

DATA-DRIVEN AND MODEL-BASED METHODS FOR VERIFICATION AND CONTROL OF PHYSICAL SYSTEMS

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus prof.dr.ir. F.P.T. Baaijens, voor een
commissie aangewezen door het College voor
Promoties, in het openbaar te verdedigen op
donderdag 16 februari 2017 om 16:00 uur

door

Sofie Haesaert

geboren te Leuven, België

Dit proefschrift is goedgekeurd door de promotoren en de samenstelling van de promotiecommissie is als volgt:

voorzitter: prof.dr.ir. A.B. Smolders
1^e promotor: prof.dr.ir. P.M.J. Van den Hof
copromotor: dr. A. Abate (University of Oxford)
leden: prof.dr.ir. J. Schoukens (Vrije Universiteit Brussel)
prof.dr.ir. A. Rantzer (Lund University of Technology)
prof.dr.ir. J.F. Groote
prof.dr. K.G. Larsen (Aalborg Universitet)
prof.dr. S. Weiland

Het onderzoek dat in dit proefschrift wordt beschreven is uitgevoerd in overeenstemming met de TU/e Gedragscode Wetenschapsbeoefening

Data-driven and
model-based methods for
verification and control of
physical systems

Sofie Haesaert



This dissertation has been completed in fulfillment of the requirements of the Dutch Institute of Systems and Control (DISC) for graduate study.



Netherlands Organisation for Scientific Research

This work is part of the research programme “DISC Graduate Programme” with project number 022.002.025, which is financed by the Netherlands Organisation for Scientific Research (NWO).

A catalogue record is available from the Eindhoven University of Technology Library
ISBN: 978-90-386-4224-6

Copyright © 2017 by Sofie Haesaert.

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilised in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the copyright owner.

Contents

1	Introduction	1
1.1	Formal methods in control engineering	1
1.2	Research questions	7
1.3	Thesis outline and publications	10
2	Verification and control of physical systems: a prelude	13
2.1	Physical systems	13
2.2	Technical preliminaries	15
2.3	Deterministic framework	16
2.4	Probabilistic framework	21
I	Control	25
3	Approximate similarity relations and controller refinement for general Markov decision processes	27
3.1	Introduction	27
3.2	Verification of general Markov decision processes: problem setup .	29
3.3	Exact (bi-)simulation relations based on lifting	32
3.4	New ϵ, δ -approximate (bi-)simulation relations via lifting	38
3.5	Case studies	43
3.6	Connections to literature and measurability issues	51
3.7	Conclusions	57
Appendices		
	List of symbols	58
	List of definitions	59
3.A	Details on case study and use of FAUST ²	60
3.B	Proofs of Theorems and Corollaries	63

4	Output-based controller synthesis	69
4.1	Introduction: correct-by-design control	69
4.2	State-based correct-by-design controller synthesis	74
4.3	The idea: output-based correct-by-design control	78
4.4	Output-based correct-by-design control	80
4.5	Case study in smart buildings	87
4.6	Conclusions	90
Appendices		
	List of symbols	91
5	A separation theorem through matrix inequalities	93
5.1	Introduction	94
5.2	Problem statement for guaranteed H_2 performance	95
5.3	Background: optimal control and linear matrix inequalities	98
5.4	Separation theorem via matrix inequalities	103
5.5	Corollaries for H_2 performance	112
5.6	Problem statement for hierarchical control	114
5.7	Computational approach to output-based control refinement	118
5.8	Future work and conclusions	125
Appendices		
	5.A Operations on linear matrix inequalities	125
II	Verification	127
6	Data-driven and model-based verification	129
6.1	Introduction	129
6.2	General framework and problem statement	131
6.3	LTL verification of LTI systems	134
6.4	Discussion on the generalisation of the results	152
6.5	Conclusions	153
Appendices		
	6.A Derivation of the bounds in Section 6.3.e	154

7	Bayesian experiment design for formal verification	157
7.1	Introduction	158
7.2	Bayesian experiment design for data-driven verification	160
7.3	Experiment design solutions in literature	164
7.4	Maximising a-posteriori confidence (offline)	169
7.5	Maximising probabilistic reachability of confidence threshold (online)	181
7.6	Conclusions	200
Appendices		
	List of symbols	201
7.A	Derivation of Bayesian inference formulae	202
8	Conclusions and future work	207
8.1	Overview & summary	207
8.2	Results	208
8.3	Future work	215
	Bibliography	217
	Summary	231
	Acknowledgments	233
	Curriculum Vitae	235

If a man will begin with certainties, he will end in doubts; but if he will be content to begin with doubts, he will end in certainties.

Francis Bacon (1561-1626), *Advancement of Learning*

1

Introduction

The design of controllers for complex, high-tech, safety-critical systems such as autonomous vehicles, intelligent robots, and cyber-physical infrastructures, demands guarantees on their correct and reliable behaviour. It is in particular of interest to synthesise the part of the computer software that controls or interacts with the physical system automatically, with low likelihood of malfunctioning. Correct functioning and reliability of systems can be attained by the use of formal methods, known from computer science. These available rigorous specification frameworks and automatised mathematical proof techniques are often bound to the requirement of having access to the system's model and state evolutions at all times. In practice, for most physical systems their evolution in time is only known in part and information can only be gathered via noisy measurements. Motivated by the need of formal methods in control engineering that still work under these circumstances, we formulate the research questions tackled in this thesis.

