\mathbf{H}_2$- and similar linear quadratic problems would, in our eyes, benefit many practical applications. One could further attempt to develop the SI method further by showing compatibility with dynamic multipliers, and hence the full IQC framework. Initial attempts by the author have shown little hope for success, however, as the linearizing change of variables utilized in the SI framework is insufficient to render the KYP matrix inequalities convex.

The nonsmooth optimization method has shown promising results, but unlocking its full potential requires in particular better tooling support. The ability to implement custom "tuning goals" for MATLAB's systune and the ability to conjugate transpose tunable LTI objects would allow much more efficient declaration of many IQC multipliers conforming to the structure (4.64). Further, for systems such as the LbSync system with many disturbance inputs, but only a scalar performance output, the frequency-domain $\mathbf{H}_2$-performance multiplier developed in Section 4.5.2 appears unnecessarily large, as it commutes with the inputs of the generalized plant. The duality of the controllability and observability Gramians (cf. Theorem 2.14) suggest that it might be possible to rephrase the entire formalism such that the performance multiplier instead commutes with the outputs of the generalized plant. Initial ideas lead in the direction of utilizing the conjugate transpose of the system, and reformulating the framework for anti-causal systems. This is reminiscent of much earlier, yet still unresolved ideas involving causality of multipliers [PF97; Fer97], but also is found in very recent publications [DHW25]. It would be interesting to see if work in this direction is also able to reduce the conservatism incurred by such multipliers for MIMO systems (cf. Remark 4.50). Finally, the possibility of warm starting the tuning algorithm from a set of promising parameter values is considered a significant step towards reducing the limitations incurred by cost functions with many local minima. Here, LMI solutions to convex approximations of the problem as well as genetic algorithms and other machine learning methods are promising leads to improve the adaptability and reliability of gradient-based methods [WF05].

6.2.2 Laser-Based Synchronization Systems

While thorough understanding of the dynamics of a system is a crucial requirement to be able to compute the optimal feedback law leading to the best possible performance that the given system can achieve, understanding the limitations incurred by e.g. interconnection structures and loop latency is equally important in order to increase the headroom for synthesized controllers to exploit. In Section 5.3.6, based on simulation results and current developments, possible improvements and recommendations were already provided for the LbSync system at European XFEL postulating sub-femtosecond precision. Nevertheless, in the following paragraphs, to conclude this thesis, a more general vision for future optical synchronization systems is presented in the realm of large photon science facilities and beyond, which eventually will require few atto- or even zeptosecond precision [KI09].

Key factors in achieving such extreme levels of stability are 1. sources with extremely low absolute jitter, 2. subsidiary oscillators with high actuation bandwidth, 3. low-loss transmission media, 4. highly sensitive (high SNR) detectors and 5. well-shielded environments. It is clear that optical technologies are the only viable option to achieve the required levels of accuracy and precision. Optical frequency combs referenced to resonant cavities in highly controlled environments already achieve world-record levels of absolute phase noise, both long and short term, and are therefore ideal reference sources. Transferring this stability to the RF domain, as needed by clients based on conventional electronics, is also well understood, but extracting sufficient RF power from the optical signal remains a challenge despite significant progress thanks to specialized photodiodes such as MUTC-PD and "tricks" such as pulse doubling schemes. Other optical precision instruments on the other hand can only benefit from the quality of the source, if the relative phase error can be detected with sufficient fidelity and compensated with sufficient bandwidth. Based on recent research papers, detectors with zeptosecond short-term precision appear well feasible given few- or sub-femtosecond optical pulses with sufficient energy density. On the other end, piezo-mounted cavity mirrors suitable for operational bandwidths beyond 500 kHz have been shown to work [Nak+20b]. For even higher bandwidths, intra-cavity electro-optic or magneto-optic modulators extend the bandwidth to the MHz regime [ESS18b]. In order to utilize such high bandwidths, suitable drivers and low-latency digital control systems are necessary, and luckily available.

The transmission of the phase information from point A to point B hence appears to be a more severe bottleneck, depending on the scale of application. With future particle accelerators expected to be much more compact, FELs may, in that respect, in fact not be the most demanding application in the future. Nevertheless, for other applications such as possibly gravitational wave observatories, the longer the transmission distance, the more severely the delay line effect will counter any improvements made on the absolute phase noise stability on the source. Similarly, the optical properties of the mode-locked laser pulses deteriorate more and more the longer the optical fibre gets, by various effects. Free-space transmission of course would remedy some of the aspects only present with fibre-couples transmission, but is also highly impractical in many settings. Another possibility to counter the delay line effect is the use of additional intermediate repeater oscillators, such that each in-between leg is upper bounded in optical path length, effectively limiting the delay line effect. This however is costly and only worth considering if the intermediate oscillators can be made to track sufficiently well. It would essentially only become feasible if current laser synchronization and link stabilization setups were engineered to be substantially more compact. This would also benefit the maintainability and reliability of optical synchronization networks. With fibre-coupled optical cross-correlators demonstrating best-in-class sensitivity, soon and significant progress on compactification however seems inevitable.

Appendix A

Definitions from Functional Analysis, Probability Theory, and Graph Theory

A.1 Functional Analysis

A.1.1 Functions

Definition A.1 (Function)**.** A *function* $f\colon X \to Y$ is a map that assigns each value in the set X exactly one element in Y. X and Y are then called *domain* and *codomain* of f, respectively. The *range* of f is the set $f(X) = \{y \mid \exists\, x \in X \text{ s.t. } y = f(x)\} \subset Y$. For a subset $Z \subset X$, $f(Z) = \{y \mid \exists\, z \in A \text{ s.t. } y = f(Z)\} \subset Y$ is the *image* of Z under f.

A.1.2 (Normed) Vector Spaces, Banach Spaces, Hilbert Spaces

Definition A.2 (Vector Space [Mar18])**.** A vector space over a scalar field $\mathbf{F}$ is a non-empty set $\mathcal{V}$ of elements, also called vectors, equipped with the two linear operations for

- vector addition: given $u, v \in \mathcal{V}$, $u, v \mapsto u + v \in \mathcal{V}$ and
- scalar multiplication: given $a \in \mathbf{F}$ and $u \in \mathcal{V}$, $u \mapsto au \in \mathcal{V}$

subject to the eight vector space axioms:

1. $u + v = v + u$ for $u, v \in \mathcal{V}$. *Commutativity*
2. $(u + v) + z = u + (v + z)$ for $u, v, z \in \mathcal{V}$. *Associativity*
3. $\exists\, 0 \in \mathcal{V} : u + 0 = u$ for $u \in \mathcal{V}$. *Existence of additive identity (zero vector)*
4. $\forall\, u \in \mathcal{V}\ \exists -u \in \mathcal{V} : u + (-u) = 0$. *Existence of additive inverses*
5. $a(bu) = (ab)u$ for $a, b \in \mathbf{F}$ and $u \in \mathcal{V}$. *Associativity of scalar multiplication*
6. $1u = u$ for $u \in \mathcal{V}$. *Neutrality of scalar identity*
7. $a(u + v) = au + av$ for $a \in \mathbf{F}$ and $u, v \in \mathcal{V}$. *Right distributivity*
8. $(a + b)u = au + bu$ for $a, b \in \mathbf{F}$ and $u \in \mathcal{V}$. *Left distributivity*

Remark A.3. Vectors in the sense of the definition above are a broader concept than conventional n-dimensional Euclidean vectors. Functions mapping into a vector space $\mathcal{V}$ themselves form a vector space, fulfilling the vector space axioms element-wise.

Definition A.4 (Normed Vector Space [Mar18]). A normed vector space is a vector space $\mathcal{V}$ equipped with a norm, i.e. a mapping

$$\|\cdot\|: \mathcal{V} \to \mathbf{R}_+ \tag{A.1}$$

subject to the four norm axioms:

1. $\|u\| \geq 0$ for $u \in \mathcal{V}$. *Nonnegativity*
2. $\|u\| = 0 \Leftrightarrow u = 0$ for $u \in \mathcal{V}$. *Separation*
3. $\|au\| = |a|\|u\|$ for $a \in \mathbf{F}$ and $u \in \mathcal{V}$. *Scalability*
4. $\|u + v\| \leq \|u\| + \|v\|$ for $u, v \in \mathcal{V}$. *Triangle inequality*

Definition A.5 (Banach Space [Mar18]). A normed vector space that is also complete is called a *Banach space*.

Definition A.6 (Inner Product Space (Pre-Hilbert Space) [Mar18]). An *inner product space* over $\mathbf{F}$ is a vector space $\mathcal{V}$ over $\mathbf{F}$ equipped with an inner product, i.e. a mapping

$$\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V} \to \mathbf{F} \tag{A.2}$$

subject to the three inner product axioms:

1. $\langle u, u \rangle \geq 0$ and $\langle u, u \rangle = 0 \Leftrightarrow u = 0$ for $u \in \mathcal{V}$. *Positive definiteness*
2. $\langle u, v \rangle = \overline{\langle v, u \rangle}$ for $u, v \in \mathcal{V}$. *Conjugate symmetry*
3. $\langle au + bv, w \rangle = a\langle u, w \rangle + b\langle v, w \rangle$ for $a, b \in \mathbf{F}$ and $u, v, w \in \mathcal{V}$. *Linearity in the first argument*

Note that an inner product space $\mathcal{V}$ is implicitly a normed vector space by the inner product norm

$$\|u\| = \langle u, u \rangle, \quad u \in \mathcal{V}. \tag{A.3}$$

Definition A.7 (Hilbert Space [Mar18]). A complete inner product space, i.e. a Banach space that has an inner product defined on it, is called a *Hilbert space*.

A.2 Probability Theory, Random Variables and Stochastic Processes

Definition A.8 (Probability Space [PP02]). A probability space is a triplet $(\Omega, \mathcal{F}, P)$ consisting of the *sample space* Ω of all possible experimental outcomes, the σ-algebra of events $\mathcal{F}$, i.e. subsets of Ω with assigned probability, and the probability function $P: \mathcal{F} \to [0, 1]$ assigning a probability to each event in $\mathcal{F}$.

Definition A.9 (Random Variable [PP02]). A random variable defined on a probability space $(\Omega, \mathcal{F}, P)$ is a function $\mathbf{x}(\xi): \Omega \to \mathbf{R}$ assigning a number to each outcome of the experiment and the event $\{\xi \in \Omega : \mathbf{x}(\xi) < x\} \in \mathcal{F}$ for all $x \in \mathbf{R}$.

Remark A.10. The above and all following definitions may be generalized to complex numbers instead of reals. In the context of this context of this work, complex-valued random variables and stochastic processes are however of limited interest and thus not explicitly covered to keep the exposition clear.

Definition A.11 (Stochastic Process [Lap17])**.** A stochastic process $\mathbf{x}(t)$ is a family of random variables indexed in $t \in \mathcal{T}$ and defined on a common probability space. In particular, if $\mathcal{T}$ is a continuous subset of $\mathbf{R}$ ($\mathbf{Z}$), $\mathbf{x}(t)$ is a continuous-time (discrete-time) stochastic process mapping each outcome of the experiment to a continuous-time (discrete-time) signal.

Going forward, to ease notation and to avoid mathematical pitfalls, we will only consider continuous-time (discrete-time) stochastic processes defined on the entire real line (all integers).

Definition A.12 (Strictly Stationary Stochastic Process [Lap17])**.** A continuous-time stochastic process $\mathbf{x}(t)$ is said to be (strictly) stationary if for every $n \in \mathbf{N}$ and any epochs $t_1, ..., t_n \in \mathbf{R}$, the joint distribution of $(\mathbf{x}(t_1), ..., \mathbf{x}(t_n))$ is independent of a shift in time $\tau \in \mathbf{R}$. The definition holds analogously for discrete-time stochastic processes.

Definition A.13 (Wide-Sense Stationary Stochastic Process [Lap17])**.** A stochastic process $\mathbf{x}(t)$ is said to be wide-sense stationary (or weakly stationary) if it is of finite variance, constant mean and the covariance between any two samples is independent of a uniform shift in time.

Definition A.14 (Autocorrelation Function [PP02])**.** The autocorrelation function $R_{\mathbf{x}}(\tau)\colon \mathbf{R} \to \mathbf{R}$ of a wide-sense stationary, continuous-time stochastic process $\mathbf{x}(t)$ is defined as

$$R_{\mathbf{x}}(\tau) := \mathrm{E}[\mathbf{x}(t+\tau)\mathbf{x}(t)]. \tag{A.4}$$

The definition holds analogously for discrete-time stochastic processes.

Remark A.15. We chose Definition A.14 to be consistent with what we feel is the most widely adapted definition in engineering literature. However, we agree with the author of [Lap17] in saying that the naming is somewhat unfortunate, as the definition is missing the normalization with respect to the (in the case of wide-sense stationary processes constant) variance found in e.g. the Pearson correlation coefficient, and as such rather measures the covariance.

The following definition is a direct result of the famous Wiener-Khinchin theorem.

Definition A.16 (power spectral density [PP02])**.** The PSD of a wide-sense stationary, continuous-time stochastic process $\mathbf{x}(t)$ is defined as the Fourier transform $S_{\mathbf{x}}(f)$ of its autocorrelation function $R_{\mathbf{x}}(\tau)$:

$$S_{\mathbf{x}}(f) := \int_{-\infty}^{\infty} R_{\mathbf{x}}(\tau) e^{-2\pi i f \tau}\, \mathrm{d}\tau. \tag{A.5}$$

The PSD of a wide-sense stationary, discrete-time stochastic process is defined analogously via the discrete time Fourier transform (DTFT).

Remark A.17. Although less common, [Lap17] defines the PSD of a wide-sense stationary process closer to the Fourier transform of the autocovariance function rather than the autocorrelation function. While it might seem inappropriate to neglect the DC component of the signal, this does not only lead to a mathematically more well-defined definition, it is also closer to engineering practice when it comes to estimation of the PSD from a time series, where the DC component is generally removed from the time domain data before applying the DFT, in order to eliminate spectral leakage from a significant DC bin to neighbouring low frequency bins.

The following lemma is a direct consequence of Plancherel's theorem.

Lemma A.18. *For a continuous-time, wide-sense stationary process* $\mathbf{x}(t)$ *with PSD* $S_{\mathbf{x}}(f)$ *it holds that*

$$\int_{-\infty}^{\infty} |\mathbf{x}(t)|^2 \, \mathrm{d}t = \int_{-\infty}^{\infty} S_{\mathbf{x}}(f) \, \mathrm{d}f. \tag{A.6}$$

A.3 Graph Theory

Definition A.19 (Directed Graph (Digraph) [Bul])**.** A *directed graph*, or *digraph*, is a pair $G = (V, E)$, where V is a set of *vertices* (or *nodes*) and E is a set of ordered pairs of vertices called *edges*, i.e. $E \subseteq V \times V$. For $u, v \in V$, the ordered pair (u, v) denotes a directed edge from u to v.

Although this work exclusively considers digraphs, we will occasionally use the word *graph*, when the directedness of the edges is not relevant to the context.

Definition A.20 ((Directed) Path [Bul])**.** A *path* in a graph is an ordered sequence of vertices such that any pair of consecutive vertices in the sequence is an edge of the graph. The concept extends analogously to directed paths in digraphs. A vertex $v \in V$ is said to be *reachable* from $u \in V$ if there exists a path from u to v.

Definition A.21 (Connectivity [Bul])**.** A digraph G is *(strongly) connected*, if there exists a directed path between any two vertices. It is *weakly connected* if there exists an undirected path in the underlying (undirected) graph of G.

Definition A.22 (Cycle [Bul])**.** A *cycle* in a (directed) graph is a (directed) path that starts and ends at the same vertex and has at least three vertices. A graph is *acyclic* if it contains no cycles.

Definition A.23 ((Ordered) Tree)**.** A connected acyclic graph is a *tree*. An *ordered tree* is a tree defined on a digraph and there exists a *root* vertex from which there exists a path to every other vertex in the tree.

Note that by definition, the root of an ordered tree is unique.

Appendix B

Supplementary Material

B.1 Calibrations and Measurement Data

Raw measurements obtained from the LbSync system at European XFEL and derived system parameters are available at the Hamburg University of Technology open access research data repository TORE.

10.15480/882.17064

B.2 Model, Analysis, and Synthesis Code

The MATLAB code that was used to prepare results presented in this work is available at the Hamburg University of Technology open access research data repository TORE.

10.15480/882.17065

Bibliography

[Abr02] D. Abramovitch. "Phase-Locked Loops: A control centric tutorial". In: *Proceedings of the 2002 American Control Conference*. Anchorage, AK, USA: IEEE, 2002, pp. 1–15. ISBN: 978-0-7803-7298-6. DOI: 10.1109/ACC.2002.1024769.

[Ack+07] W. Ackermann, G. Asova, V. Ayvazyan, A. Azima, N. Baboi, J. Bähr, V. Balandin, B. Beutner, A. Brandt, A. Bolzmann, R. Brinkmann, O. I. Brovko, M. Castellano, P. Castro, L. Catani, E. Chiadroni, S. Choroba, A. Cianchi, J. T. Costello, D. Cubaynes, J. Dardis, W. Decking, H. Delsim-Hashemi, A. Delserieys, G. Di Pirro, M. Dohlus, S. Düsterer, A. Eckhardt, H. T. Edwards, B. Faatz, J. Feldhaus, K. Flöttmann, J. Frisch, L. Fröhlich, T. Garvey, U. Gensch, C. Gerth, M. Görler, N. Golubeva, H.-J. Grabosch, M. Grecki, O. Grimm, K. Hacker, U. Hahn, J. H. Han, K. Honkavaara, T. Hott, M. Hüning, Y. Ivanisenko, E. Jaeschke, W. Jalmuzna, T. Jezynski, R. Kammering, V. Katalev, K. Kavanagh, E. T. Kennedy, S. Khodyachykh, K. Klose, V. Kocharyan, M. Körfer, M. Kollewe, W. Koprek, S. Korepanov, D. Kostin, M. Krassilnikov, G. Kube, M. Kuhlmann, C. L. S. Lewis, L. Lilje, T. Limberg, D. Lipka, F. Löhl, H. Luna, M. Luong, M. Martins, M. Meyer, P. Michelato, V. Miltchev, W.-D. Möller, L. Monaco, W. F. O. Müller, O. Napieralski, O. Napoly, P. Nicolosi, D. Nölle, T. Nuñez, A. Oppelt, C. Pagani, R. Paparella, N. Pchalek, J. Pedregosa-Gutierrez, B. Petersen, B. Petrosyan, G. Petrosyan, L. Petrosyan, J. Pflüger, E. Plönjes, L. Poletto, K. Pozniak, E. Prat, D. Proch, P. Pucyk, P. Radcliffe, H. Redlin, K. Rehlich, M. Richter, M. Roehrs, J. Roensch, R. Romaniuk, M. Ross, J. Rossbach, V. Rybnikov, M. Sachwitz, E. L. Saldin, W. Sandner, H. Schlarb, B. Schmidt, M. Schmitz, P. Schmüser, J. R. Schneider, E. A. Schneidmiller, S. Schnepp, S. Schreiber, M. Seidel, D. Sertore, A. V. Shabunov, C. Simon, S. Simrock, E. Sombrowski, A. A. Sorokin, P. Spanknebel, R. Spesyvtsev, L. Staykov, B. Steffen, F. Stephan, F. Stulle, H. Thom, K. Tiedtke, M. Tischer, S. Toleikis, R. Treusch, D. Trines, I. Tsakov, E. Vogel, T. Weiland, H. Weise, M. Wellhöfer, M. Wendt, I. Will, A. Winter, K. Wittenburg, W. Wurth, P. Yeates, M. V. Yurkov, I. Zagorodnov and K. Zapfe. "Operation of a free-electron laser from the extreme ultraviolet to the water window". In: *Nature Photonics* 1.6 (2007), pp. 336–342. ISSN: 1749-4885. DOI: 10.1038/nphoton.2007.76.

[AGB14] P. Apkarian, P. Gahinet and C. Buhr. "Multi-model, multi-objective tuning of fixed-structure controllers". In: *2014 European Control Conference (ECC'14)*. Strasbourg, France: IEEE, 2014, pp. 856–861. ISBN: 978-3-9524269-1-3. DOI: 10.1109/ECC.2014.6862200.

[Ahn+22] C. Ahn, Y. Na, M. Hyun, J. Bae and J. Kim. "Synchronization of an optical frequency comb and a microwave oscillator with $53\,\mathrm{zs}/\sqrt{\mathrm{Hz}}$ resolution and 10^{-20}-level stability". In: *Photonics Research* 10.2 (2022), pp. 365–372. DOI: 10.1364/PRJ.443316.

[AM06] P. J. Antsaklis and A. N. Michel. *Linear systems*. Birkhäuser, 2006. ISBN: 978-0-8176-4434-5. DOI: 10.1007/0-8176-4435-0.

[ÅM21] K. J. Åström and R. M. Murray. *Feedback systems: An introduction for scientists and engineers*. Second edition. Princeton University Press, 2021. ISBN: 978-0-691-19398-4.

[AM89] B. D. Anderson and J. B. Moore. *Optimal control: Linear quadratic methods*. Prentice Hall information and system sciences series. Prentice Hall, 1989. ISBN: 978-0-13-638651-3.

[AN06a] P. Apkarian and D. Noll. "IQC analysis and synthesis via nonsmooth optimization". In: *Systems & Control Letters* 55.12 (2006), pp. 971–981. ISSN: 0167-6911. DOI: `10.1016/j.sysconle.2006.06.008`.

[AN06b] P. Apkarian and D. Noll. "Nonsmooth H_∞ Synthesis". In: *IEEE Transactions on Automatic Control* 51.1 (2006), pp. 71–86. ISSN: 1558-2523. DOI: `10.1109/TAC.2005.860290`.

[ANR08] P. Apkarian, D. Noll and A. Rondepierre. "Mixed H_2/H_∞ Control via Nonsmooth Optimization". In: *SIAM Journal on Control and Optimization* 47.3 (2008), pp. 1516–1546. ISSN: 0363-0129. DOI: `10.1137/070685026`.

[Ass+20] R. W. Assmann, M. K. Weikum, T. Akhter, D. Alesini, A. S. Alexandrova, M. P. Anania, N. E. Andreev, I. Andriyash, M. Artioli, A. Aschikhin, T. Audet, A. Bacci, I. F. Barna, S. Bartocci, A. Bayramian, A. Beaton, A. Beck, M. Bellaveglia, A. Beluze, A. Bernhard, A. Biagioni, S. Bielawski, F. G. Bisesto, A. Bonatto, L. Boulton, F. Brandi, R. Brinkmann, F. Briquez, F. Brottier, E. Bründermann, M. Büscher, B. Buonomo, M. H. Bussmann, G. Bussolino, P. Campana, S. Cantarella, K. Cassou, A. Chancé, M. Chen, E. Chiadroni, A. Cianchi, F. Cioeta, J. A. Clarke, J. M. Cole, G. Costa, M.-E. Couprie, J. Cowley, M. Croia, B. Cros, P. A. Crump, R. D'Arcy, G. Dattoli, A. Del Dotto, N. Delerue, M. Del Franco, P. Delinikolas, S. de Nicola, J. M. Dias, D. Di Giovenale, M. Diomede, E. Di Pasquale, G. Di Pirro, G. Di Raddo, U. Dorda, A. C. Erlandson, K. Ertel, A. Esposito, F. Falcoz, A. Falone, R. Fedele, A. Ferran Pousa, M. Ferrario, F. Filippi, J. Fils, G. Fiore, R. Fiorito, R. A. Fonseca, G. Franzini, M. Galimberti, A. Gallo, T. C. Galvin, A. Ghaith, A. Ghigo, D. Giove, A. Giribono, L. A. Gizzi, F. J. Grüner, A. F. Habib, C. Haefner, T. Heinemann, A. Helm, B. Hidding, B. J. Holzer, S. M. Hooker, T. Hosokai, M. Hübner, M. Ibison, S. Incremona, A. Irman, F. Iungo, F. J. Jafarinia, O. Jakobsson, D. A. Jaroszynski, S. Jaster-Merz, C. Joshi, M. Kaluza, M. Kando, O. S. Karger, S. Karsch, E. Khazanov, D. Khikhlukha, M. Kirchen, G. Kirwan, C. Kitégi, A. Knetsch, D. Kocon, P. Koester, O. S. Kononenko, G. Korn, I. Kostyukov, K. O. Kruchinin, L. Labate, C. Le Blanc, C. Lechner, P. Lee, W. Leemans, A. Lehrach, X. Li, Y. Li, V. Libov, A. Lifschitz, C. A. Lindstrøm, V. Litvinenko, W. Lu, O. Lundh, A. R. Maier, V. Malka, G. G. Manahan, S. P. D. Mangles, A. Marcelli, B. Marchetti, O. Marcouillé, A. Marocchino, F. Marteau, A. La Martinez de Ossa, J. L. Martins, P. D. Mason, F. Massimo, F. Mathieu, G. Maynard, Z. Mazzotta, S. Mironov, A. Y. Molodozhentsev, S. Morante, A. Mosnier, A. Mostacci, A.-S. Müller, C. D. Murphy, Z. Najmudin, P. A. P. Nghiem, F. Nguyen, P. Niknejadi, A. Nutter, J. Osterhoff, D. Oumbarek Espinos, J.-L. Paillard, D. N. Papadopoulos, B. Patrizi, R. Pattathil, L. Pellegrino, A. Petralia, V. Petrillo, L. Piersanti, M. A. Pocsai, K. Poder, R. Pompili, L. Pribyl, D. Pugacheva, B. A. Reagan, J. Resta-Lopez, R. Ricci, S. Romeo, M. Rossetti Conti, A. R. Rossi, R. Rossmanith, U. Rotundo, E. Roussel, L. Sabbatini, P. Santangelo, G. Sarri, L. Schaper, P. Scherkl, U. Schramm, C. B. Schroeder, J. Scifo, L. Serafini, G. Sharma, Z. M. Sheng, V. Shpakov, C. W. Siders, L. O. Silva, T. Silva, C. Simon, C. Simon-Boisson, U. Sinha, E. Sistrunk, A. Specka, T. M. Spinka, A. Stecchi, A. Stella, F. Stellato, M. J. V. Streeter, A. Sutherland, E. N. Svystun, D. Symes, C. Szwaj, G. E. Tauscher, D. Terzani, G. Toci, P. Tomassini, R. Torres, D. Ullmann, C. Vaccarezza, M. Valléau, M. Vannini, A. Vannozzi, S. Vescovi, J. M. Vieira, F. Villa, C.-G. Wahlström, R. Walczak, P. A. Walker, K. Wang, A. Welsch, C. P. Welsch, S. M. Weng, S. M. Wiggins,