1.1 Formal methods in control engineering: motivation

Control engineering is the enhancement of physical systems with stabilising or governing modules. A classical example is the fly-ball governor invented in 1788 [Franklin et al., 2001]. This part of a steam engine is a purely mechanical system that counteracts irregularities in the speed of the engine by regulating the valve

of a steam engine. The automatic correction of the operation of physical systems is referred to as control engineering and, still pervasive in virtually all fields of engineering, it has grown with the digital age. Control systems are currently implemented digitally and among others used to regulate complex industrial processes, control the power grid, and increase the precision of mechatronic designs. These control systems measure variables of the physical systems and send an appropriate input signal to the actuators. Recent scientific advances are pushing the size of sensors down and are leading to decreasing costs of computation power. Consequently an ubiquitous embedding of digital components has been observed. Leaning on this increased amount of information and computational power available for real-time decision making, new and more intelligent autonomous systems are expected to go beyond the state of the art and potentially replace humans in many complex control tasks.

Recent history contains numerous examples where small oversights in control software of physical systems have had big impact [Kreiker et al., 2011], causing not only economic losses but, in some cases, also human casualties [Zhivich and Cunningham, 2009]. With the digitalisation of the modern society, repercussions of software bugs reach beyond the digital domain. When errors appear in software code of control systems, whether or not they are due to mistakes in the controller design or pure software errors, they often have grave consequences. A pipeline rupture in 1999 resulted in the death of 3 persons and property damage of at least 45 million US Dollars. The malfunctioning of its supervisory control and data acquisition (SCADA) system was one of the major causes of this disaster. The increasing interleaving of the physical domain and the digital or cyber one pushes this type of errors to the forefront. Avoiding faults or malfunctions is especially key for safety-critical systems such as vehicles, infrastructures, and systems in health-care. Unlike when the first control systems were designed, correct functioning of control systems can no longer be solely defined by simple, local requirements on noise attenuation and stability. Instead, in support of further automatisation, requirements should define the functioning of the system in its entirety. Consider as examples: the sequential and conditional performance of tasks, often present in robot control; or the capability to recover from sensor faults, as demanded in the automotive and aviation domains. Thus while an increase in automatisation is expected to reduce the number of human made errors during employment, the requirements posed on these automated processes are more complex and the required guarantees more crucial. Additionally the automatisation and the related increase in complexity of the regulatory system is generally discordant with the requirements of reliability. As a consequence, we expect an accumulation of (human made) errors in the design phase. To counter-act design errors, control engineers need design tools to support reliability in the design of controllers by means of verification and automation of control design synthesis.

Let us take a small detour in recent automotive history. The car industry has been exemplar for the ubiquitous embedding of digital modules; mechanical and electric circuits in cars have been increasingly replaced by digital components with wireless transmission [Transportation Research Board et al., 2012]. The transport of humans is inherently safety critical. Hence this has been also one of the

first engineering domains where the need for consistent designs and tools verifying the behaviour of control modules was realised. The necessity for this was also emphasised during an investigation into the electronic throttle control system (ETC) done by the Transportation Research Board et al. [2012] in the US. In current cars the classical mechanical throttle systems have been replaced by an electronic throttle control system. Around 2009 Toyota Motor Corporation was forced to recall millions of vehicles due to complaints of throttle malfunctions causing unintended accelerations. By then the ETC system had already become one of the most mainstream and simplest modules of a car. Still most theories on the cause of these highly publicised malfunctions considered some software or hardware bug in the ETC to be the source. These theories could neither be proved or disproved easily. After all, major software failures and hardware faults do not always leave physical evidence. This is especially an issue when there are multiple potential causes. Also in the present case, no evidence could be found supporting the theory that the accidents were caused by a software bug. The classical approach of testing the behaviour of a system over a large variety of scenarios to find faults can in general only provide evidence of software bugs, but does not prove absence of them [Dijkstra, 1970, p.7]. To find an answer to the question whether a software fault caused an accident, the exact same operating conditions have to be repeated during a test. Mass produced electronic modules like the ETC should not only operate well under ideal operating conditions, but they should also function safely when sensor readings are disturbed, when communication issues arise, or under minor system changes. This type of resilience is necessary as the multitude of modules are expected to interfere (electromagnetic interference) and since the likelihood of total system failure should be considerably smaller than the degradation of mechanical and electrical parts or the failure of some individual sensors. Generally no finite amount of testing can cover all these required operating conditions while checking for the imposed requirements. To still gain guarantees on the absence of errors, formal verification techniques can be used instead. These methods use mathematical reasonings on the behaviour of the systems to prove the absence of errors. Let us return to the ETC case. In the end, scientists of NASA familiar with formal verification for software systems found no real software or hardware bugs. Still, after a lengthy investigation they uncovered weak spots [Transportation Research Board et al., 2012] in the ETC. Not long after these events, new agreements were made to make the use of formal verification tools during the design phase a standard in the automotive industry.

With the above case, we have highlighted the growing importance of formal methods within a single application domain of control engineering. Similar movements have been observed in other engineering domains. Examples include the railway infrastructure, aerospace engineering, healthcare, etcetera [Kreiker et al., 2011]. Thus more generally, we can state the following.