J. Wolfenden, G. Xia, M. Yabashi, H. Zhang, Y. Zhao, J. Zhu and A. Zigler. "EuPRAXIA Conceptual Design Report". In: *The European Physical Journal Special Topics* 229.24 (2020), pp. 3675–4284. ISSN: 1951-6355. DOI: 10.1140/epjst/e2020-000127-8.

[Ati+24] Atia-tul-noor, S. Kumar, N. Schirmel, B. Erk, B. Manschwetus, S. Alisaukas, M. Braune, G. Cirmi, M. K. Czwalinna, U. Frühling, U. Grosse-Wortmann, N. Kschuev, F. Kuschewski, T. Lang, H. Lindenblatt, I. Litvinyuk, S. Meister, R. Moshammer, C. C. Papadopoulou, C. Passow, J. Roensch-Schulenburg, F. Trost, I. Hartl, S. Düsterer and S. Schulz. "Sub-50 fs temporal resolution in an FEL-optical laser pump-probe experiment at FLASH2". In: *Optics Express* 32.4 (Feb. 2024), pp. 6597–6608. DOI: 10.1364/OE.513714.

[Aun+00] B. Aune, R. Bandelmann, D. Bloess, B. Bonin, A. Bosotti, M. Champion, C. Crawford, G. Deppe, B. Dwersteg, D. A. Edwards, H. T. Edwards, M. Ferrario, M. Fouaidy, P.-D. Gall, A. Gamp, A. Gössel, J. Graber, D. Hubert, M. Hüning, M. Juillard, T. Junquera, H. Kaiser, G. Kreps, M. Kuchnir, R. Lange, M. Leenen, M. Liepe, L. Lilje, A. Matheisen, W.-D. Möller, A. Mosnier, H. Padamsee, C. Pagani, M. Pekeler, H.-B. Peters, O. Peters, D. Proch, K. Rehlich, D. Reschke, H. Safa, T. Schilcher, P. Schmüser, J. Sekutowicz, S. Simrock, W. Singer, M. Tigner, D. Trines, K. Twarowski, G. Weichert, J. Weisend, J. Wojtkiewicz, S. Wolff and K. Zapfe. "Superconducting TESLA cavities". In: *Physical Review Special Topics - Accelerators and Beams* 3.9 (2000). ISSN: 1098-4402. DOI: 10.1103/PhysRevSTAB.3.092001.

[Ban17] D. Banerjee. *PLL Performance, Simulation, and Design*. 5th ed. Dog Ear Publishing, 2017. ISBN: 978-1-4575-5177-2.

[Bar+05] A. Bartels, S. A. Diddams, C. W. Oates, G. Wilpers, J. C. Bergquist, W. H. Oskay and L. Hollberg. "Femtosecond-laser-based synthesis of ultrastable microwave signals from optical frequency references". In: *Optics letters* 30.6 (2005), pp. 667–669. ISSN: 0146-9592. DOI: 10.1364/OL.30.000667.

[Başı00] T. Başar, ed. *Control theory: Twenty-five seminal papers (1931-1981)*. Wiley-IEEE Press, 2000. ISBN: 978-0-7803-6021-1.

[BD81] I. R. Bartky and S. J. Dick. "The First Time Balls". In: *Journal for the History of Astronomy* 12.3 (1981), pp. 155–164. ISSN: 0021-8286. DOI: 10.1177/002182868101200301.

[BD94] K. P. Birch and M. J. Downs. "Correction to the Updated Edlén Equation for the Refractive Index of Air". In: *Metrologia* 31.4 (1994), pp. 315–316. ISSN: 0026-1394. DOI: 10.1088/0026-1394/31/4/006.

[Beh21] T. Behne. "Adaptive Vorwärtskompensation in einer Laser-Phasenregelschleife innerhalb des optischen Synchronisationssytems des European XFEL". Bachelor's Thesis. Hamburg, Germany: Hamburg University of Technology, 2021.

[Bey+23] M. Beye, M. Gühr, I. Hartl, E. Plönjes, L. Schaper, S. Schreiber, K. Tiedtke and R. Treusch. "FLASH and the FLASH2020+ project – current status and upgrades for the free-electron laser in Hamburg at DESY". In: *The European Physical Journal Plus* 138.3 (2023). DOI: 10.1140/epjp/s13360-023-03814-8.

[BJ65] A. Bryson and D. Johansen. "Linear filtering for time-varying systems using measurements containing colored noise". In: *IEEE Transactions on Automatic Control* 10.1 (1965), pp. 4–10. ISSN: 1558-2523. DOI: 10.1109/TAC.1965.1098063.

[Bra18] B. Brannon. "Some recent developments in the art of receiver technology: A selected history on receiver innovations over the last 100 years". In: *Analog Dialogue* 52.3 (2018), pp. 21–29. ISSN: 1552-3284.

[BT58] R. B. Blackman and J. W. Tukey. “The measurement of power spectra from the point of view of communications engineering — Part I”. In: *The Bell System Technical Journal* 37.1 (1958), pp. 185–282. ISSN: 0005-8580. DOI: 10.1002/j.1538-7305.1958.tb03874.x.

[Bul] F. Bullo. *Lectures on Network Systems*. 1.7. Kindle Direct Publishing. ISBN: 978-1-986425-64-3. URL: https://fbullo.github.io/lns/ (visited on 14/02/2025).

[But+15] L. Butkowski, T. Kozak, P. Prędki, R. Rybaniec and B. Y. Yang. “FPGA firmware framework for MTCA. 4 AMC modules”. In: *Proceedings of the 15th International Conference on Accelerator & Large Experimental Physics Control Systems (ICALEPCS 2015)*. Melbourne, Australia: JACoW, 2015, pp. 17–23. ISBN: 978-3-95450-148-9. URL: https://bib-pubdb1.desy.de/record/292499/files/wepgf074.pdf (visited on 05/02/2025).

[CA78] D. J. Clements and B. D. O. Anderson. *Singular Optimal Control: The Linear-Quadratic Problem*. Vol. 5. Springer, 1978. ISBN: 978-3-540-08694-9. DOI: 10.1007/BFb0004988.

[CG96] M. Chilali and P. Gahinet. “H/sub ∞/ design with pole placement constraints: an LMI approach”. In: *IEEE Transactions on Automatic Control* 41.3 (1996), pp. 358–367. ISSN: 1558-2523. DOI: 10.1109/9.486637.

[Cha+20] M. Chamanbaz, M. Sznaier, C. Lagoa and F. Dabbene. “Probabilistic Discrete Time Robust H2 Controller Design”. In: *Proceedings of the 59th IEEE Conference on Decision and Control (CDC’20)*. Jeju, Korea (South): IEEE, 2020, pp. 2240–2245. ISBN: 978-1-7281-7447-1. DOI: 10.1109/CDC42340.2020.9304278.

[Cho+10] C. W. Chou, D. B. Hume, T. Rosenband and D. J. Wineland. “Optical clocks and relativity”. In: *Science (New York, N.Y.)* 329.5999 (2010), pp. 1630–1633. DOI: 10.1126/science.1192720.

[Chr+23] J. Christie, L. Corner, J. Henderson and E. Snedden. “Developing a two-colour all-fibre balanced optical cross-correlator for sub-femtosecond synchronisation: JACoW Publishing”. In: *Proceedings of the 14th International Particle Accelerator Conference*. Ed. by I. Andrian. Venice, Italy: JACoW, 2023. ISBN: 978-3-95450-231-8. DOI: 10.18429/JACoW-IPAC2023-THPA073.

[Cid96] P. E. Ciddor. “Refractive index of air: new equations for the visible and near infrared”. In: *Applied optics* 35.9 (1996), pp. 1566–1573. DOI: 10.1364/AO.35.001566.

[Cla90] F. H. Clarke. *Optimization and Nonsmooth Analysis*. Society for Industrial and Applied Mathematics, 1990. ISBN: 978-0-89871-256-8. DOI: 10.1137/1.9781611971309.

[Con+12] C. Conte, N. R. Voellmy, M. N. Zeilinger, M. Morari and C. N. Jones. “Distributed synthesis and control of constrained linear systems”. In: *2012 American Control Conference (ACC)*. Montreal, QC: IEEE, 2012, pp. 6017–6022. ISBN: 978-1-4577-1096-4. DOI: 10.1109/ACC.2012.6314654.

[Cou14] M. E. Couprie. “New generation of light sources: Present and future”. In: *Journal of Electron Spectroscopy and Related Phenomena* 196 (2014), pp. 3–13. ISSN: 0368-2048. DOI: 10.1016/j.elspec.2013.12.007.

[CS06] J.-F. Cliche and B. Shillue. “Precision Timing Control for Radioastronomy: Maintaining Femtosecond Synchronization in the Atacama Large Millimeter Array”. In: *IEEE Control Systems* 26.1 (2006), pp. 19–26. ISSN: 1066-033X. DOI: 10.1109/MCS.2006.1580149.

[CS11] K. Czuba and D. Sikora. “Temperature Stability of Coaxial Cables”. In: *Acta Phys. Pol. A* 119 (2011), p. 553. URL: https://cds.cern.ch/record/1349292.

[CS20] V. M. G. B. Cavalcanti and A. M. Simões. “IQC-synthesis under structural constraints”. In: *International Journal of Robust and Nonlinear Control* 30.13 (2020), pp. 4880–4905. ISSN: 1049-8923. DOI: 10.1002/rnc.5022.

[Czu+13] K. Czuba, D. Sikora, L. Zembala, J. Branlard, F. Ludwig, H. Schlarb and H. Weddig. "Overview of the RF Synchronization System for the European XFEL". In: *Proceedings of 4th International Particle Accelerator Conference, IPAC 2013, May 13-17, 2013, Shanghai, China*. Ed. by Z. Dai. [S. l.]: JACoW, 2013, pp. 3002–3003. ISBN: 978-3-95450-122-9. URL: `https://accelconf.web.cern.ch/IPAC2013/papers/wepme035.pdf` (visited on 09/09/2024).

[Czw+21] M. Czwalinna, R. Boll, H. Kirkwood, J. Koliyadu, J. Kral, B. Lautenschlager, R. Letrun, J. Liu, J. Müller, F. Pallas, D. Rivas, T. Sato, H. Schlarb, S. Schulz and B. Steffen, eds. *Beam Arrival Stability at the European XFEL: JACoW Publishing, Geneva, Switzerland*. 2021. DOI: `10.18429/JACoW-IPAC2021-THXB02`.

[Czw+24] M. K. Czwalinna, B. Lautenschlager, F. Ludwig, H. Schlarb, J. Branlard and S. Schulz. *Advances in fs synchronization*. 2024. DOI: `10.18429/JACoW-LINAC2024-WEXA002`.

[DD03] R. D'Andrea and G. E. Dullerud. "Distributed control design for spatially interconnected systems". In: *IEEE Transactions on Automatic Control* 48.9 (2003), pp. 1478–1495. ISSN: 1558-2523. DOI: `10.1109/TAC.2003.816954`.