To guarantee that safety-critical control systems behave under all plausible circumstances it is necessary to have engineering tools that can guarantee the absence of errors by

- avoiding the introduction of design mistakes;
- proving the absence of errors; or
- showing their existence during the design phase.

These should exploit proof techniques that cover all operating conditions and can be algorithmically implemented.

By the time that the need for these tools was raised in (control) engineering, software engineers had already developed them for applications within computer science [Clarke and Wing, 1996]. Their formal methods target the quantification of (software) system properties and the proving of correct functioning. As summarised on the left side of Figure 1.1, formal methods leverage the use of a mathematically based language to specify properties of systems. This mathematical basis, built upon a mathematical sound syntax and semantics, enables these methods to be *mechanised*, or *automated* into a computer-aided set of tools for the analysis, verification, and automated synthesis of the time domain behaviour of the system. These methods can thereby reduce the number of human errors made in the design phase. Whilst this type of verification has been developed for software programs with finite memory (representable by discrete-state models), with the ever more ubiquitous embedding of digital components into physical systems, the focus has been pushed towards richer, cyber-physical, modelling frameworks able to express the combined behaviour of both the digital computer and the surrounding physical system [Clarke and Wing, 1996, Vardi, 2009, Woodcock et al., 2009].

The control of physical systems includes the monitoring, prediction, and actuation of processes evolving over continuous (or more generally uncountable) state spaces with uncertainty in the transitions, model knowledge and state information. Classically, the design of control systems tackles performance objectives either expressed in frequency domain, or in time domain, including reference tracking, disturbance attenuation, stabilisation, etc [Franklin et al., 2001] and assumes only constraints on the system that are easily satisfied. In contrast, formal methods, originally developed for software verification, focus on the functional requirements of the systems; examples include the reaching of targets, or sequential performance of tasks. Employed during the system-development process, they are used to derive guarantees on the requirements imposed on the system. Of interest to us is the extension of this theory originating from computer science to control engineering. Such formal methods for the development and verification of control systems requires the broadening of the available theory to also work for (cyber-)physical systems. As detailed and summarised in Figure 1.1, these are systems whose mathematical description includes uncertainties and whose variables evolve over continuous spaces.

Computer scientists generally consider systems modelled by finite state models

Formal methods

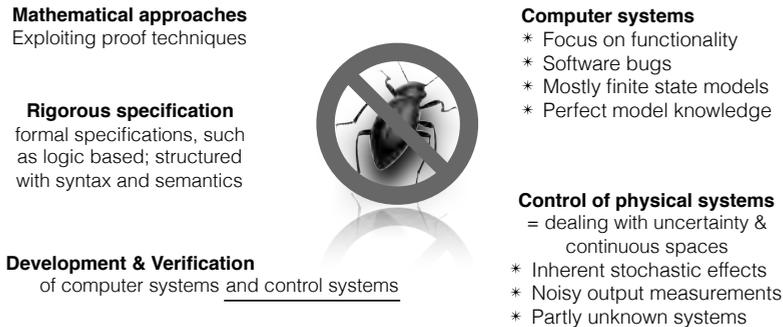


Figure 1.1: Towards the bug-free development of computer systems and control systems, the left side depicts the key aspects on which Formal methods are based. Its development has up to now focused on computer systems. On the right, the implications of this focus on computer systems, are contrasted with the fundamental aspects in the control of physical systems.

for which perfect model knowledge is a given. As mentioned before, research steps to extend this to models common in engineering, which evolve over uncountable state spaces, have already yielded promising results. This includes, but is not limited to results

- in *Robotics*, for the symbolic planning of robot motion in a continuous space [Belta et al., 2007, Lahijanian et al., 2010] and the verification of surgical robots [Kouskoulas et al., 2013];
- in *Biology*, for the analysis of biological systems such as genetic networks [Yordanov et al., 2007, Gössler, 2011]; and
- in the more general set of *Cyber-physical systems* and *Hybrid systems* for specifying complex tasks over the interleaved discrete/continuous behavioural space [Saeedloei and Gupta, 2011].

Still these are only preliminary results to go from methods for computer systems towards (cyber-)physical systems; current methods suffer under the computational load, and have difficulty handling inherent uncertainty or the uncertainty caused by lack of model knowledge.

Challenges

With the increasing amount of information and computational power available for the control of physical systems, future controlled systems are expected to be

more intelligent and autonomous. To counter-act design errors the functioning of these systems should be proven. We identify the several challenges that have to be solved such that formal, proof-based, methods can effectively be used in the design phase of these intelligent and autonomous control systems.

The control of physical systems is modelled most accurately by continuous evolutions in time interleaved with discrete transitions. The discrete operations are defined by the computerised steering or controlling module, similar to any other software program, whereas the continuous evolutions in space and time relate to the behaviour of the physical system. The verification of these cyber-physical systems or hybrid systems is even under perfect model knowledge a daunting task. Moreover, for general hybrid systems this is undecidable ¹ [Kouskoulas et al., 2013]. Currently available results on the verification of hybrid systems are targeting specific subclasses of hybrid models [Puri and Varaiya, 1994, Schupp et al., 2015]. Methods for high order models impose restrictive assumptions on the dynamics of the models such as monotonicity or partial order requirements [Desaraju et al., 2009]. A general framework for the verification of hybrid models that can handle higher order models and different model types in a computationally feasible manner is a challenging topic. This is also currently actively researched within the computer science community [Schupp et al., 2015].