[Dem06] A. Demir. "Computing Timing Jitter From Phase Noise Spectra for Oscillators and Phase-Locked Loops With White and 1/f Noise". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 53.9 (2006), pp. 1869–1884. ISSN: 1558-0806. DOI: `10.1109/TCSI.2006.881184`.

[Dena] P. Denisowski. *White Paper: Understanding phase noise fundamentals*. Ed. by Rohde & Schwarz. Munich, Germany.

[Denb] P. Denisowski. *White Paper: Understanding phase noise measurement techniques*. Ed. by Rohde & Schwarz. Munich, Germany.

[Den+22] Y. Deng, V. Léchappé, E. Moulay, Z. Chen, B. Liang, F. Plestan and Q.-L. Han. "Predictor-based control of time-delay systems: a survey". In: *International Journal of Systems Science* 53.12 (2022), pp. 2496–2534. ISSN: 0020-7721. DOI: `10.1080/00207721.2022.2056654`.

[DES24] DESY. *Accelerators 2023: highlights and annual report*. Deutsches Elektronen-Synchrotron (DESY), 2024. DOI: `10.3204/PUBDB-2024-06032`.

[DFT92] J. C. Doyle, B. A. Francis and A. Tannenbaum. *Feedback control theory*. Macmillan, 1992. ISBN: 978-0-02-330011-0.

[DHW25] A. Datar, C. Hespe and H. Werner. "Robust Performance Analysis of Cooperative Control Dynamics via Integral Quadratic Constraints". In: *IEEE Transactions on Automatic Control* 70.1 (2025), pp. 557–564. ISSN: 1558-2523. DOI: `10.1109/TAC.2024.3440840`.

[Dim+24] N. Dimarcq, M. Gertsvolf, G. Mileti, S. Bize, C. W. Oates, E. Peik, D. Calonico, T. Ido, P. Tavella, F. Meynadier, G. Petit, G. Panfilo, J. Bartholomew, P. Defraigne, E. A. Donley, P. O. Hedekvist, I. Sesia, M. Wouters, P. Dubé, F. Fang, F. Levi, J. Lodewyck, H. S. Margolis, D. Newell, S. Slyusarev, S. Weyers, J.-P. Uzan, M. Yasuda, D.-H. Yu, C. Rieck, H. Schnatz, Y. Hanado, M. Fujieda, P.-E. Pottie, J. Hanssen, A. Malimon and N. Ashby. "Roadmap towards the redefinition of the second". In: *Metrologia* 61.1 (2024), p. 012001. ISSN: 0026-1394. DOI: `10.1088/1681-7575/ad17d2`.

[Din18] H. Dinter. "Longitudinal Diagnostics for Beam-Based Intra Bunch-Train Feedback at FLASH and the European XFEL". Dissertation. Hamburg: Universität Hamburg, 2018. DOI: `10.3204/PUBDB-2018-03166`. URL: `https://ediss.sub.uni-hamburg.de/handle/ediss/7794` (visited on 01/02/2025).

[DMR98] A. Demir, A. Mehrotra and J. Roychowdhury. "Phase noise in oscillators: a unifying theory and numerical methods for characterisation". In: *Proceedings of the 35th Annual Conference on Design Automation*. Ed. by B. R. Chawla, R. E. Bryant, J. M. Rabaey and M. J. Irwin. San Francisco, California, United States: ACM Press, 1998, pp. 26–31. ISBN: 978-0-89791-964-7. DOI: 10.1145/277044.277050.

[Doy+89] J. C. Doyle, K. Glover, P. P. Khargonekar and B. A. Francis. "State-space solutions to standard H_2 and H_∞ control problems". In: *IEEE Transactions on Automatic Control* 34.8 (1989), pp. 831–847. ISSN: 1558-2523. DOI: 10.1109/9.29425.

[Doy78] J. Doyle. "Guaranteed margins for LQG regulators". In: *IEEE Transactions on Automatic Control* 23.4 (1978), pp. 756–757. ISSN: 1558-2523. DOI: 10.1109/TAC.1978.1101812.

[DP00] G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory: A Convex Approach*. Vol. 36. Texts in Applied Mathematics. Springer, 2000. ISBN: 978-0-387-98945-7. DOI: 10.1007/978-1-4757-3290-0.

[DS18] N. Da Dalt and A. Sheikholeslami. *Understanding Jitter and Phase Noise*. Cambridge University Press, 2018. ISBN: 978-1-107-18857-0. DOI: 10.1017/9781316981238.

[DV75] C. A. Desoer and M. Vidyasagar. *Feedback Systems: Input-Output Properties*. 1st ed. Elsevier, 1975. ISBN: 978-0-12-212050-3. DOI: 10.1016/B978-0-12-212050-3.X5001-4.

[EHW14] A. Eichler, C. Hoffmann and H. Werner. "Robust control of decomposable LPV systems". In: *Automatica* 50.12 (2014), pp. 3239–3245. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2014.10.046. (Visited on 11/02/2021).

[Eic16] A. Eichler. *Interconnected Systems: Analysis and Distributed Controller Synthesis*. Dissertation. Verlag Dr. Hut, 2016. ISBN: 978-3-8439-2880-9. URL: https://katalog.tub.tuhh.de/Record/874994829 (visited on 04/02/2025).

[Ein05] A. Einstein. "Zur Elektrodynamik bewegter Körper". In: *Annalen der Physik* 322.10 (1905), pp. 891–921. ISSN: 0003-3804. DOI: 10.1002/andp.19053221004.

[ESS18a] M. Endo, T. D. Shoji and T. R. Schibli. "High-sensitivity optical to microwave comparison with dual-output Mach-Zehnder modulators". In: *Scientific reports* 8.1 (2018), p. 4388. DOI: 10.1038/s41598-018-22621-1.

[ESS18b] M. Endo, T. D. Shoji and T. R. Schibli. "Ultralow Noise Optical Frequency Combs". In: *IEEE Journal of Selected Topics in Quantum Electronics* 24.5 (2018), pp. 1–13. ISSN: 1077-260X. DOI: 10.1109/JSTQE.2018.2818461.

[Eur] European XFEL GmbH, ed. *European XFEL in international comparison*. Webpage. URL: https://www.xfel.eu/facility/comparison/index_eng.html (visited on 04/02/2025).

[Eur24] European X-Ray Free-Electron Laser Facility GmbH. *Public Corporate Governance Report 2023*. Schenefeld, Germany, 2024. URL: https://www.xfel.eu/compliance/index_eng.html (visited on 04/02/2025).

[Faz+17] G. Fazelnia, R. Madani, A. Kalbat and J. Lavaei. "Convex Relaxation for Optimal Distributed Control Problems". In: *IEEE Transactions on Automatic Control* 62.1 (2017), pp. 206–221. ISSN: 1558-2523. DOI: 10.1109/TAC.2016.2562062.

[FB19] T. Fortier and E. Baumann. "20 years of developments in optical frequency comb technology and applications". In: *Communications Physics* 2.1 (2019). DOI: 10.1038/s42005-019-0249-y.

[Fel+15] M. Felber, M. K. Czwalinna, H. T. Duhme, M. Fenner, C. Gerth, M. Heuer, T. Lamb, U. Mavric, J. Müller, P. Peier, H. Schlarb, S. Schulz, B. Steffen, C. Sydlo, M. Titberidze, T. Walter, R. Wedel, F. Zummack, J. Szewinski, T. Kozak, P. Prędki, K. Przygoda and E. Janas. "New

MTCA.4-Based Hardware Developments for the Control of the Optical Synchronization Systems at DESY". In: *Proceedings of the 3rd International Beam Instrumentation Conference*. Monterey, California: JACoW, 2015, pp. 152–155. ISBN: 978-3-95450-141-0. URL: https://bib-pubdb1.desy.de/record/203529/files/PUBDB-2014-04566.PDF (visited on 06/10/2020).

[Fer97] E. Feron. "Analysis of Robust H2 Performance Using Multiplier Theory". In: *SIAM Journal on Control and Optimization* 35.1 (1997), pp. 160–177. ISSN: 0363-0129. DOI: 10.1137/S0363012994266504.

[FK18] L. Furieri and M. Kamgarpour. "Robust Distributed Control Beyond Quadratic Invariance". In: *Proceedings of the 57th IEEE Conference on Decision and Control (CDC'18)*. Miami Beach, FL, USA: IEEE, 2018, pp. 3728–3733. ISBN: 978-1-5386-1395-5. DOI: 10.1109/CDC.2018.8619137.

[For+07] S. M. Foreman, K. W. Holman, D. D. Hudson, D. J. Jones and J. Ye. "Remote transfer of ultrastable frequency references via fiber networks". In: *The Review of scientific instruments* 78.2 (2007), p. 021101. ISSN: 0034-6748. DOI: 10.1063/1.2437069.

[For+13] T. M. Fortier, F. Quinlan, A. Hati, C. Nelson, J. A. Taylor, Y. Fu, J. Campbell and S. A. Diddams. "Photonic microwave generation with high-power photodiodes". In: *Optics letters* 38.10 (2013), pp. 1712–1714. ISSN: 0146-9592. DOI: 10.1364/OL.38.001712.

[FPE19] G. F. Franklin, J. D. Powell and A. F. Emami-Naeini. *Feedback Control of Dynamic Systems*. 8th ed. Pearson Education Limited, 2019. ISBN: 978-1-292-27452-2. URL: https://elibrary.pearson.de/book/99.150005/9781292274546 (visited on 01/02/2025).

[Fra79] B. Francis. "The Optimal Linear-Quadratic Time-Invariant Regulator with Cheap Control". In: *IEEE Transactions on Automatic Control* 24.4 (1979), pp. 616–621. ISSN: 1558-2523. DOI: 10.1109/TAC.1979.1102097.

[Fri14] E. Fridman. *Introduction to Time-Delay Systems*. Springer International Publishing, 2014. ISBN: 978-3-319-09392-5. DOI: 10.1007/978-3-319-09393-2.

[Fur+19] L. Furieri, Y. Zheng, A. Papachristodoulou and M. Kamgarpour. "An Input-Output Parametrization of Stabilizing Controllers: amidst Youla and System Level Synthesis". In: *IEEE Control Systems Letters* 3.4 (2019), pp. 1014–1019. ISSN: 2475-1456. DOI: 10.1109/LCSYS.2019.2920205. URL: http://arxiv.org/pdf/1903.03828v2.

[Fur+20] L. Furieri, Y. Zheng, A. Papachristodoulou and M. Kamgarpour. "Sparsity Invariance for Convex Design of Distributed Controllers". In: *IEEE Transactions on Control of Network Systems* 7.4 (2020), pp. 1836–1847. DOI: 10.1109/TCNS.2020.3002429. (Visited on 11/02/2021).

[FW06] A. Farag and H. Werner. "Fixed-structure μ-synthesis - an evolutionary approach". In: *Proceedings of the 2006 American Control Conference (ACC'06)*. Minneapolis, MN, USA: IEEE, 2006, pp. 4332–4337. ISBN: 978-1-4244-0209-0. DOI: 10.1109/ACC.2006.1657400.

[GA11] P. Gahinet and P. Apkarian. "Structured H_∞ Synthesis in MATLAB". In: *IFAC Proceedings Volumes* 44.1 (2011), pp. 1435–1440. ISSN: 1474-6670. DOI: 10.3182/20110828-6-IT-1002.00708.

[GA94] P. Gahinet and P. Apkarian. "A linear matrix inequality approach to H_∞ control". In: *International Journal of Robust and Nonlinear Control* 4.4 (1994), pp. 421–448. ISSN: 1049-8923. DOI: 10.1002/rnc.4590040403.

[Gar05] F. M. Gardner. *Phaselock techniques*. 3rd ed. Wiley-Interscience, 2005. ISBN: 978-0-471-73268-6. DOI: 10.1002/0471732699.

[Gas+20] B. Gasowski, M. Urbanski, K. Czuba, H. Pryschelski, F. Ludwig and J. Branlard. "Concept of Master Oscillator Upgrade for FLASH". In: *2020 23rd International Microwave and Radar Conference (MIKON)*. Warsaw, Poland: IEEE, 2020, pp. 165–168. ISBN: 978-83-949421-7-5. DOI: 10.23919/MIKON48703.2020.9253858.

[Gee89] T. Geerts. "All optimal controls for the singular linear-quadratic problem without stability; a new interpretation of the optimal cost". In: *Linear Algebra and its Applications* 116 (1989), pp. 135–181. ISSN: 0024-3795. DOI: 10.1016/0024-3795(89)90403-5.

[Grü+19] J. Grünert, M. P. Carbonell, F. Dietrich, T. Falk, W. Freund, A. Koch, N. Kujala, J. Laksman, J. Liu, T. Maltezopoulos, K. Tiedtke, U. F. Jastrow, A. Sorokin, E. Syresin, A. Grebentsov and O. Brovko. "X-ray photon diagnostics at the European XFEL". In: *Journal of synchrotron radiation* 26.Pt 5 (2019), pp. 1422–1431. DOI: 10.1107/S1600577519006611.

[Gru+20] Y. Gruson, A. Rus, U. L. Rohde, A. Roth and E. Rubiola. "Artifacts and errors in cross-spectrum phase noise measurements". In: *Metrologia* 57.5 (2020), p. 055010. ISSN: 0026-1394. DOI: 10.1088/1681-7575/ab8d7b.

[Grü+25] A. Grünhagen, M. Schütte, T. Lamb, S. Schulz, A. Eichler, M. Tropmann-Frick, G. Fey and H. Schlarb. *Impact of Seismic Activities on the Optical Synchronization System of the European X-ray Free-Electron Laser*. Submitted to High Power Laser Science and Engineering. 2025. DOI: 10.48550/arXiv.2502.02453. URL: http://arxiv.org/pdf/2502.02453v1.

[Gv13] G. H. Golub and C. F. van Loan. *Matrix computations*. 4th ed. Johns Hopkins studies in the mathematical sciences. Johns Hopkins University Press, 2013. ISBN: 978-1-4214-0794-4.

[Hän06] T. W. Hänsch. "Nobel Lecture: Passion for precision". In: *Reviews of Modern Physics* 78.4 (2006), pp. 1297–1309. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.78.1297.

[Hes18] J. P. Hespanha. *Linear systems theory*. Second edition. Princeton University Press, 2018. ISBN: 978-0691179575.

[Heu18] M. Heuer. "Identification and Control of the Laser-based Synchronization System for the European X-ray Free Electron Laser". Dissertation. Hamburg, Germany: Hamburg University of Technology, 2018. DOI: 10.3204/PUBDB-2018-02247.

[Hew85] Hewlett Packard. *Phase Noise Characterization of Microwave Oscillators: Frequency Discriminator Method*. 1985. URL: https://www.hpmemoryproject.org/an/pdf/pn11729C-2.pdf (visited on 01/02/2025).

[Hid+14] B. Hidding, G. G. Manahan, O. Karger, A. Knetsch, G. Wittig, D. A. Jaroszynski, Z.-M. Sheng, Y. Xi, A. Deng, J. B. Rosenzweig, G. Andonian, A. Murokh, G. Pretzler, D. L. Bruhwiler and J. Smith. "Ultrahigh brightness bunches from hybrid plasma accelerators as drivers of 5th generation light sources". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 47.23 (2014), p. 234010. DOI: 10.1088/0953-4075/47/23/234010.

[HJK21] M. Hyun, C.-G. Jeon and J. Kim. "Ultralow-noise microwave extraction from optical frequency combs using photocurrent pulse shaping with balanced photodetection". In: *Scientific reports* 11.1 (2021), p. 17809. DOI: 10.1038/s41598-021-97378-1.

[HP09] V. Hanta and A. Procházka. *Rational Approximation of Time Delay*. 2009. URL: https://www.researchgate.net/profile/Ales-Prochazka/publication/237396133_RATIONAL_APPROXIMATION_OF_TIME_DELAY (visited on 11/06/2024).

[HRS02] G. Heinzel, A. Rüdiger and R. Schilling. *Spectrum and spectral density estimation by the Discrete Fourier Transform (DFT): Including a comprehensive list of window functions and some new flat-top windows*. 2002. URL: https://hdl.handle.net/11858/00-001M-0000-0013-5579-7 (visited on 01/02/2025).

[HS83] M. Hautus and L. M. Silverman. "System structure and singular control". In: *Linear Algebra and its Applications* 50 (1983), pp. 369–402. ISSN: 0024-3795. DOI: 10.1016/0024-3795(83)90062-9.

[Hua+21] N. Huang, H. Deng, B. Liu, D. Wang and Z. Zhao. "Features and futures of X-ray free-electron lasers". In: *Innovation (Cambridge (Mass.))* 2.2 (2021), p. 100097. DOI: 10.1016/j.xinn.2021.100097.

[HWS18] S. Haesaert, S. Weiland and C. W. Scherer. "A separation theorem for guaranteed H_2 performance through matrix inequalities". In: *Automatica* 96 (2018), pp. 306–313. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2018.07.002.

[IEEa] IEEE. *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology – Random Instabilities*. DOI: 10.1109/IEEESTD.2022.9973001.

[IEEb] IEEE. *IEEE Standard for Jitter and Phase Noise*. DOI: 10.1109/IEEESTD.2021.9364950.

[Jae+20] E. J. Jaeschke, S. Khan, J. R. Schneider and J. B. Hastings. *Synchrotron Light Sources and Free-Electron Lasers*. Springer International Publishing, 2020. ISBN: 978-3-030-23200-9. DOI: 10.1007/978-3-030-23201-6.

[Jan+15] E. Janas, K. Czuba, M. Felber, M. Heuer, U. Mavric and H. Schlarb. "MTCA.4 Phase Detector for Femtosecond-Precision Laser Synchronization". In: *Proceedings of the 37th International Free Electron Laser Conference*. Daejeon, Korea: JACoW, 2015, pp. 474–477. ISBN: 978-3-95450-134-2. URL: https://bib-pubdb1.desy.de/record/292517/files/tup045.pdf (visited on 05/02/2025).

[JG01] M. Jun and M. G. Safonov. "Multiplier IQCs for uncertain time delays". In: *Journal of the Franklin Institute* 338.2-3 (2001), pp. 335–351. ISSN: 0016-0032. DOI: 10.1016/S0016-0032(00)00079-X.

[JM99] U. Jonsson and A. Megretski. "IQC characterizations of signal classes". In: *Proceedings of the 5th European Control Conference (ECC'99)*. Karlsruhe: IEEE, 1999, pp. 1481–1486. ISBN: 978-3-9524173-5-5. DOI: 10.23919/ECC.1999.7099521.

[JO75] A. Jameson and R. E. O'Malley. "Cheap Control of the Time-Invariant Regulator". In: *Applied Mathematics and Optimization* 1.4 (1975), pp. 337–354. ISSN: 0095-4616. DOI: 10.1007/BF01447957.

[Jön01] U. Jönsson. *Lecture Notes on Integral Quadratic Constraints*. Stockholm, Sweden, 2001.

[JS02] M. Jun and M. G. Safonov. "Rational multiplier IQCs for uncertain time-delays and LMI stability conditions". In: *IEEE Transactions on Automatic Control* 47.11 (2002), pp. 1871–1875. ISSN: 1558-2523. DOI: 10.1109/TAC.2002.804477.

[JVM11] E. Jarlebring, J. Vanbiervliet and W. Michiels. "Characterizing and Computing the H_2 Norm of Time-Delay Systems by Solving the Delay Lyapunov Equation". In: *IEEE Transactions on Automatic Control* 56.4 (2011), pp. 814–825. ISSN: 1558-2523. DOI: 10.1109/TAC.2010.2067510.

[Kal+21] C. Kalha, N. K. Fernando, P. Bhatt, F. O. L. Johansson, A. Lindblad, H. Rensmo, L. Z. Medina, R. Lindblad, S. Siol, L. P. H. Jeurgens, C. Cancellieri, K. Rossnagel, K. Medjanik, G. Schönhense, M. Simon, A. X. Gray, S. Nemšák, P. Lömker, C. Schlueter and A. Regoutz. "Hard x-ray photoelectron spectroscopy: a snapshot of the state-of-the-art in 2020". In: *Journal of physics. Condensed matter* 33.23 (2021). ISSN: 1361-648X. DOI: 10.1088/1361-648X/abeacd.

[KI09] F. Krausz and M. Ivanov. "Attosecond physics". In: *Reviews of Modern Physics* 81.1 (2009), pp. 163–234. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.81.163.

[Kim+08] J. Kim, J. A. Cox, J. Chen and F. X. Kärtner. "Drift-free femtosecond timing synchronization of remote optical and microwave sources". In: *Nature Photonics* 2.12 (2008), pp. 733–736. ISSN: 1749-4893. DOI: 10.1038/nphoton.2008.225.

[KKP04] J. Kim, F. X. Kärtner and M. H. Perrott. "Femtosecond synchronization of radio frequency signals with optical pulse trains". In: *Optics letters* 29.17 (2004), pp. 2076–2078. ISSN: 0146-9592. DOI: 10.1364/OL.29.002076.

[KL04] C.-Y. Kao and B. Lincoln. "Simple stability criteria for systems with time-varying delays". In: *Automatica* 40.8 (2004), pp. 1429–1434. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2004.03.011.

[Kol+16] S. Kolkowitz, I. Pikovski, N. Langellier, M. D. Lukin, R. L. Walsworth and J. Ye. "Gravitational wave detection with optical lattice atomic clocks". In: *Physical Review D* 94.12 (2016). ISSN: 2470-0010. DOI: 10.1103/PhysRevD.94.124043.

[KR07] C.-Y. Kao and A. Rantzer. "Stability analysis of systems with uncertain time-varying delays". In: *Automatica* 43.6 (2007), pp. 959–970. ISSN: 0005-1098. DOI: 10.1016/j.automatica.2006.12.006.

[KS16] J. Kim and Y. Song. "Ultralow-noise mode-locked fiber lasers and frequency combs: principles, status, and applications". In: *Advances in Optics and Photonics* 8.3 (2016), p. 465. DOI: 10.1364/AOP.8.000465.

[Lab+11] M. Labat, M. Bellaveglia, M. Bougeard, B. Carré, F. Ciocci, E. Chiadroni, A. Cianchi, M. E. Couprie, L. Cultrera, M. Del Franco, G. Di Pirro, A. Drago, M. Ferrario, D. Filippetto, F. Frassetto, A. Gallo, D. Garzella, G. Gatti, L. Giannessi, G. Lambert, A. Mostacci, A. Petralia, V. Petrillo, L. Poletto, M. Quattromini, J. V. Rau, C. Ronsivalle, E. Sabia, M. Serluca, I. Spassovsky, V. Surrenti, C. Vaccarezza and C. Vicario. "High-gain harmonic-generation free-electron laser seeded by harmonics generated in gas". In: *Physical review letters* 107.22 (2011), p. 224801. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.107.224801.

[Lam+11] T. Lamb, M.-K. Bock, M. Bousonville, M. Felber, P. Gessler, F. Ludwig, S. Ruzin, H. Schlarb, B. Schmidt, S. Schulz and E. Janas. "Femtosecond Stable Laser-To-RF Phase Detection Using Optical Modulators". In: *The 33rd International Free-Electron Laser Conference*. Shanghai, China, 2011, pp. 551–554. ISBN: 978-3-95450-117-5. URL: https://accelconf.web.cern.ch/FEL2011/papers/thpa32.pdf (visited on 12/02/2025).

[Lam+19] T. Lamb, M. Czwalinna, M. Felber, C. Gerth, T. Kozak, J. Müller, H. Schlarb, S. Schulz, C. Sydlo, M. Titberidze and F. Zummack. "Large-Scale Optical Synchronization System of the European XFEL with Femtosecond Precision". In: *Proceedings of the 10th International Particle Accelerator Conference (IPAC'19)*. Ed. by M. Boland, H. Tanaka and D. Button. Melbourne, Australia: JACoW, 2019, pp. 3835–3838. ISBN: 978-3-95450-208-0. DOI: 10.18429/JACOW-IPAC2019-THPRB018.

[Lam+24] T. Lamb, M. Schütte, F. Ludwig, H. Pryschelski, H. Schlarb and S. Schulz. "Reliable, Optically Augmented RF Reference Distribution with Femtosecond Stability". In: *LLRF Topical Workshop on Timing, Synchronization, Measurements and Calibration*. Frascati, Italy, 2024. URL: https://bib-pubdb1.desy.de/record/619043 (visited on 10/02/2025).

[Lam17] T. Lamb. *Laser-to-RF Phase Detection with Femtosecond Precision for Remote Reference Phase Stabilization in Particle Accelerators*. Deutsches Elektronen-Synchrotron (DESY), 2017. DOI: 10.3204/PUBDB-2017-02117.

[Lap17] A. Lapidoth. *A foundation in digital communication*. Second edition. Cambridge University Press, 2017. ISBN: 978-1-107-17732-1.

[LG17] J. Liu and J. Grünert. "Advanced temporal characterization of free electron laser pulses at European XFEL by THz photoelectron spectroscopy". In: *Terahertz, RF, Millimeter, and Submillimeter-Wave Technology and Applications X*. Ed. by L. P. Sadwick and T. Yang. SPIE Proceedings. San Francisco, California, United States: SPIE, 2017, 101030E. DOI: 10.1117/12.2253138.