When verifying the behaviour of airplanes in air traffic, external influences such as weather and other airplanes need to be taken into account. Accidents as a consequence of very specific weather conditions are partly unavoidable, hence it is impossible to verify the behaviour uniformly over all weather conditions. Instead weather can be modelled as a stochastic process and it can be verified that the probability of accidents is negligible. Thereby the failure is expressed in relation to how likely the scenario is that caused it. The expression of quantitative and qualitative probabilistic properties via formal languages is well-known, as is their verification and control synthesis over finite state Markov decision processes [Baier and Katoen, 2008]. However the extension to continuous, or even hybrid stochastic systems hinges on approximate methods that do not scale to the complexity of realistic systems or should be build on model reductions or abstractions for which only in some cases error bounds are available.

The behaviour of systems composed of physical components is generally only known up to a certain point. Within control engineering, a controller for a partly unknown plant or system is generally developed as follows. First, while applying input signals to the physical system, its behaviour over time is measured. The data is collected and then used to estimate a model of the system [Ljung, 1999, Peterka, 1981]. Then a controller is designed with respect to a given performance criterion. The cost effective gathering of data and the certification of controllers with respect to performance criteria classically present in systems and control is a mature research topic [Gevers, 2005]. Due to the presence of uncertainty, the use and interpretation of data obtained from system measurements is fundamental for the control of physical systems. Still, there is little known on how to combine

¹If it is known to be impossible to construct a single algorithm answering a decision problem (e.g. the verification of a hybrid system with respect to a property) with the correct yes-or-no answer, then that decision problem is undecidable.

formal methods to prove functional requirements in an automatised way together with uncertain system descriptions and measurement data. For systems evolving over finite state spaces, this has already triggered research on data-based formal methods. This includes data-driven modelling with as goal verification [Chen and Nielsen, 2012] and learning-based controller synthesis [Fu and Topcu, 2014]. For systems evolving (stochastically) over continuous spaces, the formal verification and controller synthesis based on data or based on models build from data lacks coverage in research.

1.2 Research questions

Of the open challenges, we perceive the handling of stochastic effects, uncertainty and the use of data as key. We pick these issues, as they cause a lack of knowledge on the behaviour of the system, which, when left unresolved, can render practical application of formal methods to (simple) physical systems infeasible. As elucidated in the following chapter, we consider a system to be a physical system, and a perfect *model* to be the exact mathematical description of the associated dynamical behaviour. An input signal represents the actions of the environment on the system. Similarly, the variables with which the system influences its environment are output signals. Measuring these signals gives additional information on the system and allows us to attenuate uncertainty.

On these premises we raise the question: How should we synthesise control modules and formally prove properties defined over physical systems, when our knowledge is limited by uncertainty?

As we cannot tackle the question in its full generality yet, we first categorise the types of uncertainty in physical systems. Then we pose sub-questions relating to combinations of the types of uncertainty. We differentiate three types of uncertainty in physical systems:

1. Inherent stochastic effects over uncountable state spaces.
Consider a power network for which we would like to quantify the likelihood of blackouts and to synthesise strategies to minimise this. This network is a stochastic system for which at every time instant the operator can change variables of this process, which are referred to as control variables or inputs. Probabilistic transitions of this system model the unpredictable changes of power demands and lead to probability distributions describing the next state based on the current state and control input. Although the probability of blackouts cannot be fully removed, by designing a control system that chooses the most ideal inputs, the probability of blackouts can be minimised.
2. Partly unknown systems.
The knowledge of the behaviour of the system is often limited or uncertain, making it impossible to analyse its dynamics by means of a "true" model as postulated in formal methods. Instead, in a Bayesian framework, a-priori

available knowledge allows to construct a model class indexed with a parameterisation quantifying the uncertainty by means of a prior probability distribution over the parameters.

3. Noisy output measurements.

Often the state of a system cannot be monitored directly. Instead only output variables or signals can be measured. Measurement data of these variables obtained from sensors is often inexact. Referred to as noisy output measurements, we consider the case that the accuracy of such measurements can be quantified via probabilistic distributions, which define coloured or white noise disturbances on the system outputs.

For cases containing at least one of the above types of uncertainty, that are perceived as simple within the systems and control domain, we question how to either automatically synthesise correct controllers or how to verify the correct functioning based on formal methods. Leaving the problem of partly unknown dynamics to the side, we first pose the research questions related to controller synthesis. We consider in detail the case where full state information is not available and measurements are noisy. Secondly, we come back to the problem of partly unknown systems and investigate formal and data-driven methods for verification.