[LHE24] J. Lübsen, C. Hespe and A. Eichler. "Towards safe multi-task Bayesian optimization". In: *Proceedings of the 6th Annual Learning for Dynamics & Control Conference*. Ed. by A. Abate, M. Cannon, K. Margellos and A. Papachristodoulou. Vol. 242. Proceedings of Machine Learning Research. ML Research Press, 2024, pp. 839–851. URL: https://proceedings.mlr.press/v242/lubsen24a.html.

[LHu24] A. L'Huillier. "Nobel Lecture: The route to attosecond pulses". In: *Reviews of Modern Physics* 96.3 (2024). DOI: 10.1103/RevModPhys.96.030503.

[Lin+98] Z. Lin, A. Saberi, P. Sannuti and Y. A. Shamash. "Direct Method of Constructing H_2-Suboptimal Controllers: Continuous-Time Systems". In: *Journal of Optimization Theory and Applications* 99.3 (1998), pp. 585–616. ISSN: 0022-3239. DOI: 10.1023/A:1021751016836.

[Lit+14] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, R. J. England, A. S. Fisher, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, N. Vafaei-Najafabadi, D. Walz, G. White, Z. Wu, V. Yakimenko and G. Yocky. "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator". In: *Nature* 515.7525 (2014), pp. 92–95. DOI: 10.1038/nature13882.

[Löf04] J. Löfberg. "YALMIP: A Toolbox for Modeling and Optimization in MATLAB". In: *Proceedings of the IEEE International Symposium on Computer Aided Control System Design (CACSD)*. Taipei, Taiwan: IEEE, 2004.

[Loh+13a] W. Loh, S. Yegnanarayanan, R. J. Ram and P. W. Juodawlkis. "Unified Theory of Oscillator Phase Noise I: White Noise". In: *IEEE Transactions on Microwave Theory and Techniques* 61.6 (2013), pp. 2371–2381. DOI: 10.1109/TMTT.2013.2260170.

[Loh+13b] W. Loh, S. Yegnanarayanan, R. J. Ram and P. W. Juodawlkis. "Unified Theory of Oscillator Phase Noise II: Flicker Noise". In: *IEEE Transactions on Microwave Theory and Techniques* 61.12 (2013), pp. 4130–4144. DOI: 10.1109/TMTT.2013.2288205.

[Löh09] F. Löhl. "Optical Synchronization of a Free-Electron Laser with Femtosecond Precision". Dissertation. Hamburg: University of Hamburg, 2009. DOI: 10.3204/DESY-THESIS-2009-031.

[LSA94] J. H. Ly, M. G. Safonov and F. Ahmad. "Positive real Parrott theorem with application to LMI controller synthesis". In: *Proceedings of 1994 American Control Conference (ACC '94)*. Baltimore, MD, USA: IEEE, 1994, pp. 50–52. ISBN: 978-0-7803-1783-3. DOI: 10.1109/ACC.1994.751691.

[Lüb+23] J. O. Lübsen, M. Schütte, S. Schulz and A. Eichler. "A Safe Bayesian Optimization Algorithm for Tuning the Optical Synchronization System at European XFEL". In: *IFAC-PapersOnLine* 56.2 (2023), pp. 3079–3085. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2023.10.1438.

[Lüb22] J. O. Lübsen. "Bayesian Optimization for the Control Parameters of the Optical Synchronization System at European XFEL". Master's Thesis. Hamburg, Germany: Technische Universität Hamburg, 2022. URL: https://bib-pubdb1.desy.de/record/485739 (visited on 05/02/2025).

[Lud+06] F. Ludwig, M. Hoffmann, H. Schlarb and S. Simrock. "Phase Stability of the Next Generation RF Field Control for VUV- and X-Ray Free-Electron Laser". In: *Proceedings of the 10th European Particle Accelerator Conference*. Edinburgh, Scotland: European Physical Society Accelerator Group, 2006, pp. 1453–1455. ISBN: 978-92-9083-278-2. URL: `https://accelconf.web.cern.ch/e06/PAPERS/TUPCH188.PDF` (visited on 04/02/2025).

[Mar+10] G. Marro, F. Morbidi, L. Ntogramatzidis and D. Prattichizzo. "Geometric control theory for linear systems: a tutorial". In: *19th International Symposium on Mathematical Theory of Networks and Systems (MTNS'10)*. Vol. 5. Budapest, 2010. ISBN: 978-963-311-370-7.

[Mar18] M. V. Markin. *Elementary functional analysis*. De Gruyter Graduate. De Gruyter, 2018. ISBN: 978-3-11-061391-9. DOI: `10.1515/9783110614039`.

[McG+19] W. F. McGrew, X. Zhang, H. Leopardi, R. J. Fasano, D. Nicolodi, K. Beloy, J. Yao, J. A. Sherman, S. A. Schäffer, J. Savory, R. C. Brown, S. Römisch, C. W. Oates, T. E. Parker, T. M. Fortier and A. D. Ludlow. "Towards the optical second: verifying optical clocks at the SI limit". In: *Physical review. X* 6.4 (2019). ISSN: 2160-3308. DOI: `10.1364/OPTICA.6.000448`.

[Mic19] S. Michizono. "The International Linear Collider". In: *Nature Reviews Physics* 1.4 (2019), pp. 244–245. DOI: `10.1038/s42254-019-0044-4`.

[MOS24] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB*. 2024. URL: `https://docs.mosek.com/10.2/toolbox/index.html`.

[MR97] A. Megretski and A. Rantzer. "System analysis via integral quadratic constraints". In: *IEEE Transactions on Automatic Control* 42.6 (1997), pp. 819–830. ISSN: 1558-2523. DOI: `10.1109/9.587335`.

[Mül+19] J. Müller, S. Schulz, L. Winkelmann, M. Czwalinna, I. Hartl and H. Schlarb. "Arrival Time Stabilization of the Photocathode Laser at the European XFEL". In: *Conference on Lasers and Electro-Optics*. San Jose, California: OSA, 2019, SF3I.6. ISBN: 978-1-943580-57-6. DOI: `10.1364/CLEO_SI.2019.SF3I.6`.

[Mur+22] S. D. Murthy, L. Doolittle, C. Xu, B. Hong, A. Benwell and J. Chen. *Installation, Commissioning and Performance of Phase Reference Line for LCLS-II*. 2022. DOI: `10.48550/arXiv.2210.05441`.

[NA05] D. Noll and P. Apkarian. "Spectral bundle methods for non-convex maximum eigenvalue functions: first-order methods". In: *Mathematical Programming* 104.2-3 (2005), pp. 701–727. ISSN: 0025-5610. DOI: `10.1007/s10107-005-0634-z`.

[Nak+20a] T. Nakamura, J. Davila-Rodriguez, H. Leopardi, J. A. Sherman, T. M. Fortier, X. Xie, J. C. Campbell, W. F. McGrew, X. Zhang, Y. S. Hassan, D. Nicolodi, K. Beloy, A. D. Ludlow, S. A. Diddams and F. Quinlan. "Coherent optical clock down-conversion for microwave frequencies with 10-18 instability". In: *Science (New York, N.Y.)* 368.6493 (2020), pp. 889–892. DOI: `10.1126/science.abb2473`.

[Nak+20b] T. Nakamura, S. Tani, I. Ito, M. Endo and Y. Kobayashi. "Piezo-electric transducer actuated mirror with a servo bandwidth beyond 500 kHz". In: *Optics express* 28.11 (2020), pp. 16118–16125. DOI: `10.1364/OE.390042`.

[Nar+22] N. V. Nardelli, T. M. Fortier, M. Pomponio, E. Baumann, C. Nelson, T. R. Schibli and A. Hati. "10 GHz generation with ultra-low phase noise via the transfer oscillator technique". In: *APL Photonics* 7.2 (2022). DOI: `10.1063/5.0073843`.

[nat] nature portfolio, ed. *Free-electron lasers: development and application*. URL: `https://www.nature.com/collections/gfcejeahgi` (visited on 01/02/2025).

[NM65] J. A. Nelder and R. Mead. "A Simplex Method for Function Minimization". In: *The Computer Journal* 7.4 (1965), pp. 308–313. ISSN: 0010-4620. DOI: 10.1093/comjnl/7.4.308.

[OF01] P. G. O'Shea and H. P. Freund. "Free-electron lasers. Status and applications". In: *Science (New York, N.Y.)* 292.5523 (2001), pp. 1853–1858. DOI: 10.1126/science.1055718.

[Pag99] F. Paganini. "Convex methods for robust H_2 analysis of continuous-time systems". In: *IEEE Transactions on Automatic Control* 44.2 (1999), pp. 239–252. ISSN: 1558-2523. DOI: 10.1109/9.746251.

[Pal80] Z. Palmor. "Stability properties of Smith dead-time compensator controllers". In: *International Journal of Control* 32.6 (1980), pp. 937–949. DOI: 10.1080/00207178008922900.

[Pas04] R. Paschotta. "Noise of mode-locked lasers (Part II): timing jitter and other fluctuations". In: *Applied Physics B* 79.2 (2004), pp. 163–173. ISSN: 0946-2171. DOI: 10.1007/s00340-004-1548-9.

[PD93] A. Packard and J. Doyle. "The complex structured singular value". In: *Automatica* 29.1 (1993), pp. 71–109. ISSN: 0005-1098. DOI: 10.1016/0005-1098(93)90175-S.

[PF00] F. Paganini and E. Feron. "Linear Matrix Inequality Methods for Robust H_2 Analysis: A Survey with Comparisons". In: *Advances in Linear Matrix Inequality Methods in Control*. Ed. by L. El Ghaoui and S.-l. Niculescu. Advances in Design and Control. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2000, pp. 129–151. ISBN: 978-0-89871-438-8. DOI: 10.1137/1.9780898719833.ch7.

[PF97] F. Paganini and E. Feron. "Analysis of robust H_2 performance: comparisons and examples". In: *Proceedings of the 36th IEEE Conference on Decision and Control (CDC'97)*. San Diego, CA, USA: IEEE, 1997, pp. 1000–1005. ISBN: 978-0-7803-4187-6. DOI: 10.1109/CDC.1997.657576.

[PP02] A. Papoulis and S. U. Pillai. *Probability, random variables, and stochastic processes*. 4. ed., internat. ed. McGraw-Hill series in electrical and computer engineering. McGraw-Hill, 2002. ISBN: 978-0073660110.

[PS12] İ. Polat and C. W. Scherer. "Stability Analysis for Bilateral Teleoperation: An IQC Formulation". In: *IEEE Transactions on Robotics* 28.6 (2012), pp. 1294–1308. ISSN: 1552-3098. DOI: 10.1109/TRO.2012.2209230.

[PS15] H. Pfifer and P. Seiler. *An Overview of Integral Quadratic Constraints for Delayed Nonlinear and Parameter-Varying Systems*. 2015. DOI: 10.48550/arXiv.1504.02502.

[PTV07] W. H. Press, S. A. Teukolsky and W. T. Vetterling. *Numerical Recipes: The art of scientific computing*. 3rd ed. Cambridge University Press, 2007. ISBN: 978-0521880688.

[PW09] D. B. Percival and A. T. Walden. *Spectral Analysis for Physical Applications*. Cambridge University Press, 2009. ISBN: 978-0-521-35532-2. DOI: 10.1017/CBO9780511622762.

[Qia22] J. Qiang. "X-ray FEL linear accelerator design via start-to-end global optimization". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1027.166294 (2022). ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.166294.

[Qui23] F. Quinlan. "The photodetection of ultrashort optical pulse trains for low noise microwave signal generation". In: *Laser & photonics reviews* 17.12 (2023). ISSN: 1863-8880. DOI: 10.1002/lpor.202200773.

[Rao72] C. R. Rao. "GENERALIZED INVERSE OF A MATRIX AND ITS APPLICATIONS". In: *Theory of Statistics*. Ed. by L. M. Le Cam, J. Neyman and E. L. Scott. University of California Press, 1972, pp. 601–620. ISBN: 978-0-520-32588-3. DOI: 10.1525/9780520325883-032.

[RH08] W. Riley and D. Howe. *Handbook of Frequency Stability Analysis*. Special Publication (NIST SP), National Institute of Standards and Technology, 2008. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=50505 (visited on 01/02/2024).

[Ric03] J.-P. Richard. "Time-delay systems: an overview of some recent advances and open problems". In: *Automatica* 39.10 (2003), pp. 1667–1694. ISSN: 0005-1098. DOI: 10.1016/S0005-1098(03)00167-5.

[RL06] M. Rotkowitz and S. Lall. "A Characterization of Convex Problems in Decentralized Control". In: *IEEE Transactions on Automatic Control* 51.2 (2006), pp. 274–286. ISSN: 1558-2523. DOI: 10.1109/TAC.2005.860365.