1.2.a Subproblems on control synthesis

For physical systems, best represented by discrete time stochastic processes evolving over general uncountable state spaces, we investigate the design of certifiable control strategies with respect to pre-specified requirements, such as probabilistic invariance, reachability, or more complex probabilistic properties. These processes can be modelled as Markov decision processes and their verification with respect to probabilistic properties can often not be attained analytically. Instead for the verification and synthesis of control strategies one can rely on reductions of the models to (finite) state space abstractions for which controller synthesis can be constructed algorithmically with guarantees on its properties [Esmail Zadeh Soudjani and Abate, 2013]. Leveraging model reductions or abstractions of this type that enable the quantitative refinement of control synthesised on the abstract model to the (original) concrete model, one can still obtain controllers that are correct-by-construction. For this it is key to provide formal guarantees on the approximation or abstraction step. Normally the abstract and concrete models are related with (approximate) similarity relations, but the available relations quantify deviations in probability only and require a kind of lumping. As such they do not scale well and are less well applicable for abstractions obtained based on classical systems and control methods. As these are often based on model order reductions that give output trajectories closely related to the original model, norm-based relations make more sense. We search for a new probabilistic similarity relation, with relaxed requirements on the allowed set of relations; which allows for deviations both in probability and output. A detailed formulation of this problem and the results can be found in *Chapter 3* entitled *Approximate similarity and controller refinement for general Markov decision processes*.

Current state-of-the-art correct-by-design controllers are designed for full-state measurable systems. For Linear Time Invariant systems (LTI) disturbed by Gaussian noise both on the state evolutions and on the output measurements, we question whether it is possible to apply state-based correct-by-design controllers by employing a state-estimator, or equivalently an observer. For this we require that the extended correct-by-design output-based controller has quantified bounds on accuracy loss. In *Chapter 4* entitled *Output-based controller synthesis: a correct-by-design approach for Gaussian LTI systems* we consider this problem in detail.

The computation and algorithmic implementation of these observer-based extensions of the state-based correct-by-design controller synthesis to the output-based case, can be shown to be related to the H_2 -optimal control problem. In the latter problem the goal is to design an output-based controller with respect to a H_2 -norm. The quantification of this H_2 -norm contains a feasibility check over matrix inequalities that are equal to those defined for the quantification of the accuracy loss. We know that the original-optimal control problem can be separated into two separate control problems, one relating to the design of a state-estimator and one relating to the design of a state-based controller. We now investigate whether we can also impose this type of separation on the set matrix inequalities. More precisely to increase the computational efficiency, we need results for these inequalities that allow the separation of estimation and control via feasible linear matrix inequalities. We take on this problem in *Chapter 5* entitled *A separation theorem for guaranteed performance through matrix inequalities*.

1.2.b Subproblems on verification of partially unknown systems

We consider partly unknown systems for which we can design experiments and obtain noisy output measurements, and we question how to verify formal properties over the system. We specifically analyse the case where the lack of knowledge on the system dynamics makes answering the verification question with a yes-or-no answer impossible.

In *Chapter 6* entitled *Data-driven and model-based verification*, we consider the case that the system knowledge can be used to construct a parameterised model class, over whose parameter space the remaining uncertainty can be quantified by means of a prior probability distribution.

Starting from a given set of system measurements, we seek a quantification conditioned on this data of the confidence or believe in the satisfaction of a property over the unknown system. Thus we pose the verification question within a Bayesian framework. When applied to nonlinearly parameterised, linear time invariant (LTI) models such a Bayesian reasoning generally introduces heavy computational issues, which can only be mitigated via statistical methods [Gyori et al., 2014]. Instead, in order to obtain reliable and numerical solutions, we investigate the use of linearly parameterised model sets defined through orthonormal basis functions.

After deriving a data-driven and model-based approach in *Chapter 6*, we subsequently question the ‘optimal’ design of experiments to be used for verification in *Chapter 7*. We first assess the difference between the use of experiments for

verification and their use for estimation. In the latter case, the accuracy of the estimated model is key. Instead, in the verification setting the overall objective is to verify or falsify a property of the underlying system of interest. As such, the end result of an experiment is the confidence in either accepting or rejecting a property. We first show that objective of the experiment (either quantifying the confidence in a property or deciding on an hypothesis) can be best tackled in a Bayesian framework. This is detailed in *Chapter 7* entitled *Bayesian experiment design for formal verification* where we further research the existence and the properties of computational solutions to the experiment design problem when the optimality criteria is either

- the expected confidence in the property of interest, or
- the probability that a given threshold confidence is reached.

We ask these questions related to the design of experiments on a very simple setting to build insight into the properties of the optimal experiment inputs. The development of methods that can solve the experiment design under realistic state dimensions and assumptions on the system and measurement set-up is perceived beyond the scope of this thesis.

1.3 Thesis outline and publications

The division in research questions, that is synthesis versus verification, is also reflected in the structure of the thesis. A collection of technical preliminaries is given in the next chapter. This is followed by the first part of the thesis, which addresses formal controller synthesis for systems affected by inherent stochastic effects and evolving over uncountable state spaces. Consecutively we consider the case that full state information is not available and output measurements are noisy. The second part of the thesis deals with the verification of properties when exact model knowledge is lacking and measurements are noisy.

The thesis has been structured in chapters, each of them targeting a specific sub-problem within either the control synthesis or verification question. Up-to-now, we have given a birds-eye perspective on the field of verification and control without giving a full embedding in literature. This will be rectified in the subsequent chapters, where together with the detailed posing of the research questions, also a relevant positioning in literature is given.

The following publications have been written related to the research in this doctoral thesis.

- Chapter 3
S. Haesaert, S. Esmaeil Zadeh Soudjani, A. Abate. *Verification of general Markov decision processes by approximate similarity relations and policy refinement*. (arXiv preprint arXiv:1605.09557), under review, 2016.
S. Haesaert, A. Abate, P.M.J Van den Hof, *Verification of general Markov decision processes by approximate similarity relations and policy refinement*. QUANTITATIVE EVALUATION OF SYSTEMS, pages 227–243. Springer, 2016.

- Chapter 4 and 5
 - S. Haesaert, P.M.J. Van den Hof, A. Abate. *Observer-based correct-by-design controller synthesis*. (arXiv preprint arXiv:1509.03427), 2016
 - S. Haesaert, A. Abate, P.M.J. Van den Hof, *Correct-by-design output feedback of LTI systems*.
IEEE CONFERENCE ON DECISION AND CONTROL, Osaka, Japan, pp.6159-6164, 2015.
 - S. Haesaert, S. Weiland, C. Scherer. *A separation theorem for guaranteed H_2 performance through matrix inequalities*. (To be submitted to a journal)
- Chapter 6 and 7
 - S. Haesaert, P.M.J. Van den Hof, A. Abate. *Data-driven and Model-based Verification: A Bayesian identification approach*. (arXiv preprint arXiv:1509.03347.)
Accepted as a Full Paper in AUTOMATICA
 - S. Haesaert, P.M.J. Van den Hof, A. Abate, *Experiment Design for Formal Verification via Stochastic Optimal Control*.
IEEE EUROPEAN CONTROL CONFERENCE, Aalborg, Denmark, 2016.
 - S. Haesaert, P.M.J. Van den Hof, A. Abate, *Data-driven and model-based verification: a Bayesian identification approach*.
IEEE CONFERENCE ON DECISION AND CONTROL, Osaka, Japan, pp. 6830-6835, 2015.
 - S. Haesaert, P.M.J. Van den Hof, A. Abate, *Data-driven property verification of grey-box systems by Bayesian experiment design*.
IEEE AMERICAN CONTROL CONFERENCE, Chicago, USA, pp. 1800-1805, 2015.

The following publication is related to the material in this thesis, but is not included in the thesis

- E. Polgreen, V.B. Wijesuriya, S. Haesaert and A. Abate. *Data-Efficient Bayesian Verification of Parametric Markov Chains*.
QUANTITATIVE EVALUATION OF SYSTEMS, pages 35-51. Springer, 2016.

Language is by its very nature a communal thing; that is, it expresses never the exact thing but a compromise – that which is common to you, me, and everybody.

Thomas Earnest Hulme, *Speculations*, 1923

2

Verification and control of physical systems: a prelude

2.1 Physical systems

All of space and its contents over all of time is referred to as the universe. The laws followed by matter and energy, subject of investigation of physics, govern its change over time. A physical system is in its origin a portion of the universe. It contains matter and energy governed by the laws of nature. We use it to isolate a unified whole of interrelated interacting matter and energy over time. The remaining part of the universe, called the environment, is seen as an entity interacting over the systems boundaries.

The division into physical system and environment by means of the system's boundary is done by us, the observer. Within the context of this thesis, the physical system represents a part of the universe that we would like to verify or control. For this we build a mathematical model of the behaviour of the physical system and its interaction over the system boundary. The model presents the dominant phenomena in the system relevant to the posed verification or control problem. We consider physical systems that are dynamically evolving, that is, systems of which the quantitative variables of the relevant phenomena change dynamically over time. Inasmuch this can be said about any subset of the universe, we particularly focus on systems and their verification and control problems, which can be best tackled with dynamical models.

Exemplar for such a physical system is an office building and more specifically its room temperatures and thermal phenomena. Leveraging thermodynamics we know that over time the temperature fluctuations depend on the heat capacity of

the office building, the heat flux of the radiators placed in the rooms and the heat generated by people working in the office together with heat transferred between the inside and the outside. The isolation of phenomena and variables of interest to us into a physical system is a choice that, together with the choice of dominant phenomena to be modelled, depends on the analysis or control questions asked.

Consider the design of a thermostat that controls the temperature in the offices; based on temperature readings it turns the radiators on or off. To design the software of the thermostat a control engineer selects the physical system such that it encapsulates the phenomena governing the heat fluctuations. More precisely, the system is defined as the content and the office building itself up to where it touches the ambient air. The changes of variables over time are modelled based on the laws of thermodynamics and depend on variables defined across the system boundary. To model this interaction of the physical system over the system boundary, we classically require that we can differentiate these variables into inputs, outputs and disturbances. In this case the inputs are the on/off signals sent to the radiators; the outputs are the temperature readings of the sensors placed in the rooms; and the disturbances are the (unpredictable) temperature fluctuations of the ambient, the heat fluctuations caused by people, and the changes in solar radiation.

Let us be more precise. With *output signals*, we indicate those variables indexed over time with which the system interacts with the environment. Furthermore, we say that an *input variable*, or the signal it defines over time, captures how the environment acts on the system. Control inputs, also referred to as actions are those variables that we can control or choose freely. This is in contrast to disturbance inputs, which contrary to control inputs can neither be manipulated nor predicted fully. As such disturbance inputs represent a part of the environment of influence to the modelled phenomena, whose dynamics are deliberately not included in the physical system and hence not modelled.