[RM12] M. C. Rotkowitz and N. C. Martins. "On the Nearest Quadratically Invariant Information Constraint". In: *IEEE Transactions on Automatic Control* 57.5 (2012), pp. 1314–1319. ISSN: 1558-2523. DOI: 10.1109/TAC.2011.2173411.

[Rou+24] R. Roussel, A. L. Edelen, T. Boltz, D. Kennedy, Z. Zhang, F. Ji, X. Huang, D. Ratner, A. S. Garcia, C. Xu, J. Kaiser, A. F. Pousa, A. Eichler, J. O. Lübsen, N. M. Isenberg, Y. Gao, N. Kuklev, J. Martinez, B. Mustapha, V. Kain, C. Mayes, W. Lin, S. M. Liuzzo, J. St. John, M. J. V. Streeter, R. Lehe and W. Neiswanger. "Bayesian optimization algorithms for accelerator physics". In: *Physical Review Accelerators and Beams* 27.8 (2024). DOI: 10.1103/PhysRevAccelBeams.27.084801.

[Rub08] E. Rubiola. *Phase Noise and Frequency Stability in Oscillators*. Cambridge University Press, 2008. ISBN: 978-0-511-81279-8. DOI: 10.1017/CBO9780511812798.

[Ryb+17] R. Rybaniec, K. Przygoda, W. Cichalewski, V. Ayvazyan, J. Branlard, L. Butkowski, S. Pfeiffer, C. Schmidt, H. Schlarb and J. Sekutowicz. "FPGA-Based RF and Piezocontrollers for SRF Cavities in CW Mode". In: *IEEE Transactions on Nuclear Science* 64.6 (2017), pp. 1382–1388. ISSN: 0018-9499. DOI: 10.1109/TNS.2017.2687981.

[Sch+03] T. R. Schibli, J. Kim, O. Kuzucu, J. T. Gopinath, S. N. Tandon, G. S. Petrich, L. A. Kolodziejski, J. G. Fujimoto, E. P. Ippen and F. X. Kaertner. "Attosecond active synchronization of passively mode-locked lasers by balanced cross correlation". In: *Optics letters* 28.11 (2003), pp. 947–949. ISSN: 0146-9592. DOI: 10.1364/OL.28.000947.

[Sch+11] S. Schreiber, M. Görler, G. Klemz, K. Klose, G. Koss, T. Schulz, M. Staack, I. Will, I. Templin and H. Willert. "Operation of the FLASH Photoinjector Laser System". In: *The 33rd International Free-Electron Laser Conference*. Shanghai, China, 2011, pp. 507–510. ISBN: 978-3-95450-117-5. URL: https://epaper.kek.jp/FEL2011/papers/thpa18.pdf (visited on 03/02/2015).

[Sch+13] H. Schlarb, T. Walter, K. Rehlich and F. Ludwig. "Novel Crate Standard MTCA.4 for Industry and Research". In: *Proceedings of 4th International Particle Accelerator Conference (IPAC'13)*. Ed. by Z. Dai. Shanghai, China: JACoW, 2013, pp. 3633–3635. ISBN: 978-3-95450-122-9. URL: https://accelconf.web.cern.ch/ipac2013/papers/thpwa003.pdf (visited on 05/02/2025).

[Sch+14] P. Schmüser, M. Dohlus, J. Rossbach and C. Behrens. *Free-Electron lasers in the ultraviolet and x-ray regime: Physical principles, experimental results, technicals realization*. 2. ed. Vol. 258. Springer tracts in modern physics. Springer, 2014. ISBN: 978-3-319-04080-6. DOI: 10.1007/978-3-319-04081-3.

[Sch+15] S. Schulz, I. Grguraš, C. Behrens, H. Bromberger, J. T. Costello, M. K. Czwalinna, M. Felber, M. C. Hoffmann, M. Ilchen, H. Y. Liu, T. Mazza, M. Meyer, S. Pfeiffer, P. Predki, S. Schefer, C. Schmidt, U. Wegner, H. Schlarb and A. L. Cavalieri. "Femtosecond all-optical synchronization of an X-ray free-electron laser". In: *Nature Communications* 6 (2015), p. 5938. DOI: 10.1038/ncomms6938.

[Sch+21] M. Schütte, A. Eichler, H. Schlarb, G. Lichtenberg and H. Werner. "Decentralized Output Feedback Control using Sparsity Invariance with Application to Synchronization at European XFEL". In: *Proceedings of the 60th IEEE Conference on Decision and Control (CDC'21)*. Austin, TX, USA: IEEE, 2021, pp. 5723–5728. ISBN: 978-1-6654-3659-5. DOI: 10.1109/CDC45484.2021.9683027.

[Sch+22] M. Schütte, A. Eichler, H. Schlarb, G. Lichtenberg and H. Werner. "Convex Synthesis of Robust Distributed Controllers for the Optical Synchronization System at European XFEL". In: *Proceedings of the 6th IEEE Conference on Control Technology and Applications (CCTA'22)*. Trieste, Italy: IEEE, 2022, pp. 1061–1067. ISBN: 978-1-6654-7338-5. DOI: 10.1109/CCTA49430.2022.9966114.

[Sch+23] C. Schmidt, M. Bousonville, J. Branlard, M. Diomede, S. Göller, D. Kostin, M. Scholz, V. Vogel and N. Walker, eds. *Operational Experience with the European XFEL SRF Linac: JACoW Publishing, Geneva, Switzerland*. 2023. DOI: 10.18429/JACoW-SRF2023-MOIXA06.

[Sch+24] D. Schwickert, S. Alisauskas, N. Kschuev, A.-L. Calendron, A. T. Seifi, G. Cirmi, S. Düsterer, H. Cankaya, I. Hartl, S. Schulz and H. Schlarb. "Sub-15 fs Jitter After Multi-Pass Cell Pulse Compression at the Beamline FL23 of the FLASH Facility". In: *Laser Congress 2024 (ASSL, LAC, LS&C)*. Osaka, Japan: Optica Publishing Group, 2024, AW1A.4. ISBN: 978-1-957171-42-5. DOI: 10.1364/ASSL.2024.AW1A.4.

[Sch00] C. W. Scherer. "Robust Mixed Control and Linear Parameter-Varying Control with Full Block Scalings". In: *Advances in Linear Matrix Inequality Methods in Control*. Ed. by L. El Ghaoui and S.-l. Niculescu. Advances in Design and Control. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2000, pp. 187–207. ISBN: 978-0-89871-438-8. DOI: 10.1137/1.9780898719833.ch10.

[Sch05] H. Schlarb. "Techniques for Pump-Probe Synchronisation of FSEC Radiation Pulses". In: *Proceedings of the 2005 Particle Accelerator Conference*. Knoxville, TN, USA: IEEE, 2005, pp. 59–63. DOI: 10.1109/PAC.2005.1590361.

[Sch24] M. Schütte. "Phase and Amplitude Noise Measurements: Fundamentals and Best Practices". In: *LLRF Topical Workshop on Timing, Synchronization, Measurements and Calibration*. Frascati, Italy, 2024. URL: https://bib-pubdb1.desy.de/record/618925 (visited on 01/02/2025).

[Sch96] C. W. Scherer. "Mixed H_2/H_∞ control for time-varying and linear parametrically-varying systems". In: *International Journal of Robust and Nonlinear Control* 6.9-10 (1996), pp. 929–952. ISSN: 1049-8923. DOI: 10.1002/(SICI)1099-1239(199611)6:9/10%3C929::AID-RNC260%3E3.0.CO;2-9.

[SCS24] S. Schulz, M. K. Czwalinna and H. Schlarb. "Latest achievements in femtosecond synchronization of large-scale facilities". In: *Proceedings of 13th International Beam Instrumentation Conference*. Beijing, China: JACoW, 2024, pp. 457–464. ISBN: 978-3-95450-249-3. DOI: 10.18429/JACoW-IBIC2024-THAT1. URL: https://epaper.kek.jp/ibic2024/pdf/THAT1.pdf (visited on 04/02/2025).

[SCS95] A. Saberi, B. M. Chen and P. Sannuti. *H_2 optimal control*. Prentice Hall, 1995. ISBN: 978-0-13-489782-0.

[Šer+26] A. Šerlat, M. Grzegrzółka, M. Urbański, B. GąSowski, K. Czuba, D. Sikora, R. Papis, P. Jatczak, I. Rutkowski, J. Holzbauer and N. Eddy. "Phase Drifts and Signal Dispersion in

Coaxial Cables”. In: *IEEE Access* 14 (2026), pp. 40261–40271. DOI: 10.1109/ACCESS.2026.3670751.

[SEW23] M. Schütte, A. Eichler and H. Werner. “Structured IQC Synthesis of Robust H_2 Controllers in the Frequency Domain”. In: *IFAC-PapersOnLine* 56.2 (2023), pp. 10408–10413. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2023.10.1055.

[SF15] S. Schreiber and B. Faatz. “The free-electron laser FLASH”. In: *High Power Laser Science and Engineering* 3 (2015). ISSN: 2095-4719. DOI: 10.1017/hpl.2015.16.

[SG18] S. Skogestad and C. Grimholt. “Should we forget the Smith Predictor?” In: *IFAC-PapersOnLine* 51.4 (2018), pp. 769–774. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2018.06.203.

[SG22] S. Simrock and Z. Geng. *Low-level radio frequency systems*. Particle acceleration and detection. Springer, 2022. ISBN: 978-3-030-94418-6. DOI: 10.1007/978-3-030-94419-3.

[SGC97] C. Scherer, P. Gahinet and M. Chilali. “Multiobjective output-feedback control via LMI optimization”. In: *IEEE Transactions on Automatic Control* 42.7 (1997), pp. 896–911. ISSN: 1558-2523. DOI: 10.1109/9.599969.

[Sha+23] K. Shao, S. Liu, P. Gao, Y. Zhang, Z. Xu, H. Wang and S. Pan. “All-Polarization-Maintained Microwave Photonic Phase Detector Based on a DP-DMZM”. In: *IEEE Photonics Technology Letters* 35.7 (2023), pp. 385–388. ISSN: 1041-1135. DOI: 10.1109/LPT.2023.3246721.

[Shi14] S. V. Shinde. “Review of Oscillator Phase Noise Models”. In: *Proceedings of the International MultiConference of Engineers and Computer Scientists 2014*. Ed. by IAENG. Hong Kong, 2014.

[Sik+20] D. Sikora, K. Czuba, P. Jatczak, M. Urbanski, H. Schlarb, F. Ludwig and H. Pryschelski. “Phase Drift Compensating RF Link for Femtosecond Synchronization of E-XFEL”. In: *IEEE Transactions on Nuclear Science* 67.9 (2020), pp. 2136–2142. ISSN: 0018-9499. DOI: 10.1109/TNS.2020.2966018.

[Sim06] S. Simrock. “Low-level radio frequency control system for the European XFEL”. In: *Proceedings of the 2006 International Conference Mixed Design of Integrated Circuits and Systems (MIXDES’06)*. Gdynia, Poland: IEEE, 2006, pp. 79–84. DOI: 10.1109/MIXDES.2006.1706542.

[SJZ00] Stone, Jack A. and J. H. Zimmerman. *Engineering Metrology Toolbox - Index of Refraction of Air: Based on Ciddor Equation*. 2000. URL: https://emtoolbox.nist.gov/Wavelength/Ciddor.asp (visited on 05/02/2025).

[SL99] T. Sato and K.-Z. Liu. “LMI solution to general suboptimal control problems”. In: *Systems & Control Letters* 36.4 (1999), pp. 295–305. ISSN: 0167-6911. DOI: 10.1016/S0167-6911(98)00102-9.

[Smi22] K. G. Smith. *The Case of the Spurious Phase Noise: Part I*. 2022. URL: https://www.skyworksinc.com/Blog/2022/09/Timing-101-06 (visited on 11/06/2024).

[Smi57] O. J. M. Smith. “Closer control of loops with dead time”. In: *Chemical Engineering Progress* 53 (1957), pp. 217–219.

[SP09] S. Skogestad and I. Postlethwaite. *Multivariable feedback control: Analysis and design*. 2nd ed. Wiley, 2009. ISBN: 978-0-470-01168-3.

[Sri+12] N. Srinivas, A. Krause, S. M. Kakade and M. W. Seeger. “Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting”. In: *IEEE Transactions on Information Theory* 58.5 (2012), pp. 3250–3265. ISSN: 0018-9448. DOI: 10.1109/TIT.2011.2182033.