In this work, we define a physical system as a part of the universe evolving in time over a continuous space, whose future behaviour depends on its causal interaction over its input and output signals. We presume that we can model the behaviour of the system, either based on our knowledge of the governing laws of nature or by observing it, but that we cannot physically change the composition of the physical system.

The definition of physical systems differs from domain to domain. Similarly, widely used words such as models, actions and policies have a different interpretation in computer science, probability theory, and control engineering each of them perceived as fundamental and standard. Therefore, analogous to how we introduced the notion of a physical system, we use this chapter to introduce the fundamental concepts on which this thesis is built. More specifically we introduce the problems we will tackle, by first introducing the models of physical systems and the properties that we will be looking at.

We will limit the scope of this thesis to models of systems controlled and measured over discrete-time. That is, systems whose behaviour is monitored and whose performance is quantified with respect to samples taken with regular time intervals.

For this we will tackle two types of systems, classified on how they can be modelled. Firstly we will look at *deterministic systems*. For these systems we model the time evolutions deterministically; given an input signal and the initial state of the system, the model of the system allows us to exactly predict the behaviour of the system. In other words, we classify physical systems and their mathematical models for which appropriate knowledge on the governing laws would allow us to predict the future evolutions of the system exactly within this *deterministic framework*, cf. Section 2.3.

Alternatively in Section 2.4, we introduce a *stochastic framework* for those systems whose evolutions are inherently uncertain and can be quantified best by probability measures. In this case some uncertainty about the future behaviour is retained. This uncertainty is quantified with a probabilistic distribution over the future paths of the system. The latter case represents the uncertainty caused by behaviours that cannot be modelled or predicted in a deterministic way. For both modelling frameworks, stochastic and deterministic, it is not assumed that all information and quantities necessary to model the system are available.

As common in engineering, we draw the system boundary such that it encapsulates well-defined physical quantities of which many can be measured by some appropriate equipment. In this thesis, we will use the measuring of data to tackle the remaining uncertainty, caused by stochastic transitions or unknown dynamics.

The Sections 2.3 and 2.4 on the deterministic and stochastic frameworks are preceded with some standard technical preliminaries.

2.2 Technical preliminaries

Sets, relations, and orders Given two sets A and B , the Cartesian product of A and B is given as $A \times B = \{(a, b) : a \in A \text{ and } b \in B\}$. As standard, we denote the union of two sets with \cup , i.e., $A \cup B := \{x : x \in A \text{ or } x \in B\}$. We also use the disjoint union of A and B , this is denoted as $A \sqcup B$ and consists of the combination of the members of A and B , where the original set membership is the distinguishing characteristic that forces the union to be disjoint, i.e., $A \sqcup B = (A \times \{0\}) \cup (B \times \{1\})$. As usual for $C \subset A \sqcup B$, we denote $C \cap A = \{a \in A : (a, 0) \in C\}$. For the sets A and B a relation $\mathcal{R} \subset A \times B$ is a subset of their Cartesian product that relates elements $x \in A$ with elements $y \in B$, denoted as $x\mathcal{R}y$. More precisely, we denote $(x, y) \in \mathcal{R}$ with $x \in A$ and $y \in B$ as $x\mathcal{R}y$. Further, the relation \mathcal{R} defines set-valued mappings $\mathcal{R}(x) := \{y \mid x\mathcal{R}y\}$ and $\mathcal{R}^{-1}(y) := \{x \mid x\mathcal{R}y\}$. We trivially extend this notation to the mappings $\mathcal{R}(\tilde{A}) := \{y : x\mathcal{R}y, x \in \tilde{A}\}$ and $\mathcal{R}^{-1}(\tilde{B}) := \{x : x\mathcal{R}y, y \in \tilde{B}\}$ for $\tilde{A} \subseteq A$ and $\tilde{B} \subseteq B$. A relation over a set defines a preorder if it is reflexive, $\forall x \in A : x\mathcal{R}x$; and transitive, $\forall x, y, z \in A : \text{if } x\mathcal{R}y \text{ and } y\mathcal{R}z \text{ then } x\mathcal{R}z$. We say that a relation $\mathcal{R} \subseteq A \times A$ is an equivalence relation if it is reflexive, transitive and symmetric, that is $\forall x, y \in A : \text{if } x\mathcal{R}y \text{ then } y\mathcal{R}x$.

An *equivalence class* is given over the set A as $[x]_{\mathcal{R}} = \{x' \in A : (x, x') \in \mathcal{R}\}$ and induces a partition over the set A also known as a quotient space $A/\mathcal{R} = \{[x]_{\mathcal{R}} : x \in A\}$.

For a given set \mathbb{X} a metric or distance function $\mathbf{d}_{\mathbb{X}}$ is a function $\mathbf{d}_{\mathbb{X}} : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}_0^+$. For Euclidean spaces the Euclidean norm, denoted $\|\cdot\|_2$, induces a metric as $\mathbf{d}(x_1, x_2) = \|x_1 - x_2\|_2$.