[SSY00] E. L. Saldin, E. A. Schneidmiller and M. V. Yurkov. *The Physics of Free Electron Lasers*. Advanced Texts in Physics. Springer, 2000. ISBN: 978-3-642-08555-0. DOI: 10.1007/978-3-662-04066-9.

[ST00] M. Sznaier and J. Tierno. "Is set modeling of white noise a good tool for robust H_2 analysis?" In: *Automatica* 36.2 (2000), pp. 261–267. ISSN: 0005-1098. DOI: 10.1016/S0005-1098(99)00137-5.

[Ste13] J. D. Stefanovski. "Algorithms for optimal control with invariant zeros on the extended imaginary axis". In: *Systems & Control Letters* 62.7 (2013), pp. 587–596. ISSN: 0167-6911. DOI: 10.1016/j.sysconle.2013.03.009.

[Ste97] J. L. Stensby. *Phase-locked loops: Theory and applications*. CRC Press, 1997. ISBN: 978-0-8493-9471-3.

[Sto92] A. A. Stoorvogel. "The singular H_2 control problem". In: *Automatica* 28.3 (1992), pp. 627–631. ISSN: 0005-1098. DOI: 10.1016/0005-1098(92)90189-M.

[Sto93] A. A. Stoorvogel. "The robust H_2 control problem: a worst-case design". In: *IEEE Transactions on Automatic Control* 38.9 (1993), pp. 1358–1371. ISSN: 1558-2523. DOI: 10.1109/9.237647.

[SV12] C. W. Scherer and J. Veenman. "Robust Controller Synthesis is Convex for Systems without Control Channel Uncertainties". In: *IFAC Proceedings Volumes* 45.13 (2012), pp. 325–330. ISSN: 1474-6670. DOI: 10.3182/20120620-3-DK-2025.00110.

[SW17] C. W. Scherer and S. Weiland. *Linear Matrix Inequalities in Control*. 2017. URL: https://www.researchgate.net/profile/Carsten-Scherer/publication/228731368_Linear_Matrix_Inequality_in_Control/ (visited on 11/06/2024).

[The24a] The MathWorks, Inc. *MATLAB*. Natick, Massachusetts, United States, 2024.

[The24b] The MathWorks, Inc. *MATLAB: Control System Toolbox*. Natick, Massachusetts, United States, 2024.

[The24c] The MathWorks, Inc. *MATLAB: Robust Control Toolbox*. Natick, Massachusetts, United States, 2024.

[Tit17] M. Titberidze. *Pilot Study of Synchronization on a Femtosecond Scale between the Electron Gun REGAE and a Laser-Plasma Accelerator*. Deutsches Elektronen-Synchrotron (DESY), 2017. DOI: 10.3204/PUBDB-2017-11374.

[TSH01] H. L. Trentelman, A. A. Stoorvogel and M. Hautus. *Control Theory for Linear Systems*. 1st ed. 2001. Communications and Control Engineering. Springer, 2001. ISBN: 978-14471-1073-6. DOI: 10.1007/978-1-4471-0339-4.

[UHH02] T. Udem, R. Holzwarth and T. W. Hänsch. "Optical frequency metrology". In: *Nature* 416.6877 (2002), pp. 233–237. DOI: 10.1038/416233a.

[Vaj00] M. Vajta. "Some remarks on Padé approximations". In: *Proceedings of the 3rd TEMPUS-INTCOM Symposium* 242 (2000), pp. 1–6.

[van28] K. S. van Dyke. "The Piezo-Electric Resonator and Its Equivalent Network". In: *Proceedings of the IRE* 16.6 (1928), pp. 742–764. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1928.221466.

[Vas+21] M. Vassholz, H. P. Hoeppe, J. Hagemann, J. M. Rosselló, M. Osterhoff, R. Mettin, T. Kurz, A. Schropp, F. Seiboth, C. G. Schroer, M. Scholz, J. Möller, J. Hallmann, U. Boesenberg, C. Kim, A. Zozulya, W. Lu, R. Shayduk, R. Schaffer, A. Madsen and T. Salditt. "Pump-probe X-ray holographic imaging of laser-induced cavitation bubbles with femtosecond FEL pulses". In: *Nature Communications* 12.1 (2021), p. 3468. DOI: 10.1038/s41467-021-23664-1.

[Vee+21] J. Veenman, C. W. Scherer, C. Ardura, S. Bennani, V. Preda and B. Girouart. "IQClab: A new IQC based toolbox for robustness analysis and control design". In: *IFAC-PapersOnLine* 54.8 (2021), pp. 69–74. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2021.08.583.

[Vog+11] M. Vogt, B. Faatz, J. Feldhaus, Honkavaara Katja, S. Schreiber and R. Treusch. "Status of the Free-Electron Laser FLASH at DESY". In: *IPAC 2011 contributions to the proceedings*. Kursaal, Spain: JACoW, 2011, pp. 3080–3082. ISBN: 978-92-9083-366-6. URL: https://accelconf.web.cern.ch/IPAC2011/papers/THPC081.PDF (visited on 03/02/2025).

[Vog+23] M. Vogt, E. Ferrari, C. Gerth, K. Honkavaara, J. Rönsch-Schulenburg, L. Schaper, S. Schreiber and J. Zemella, eds. *The FLASH 2020+ Upgrade Project: JACoW Publishing, Geneva, Switzerland*. 2023. DOI: 10.18429/JACoW-SRF2023-TUIAA02.

[VS14] J. Veenman and C. W. Scherer. "IQC-synthesis with general dynamic multipliers". In: *International Journal of Robust and Nonlinear Control* 24.17 (2014), pp. 3027–3056. ISSN: 1049-8923. DOI: 10.1002/rnc.3042.

[VSK16] J. Veenman, C. W. Scherer and H. Köroğlu. "Robust stability and performance analysis based on integral quadratic constraints". In: *European Journal of Control* 31 (2016), pp. 1–32. ISSN: 0947-3580. DOI: 10.1016/j.ejcon.2016.04.004.

[Wan+23] X. Wang, Y. Jin, S. Schmitt and M. Olhofer. "Recent Advances in Bayesian Optimization". In: *ACM Computing Surveys* 55.13s (2023), pp. 1–36. ISSN: 0360-0300. DOI: 10.1145/3582078.

[Wer21] H. Werner. *Optimal and Robust Control (Lecture Notes)*. 2021. URL: https://collaborating.tuhh.de/ICS/ics-public/lecture-files/-/raw/master/ORC/OptimalAndRobustControl.pdf (visited on 01/02/2025).

[WF05] H. Werner and A. Farag. "A Hybrid Evolutionary-Algebraic Approach to Optimal and Robust Controller Design". In: *at - Automatisierungstechnik* 53.11 (2005), pp. 546–555. DOI: 10.1524/auto.2005.53.11.546.

[WFS24] Y. Watanabe, S. Fushimi and K. Sakurama. *Convex Reformulation of LMI-Based Distributed Controller Design with a Class of Non-Block-Diagonal Lyapunov Functions*. 2024. DOI: 10.48550/arXiv.2404.04576.

[Wil+11] I. Will, H. I. Templin, S. Schreiber and W. Sandner. "Photoinjector drive laser of the FLASH FEL". In: *Optics express* 19.24 (2011), pp. 23770–23781. DOI: 10.1364/OE.19.023770.

[Wil+17] T. Wilksen, A. Aghababyan, R. Bacher, P. Bartkiewicz, C. Behrens, T. Delfs, P. Duval, L. Fröhlich, W. Gerhardt, C. Gindler, O. Hensler, K. Hinsch, J. JÃ¤Ger, R. Kammering, S. Karstensen, H. Kay, V. Kocharyan, A. Labudda, T. Limberg, S. Meykopff, A. Petrosyan, G. Petrosyan, L. Petrosyan, V. Petrosyan, P. Pototzki, K. Rehlich, G. Schlesselmann, E. Sombrowski, M. Staack, J. Szczesny, M. Walla, J. Wilgen and H. Wu. "The Control System for the Linear Accelerator at the European XFEL: Status and First Experiences". In: *Proceedings of the 16th International Conference on Accelerators and Large Experimental Physics Control Systems (ICALEPCS'17)*. Ed. by I. Costa, D. Fernández, Ó. Matilla and V. R. W. Schaa. Barcelona, Spain: JACoW, 2017, pp. 1–5. ISBN: 978-3-95450-193-9. DOI: 10.18429/JACoW-ICALEPCS2017-MOAPL01.

[Wil92] J. E. D. Williams. *From sails to satellites: The origin and development of navigational science*. Oxford Univ. Press, 1992. ISBN: 978-0-19-856387-7.

[Win98] H. Winick. "Fourth generation light sources". In: *Proceedings of the 1997 Particle Accelerator Conference*. Vancouver, BC, Canada: IEEE, 1998, pp. 37–41. ISBN: 978-0-7803-4376-4. DOI: 10.1109/PAC.1997.749539.

[WKS86] J. C. Willems, A. Kitapçi and L. M. Silverman. "Singular Optimal Control: A Geometric Approach". In: *SIAM Journal on Control and Optimization* 24.2 (1986), pp. 323–337. ISSN: 0363-0129. DOI: 10.1137/0324018.

[WLS18] Y. Wang, J. A. Lopez and M. Sznaier. "Convex Optimization Approaches to Information Structured Decentralized Control". In: *IEEE Transactions on Automatic Control* 63.10 (2018), pp. 3393–3403. ISSN: 1558-2523. DOI: 10.1109/TAC.2018.2830112.

[Won68] W. M. Wonham. "On the Separation Theorem of Stochastic Control". In: *SIAM Journal on Control* 6.2 (1968), pp. 312–326. ISSN: 0036-1402. DOI: 10.1137/0306023.

[Won85] W. M. Wonham. *Linear multivariable control: A geometric approach*. 3rd ed. Vol. 10. Applications of mathematics. Springer, 1985. ISBN: 978-0-387-96071-5.

[Xie+17] X. Xie, R. Bouchand, D. Nicolodi, M. Giunta, W. Hänsel, M. Lezius, A. Joshi, S. Datta, C. Alexandre, M. Lours, P.-A. Tremblin, G. Santarelli, R. Holzwarth and Y. Le Coq. "Photonic microwave signals with zeptosecond-level absolute timing noise". In: *Nature Photonics* 11.1 (2017), pp. 44–47. ISSN: 1749-4885. DOI: 10.1038/nphoton.2016.215.

[Yan+23] J. Yan, S. Serkez, C. Lechner, E. A. Schneidmiller, C. Heyl, M. Guetg, G. Geloni and Y. Chen, eds. *Simulation Studies for the ASPECT Project at European XFEL: JACoW Publishing*. 2023. DOI: 10.18429/JACoW-FEL2022-MOP22.

[YBJ76] D. Youla, J. Bongiorno and H. Jabr. "Modern Wiener–Hopf design of optimal controllers Part I: The single-input-output case". In: *IEEE Transactions on Automatic Control* 21.1 (1976), pp. 3–13. ISSN: 1558-2523. DOI: 10.1109/TAC.1976.1101139.

[YP93] Young and Peter M. "Robustness with Parametric and Dynamic Uncertainty". Dissertation. Pasadena, California: California Institute of Technology, 1993. URL: https://thesis.library.caltech.edu/4695/1/Young_pm_1993.pdf (visited on 01/02/2025).

[Zas+21] U. Zastrau, K. Appel, C. Baehtz, O. Baehr, L. Batchelor, A. Berghäuser, M. Banjafar, E. Brambrink, V. Cerantola, T. E. Cowan, H. Damker, S. Dietrich, S. Di Dio Cafiso, J. Dreyer, H. O. Engel, T. Feldmann, S. Findeisen, M. Foese, D. Fulla-Marsa, S. Göde, M. Hassan, J. Hauser, T. Herrmannsdörfer, H. Höppner, J. Kaa, P. Kaever, K. Knöfel, Z. Konôpková, A. Laso García, H. P. Liermann, J. Mainberger, M. Makita, E. C. Martens, E. E. McBride, D. Möller, M. Nakatsutsumi, A. Pelka, C. Plueckthun, C. Prescher, T. R. Preston, M. Röper, A. Schmidt, W. Seidel, J. P. Schwinkendorf, M. O. Schoelmerich, U. Schramm, A. Schropp, C. Strohm, K. Sukharnikov, P. Talkovski, I. Thorpe, M. Toncian, T. Toncian, L. Wollenweber, S. Yamamoto and T. Tschentscher. "The High Energy Density Scientific Instrument at the European XFEL". In: *Journal of synchrotron radiation* 28.Pt 5 (2021), pp. 1393–1416. DOI: 10.1107/S1600577521007335.

[ZDG96] K. Zhou, J. C. Doyle and K. Glover. *Robust and optimal control*. Prentice Hall, 1996. ISBN: 978-0-13-456567-5.