Vectors, matrices and matrix inequalities The notion of vectors $x \in \mathbb{R}^n$ is used in the remainder to define n -dimensional column vectors, specifically $x \in \mathbb{R}^{n \times 1}$. The trace of a matrix $A \in \mathbb{R}^{n \times n}$ is denoted as $\text{tr}(A)$. Remember that a matrix A is symmetric if $A = A^T$, where $(\cdot)^T$ denotes the matrix transpose. For a scalar $c = a + bi$ with $a, b \in \mathbb{R}$ its complex conjugate is given as $\bar{c} = a - bi$. The conjugate transpose (\dagger) of a matrix is defined as $A^\dagger = (\bar{A})^T$, where \bar{A} denotes the element-wise extension of the scalar conjugate to the matrix. We say that a matrix A is Hermitian if it is square and $A = A^\dagger$. Denote the set of $n \times n$ Hermitian matrices as \mathbb{H}^n and the set of symmetric matrices as \mathbb{S}^n . Positive or negative definiteness of matrices in \mathbb{S}^n and \mathbb{H}^n is defined for a matrix $A \in \mathbb{S}^n$, respectively $A \in \mathbb{H}^n$, as follows

A is positive definite if $x^\dagger Ax > 0$ for all $x \in \mathbb{C}^n$ with $x \neq 0$;

A is positive semi-definite if $x^\dagger Ax \geq 0$ for all $x \in \mathbb{C}^n$;

similarly, A is negative definite if $x^\dagger Ax < 0$ for all $x \in \mathbb{C}^n$ with $x \neq 0$; A is negative semi-definite if $x^\dagger Ax \leq 0$ for all $x \in \mathbb{C}^n$. For $A, B \in \mathbb{H}^n$ or \mathbb{S}^n , we denote $A \prec B$ and $A \preceq B$ if $A - B$ is negative definite or negative semi-definite, respectively. Similarly, the corresponding inequalities $A \succ B$ and $A \succeq B$ are true if $A - B$ is positive definite, respectively positive semi-definite. Note that $A \prec 0$ ($A \succ 0$) if and only if all eigenvalues of A are strictly negative (positive).

As is standard, c.f. [Scherer and Weiland, 2000], we define a linear matrix inequality (LMI) as an inequality $F(x) \prec 0$ where F is an affine function¹ mapping of a finite dimensional vector space \mathbb{R}^n with elements $x = [x_1 \ x_2 \ \dots \ x_n]^T$ to either the set \mathbb{H} or the set \mathbb{S} of symmetric matrices, i.e.,

$$F(x) := F_0 + F_1 x_1 + \dots + F_n x_n$$

with $F_i \in \mathbb{H}$ or $F_i \in \mathbb{S}$. We will work with linear matrix inequalities described by functions of matrix variables $X \in \mathbb{R}^{n_1 \times n_2}$ rather than scalar design variables x_i . Of course these matrix functions $F(X)$ can again be rewritten as functions of scalar design variables.

2.3 Deterministic framework

Consider a class of finite-dimensional dynamical models that evolve in discrete-time and that are linear time-invariant (LTI). These models depend on input and output signals ranging over \mathbb{R}^m and \mathbb{R}^p , respectively, and on variables $x(t)$ taking values in an Euclidean space, $x(t) \in \mathbb{X} \subseteq \mathbb{R}^n$, where n , the state dimension, is the

¹An affine function in x is a linear function in x plus a constant translation.

model order. The behaviour of such a system can be described by a state-space model, denoted (A, B, C, D) and given as

$$\begin{aligned}x(t+1) &= Ax(t) + Bu(t), \\y(t) &= Cx(t) + Du(t),\end{aligned}\tag{2.1}$$

where matrices A, B, C, D are of appropriate dimensions. At $t = 0$ the state $x(0)$ of the model is initialised from \mathbb{X}_0 , that is $x(0) \in \mathbb{X}_0$, where $\mathbb{X}_0 \subset \mathbb{X}$ represents the set of possible initial states. A signal, that is, a time-indexed sequence of variables, given as $\{x(t)\}_{t \geq 0}$ is a possible trajectory of the LTI model (2.1) if $x(0) \in \mathbb{X}_0$ and if for every $t \geq 0$ there exists a $u(t)$ such that $x(t+1) = Ax(t) + Bu(t)$. Similarly, a signal is an output trace $\{y(t)\}_{t \geq 0}$ if there exists a trajectory $\{x(t)\}_{t \geq 0}$ of the LTI model (2.1) such that $y(t) = Cx(t) + Du(t)$ for all $t \geq 0$. When referring to state evolutions, we talk about the transition of $x(t)$ to $x(t+1)$ governed by the equation $x(t+1) = Ax(t) + Bu(t)$.

Relations between input and output signals can also be used to define discrete-time systems. For this consider the forward shift operator $qu(t) = u(t+1)$ and the backward shift operator q^{-1} and $q^{-1}u(t) = u(t-1)$. By extension, we also allow for polynomial functions of q that induce signal transformations and compositions as $\sum_{i=1}^n d_i q^i y(t) = \sum_{i=1}^n d_i y(t+i)$. We can now use such polynomials to verify whether two signals $u(t)$ and $y(t)$ belong to the behaviour of a linear discrete-time system. Let $D(q)$ and $N(q)$ be polynomial functions in the q operator. Suppose that the behaviour of the discrete-time systems contains only those input and output signals which satisfy

$$D(q)y(t) = N(q)u(t).$$

We can now define a relation over the space of input and output sequences with elements $(\{u(t)\}_{t \geq 0}, \{y(t)\}_{t \geq 0}) \in \mathcal{R}_{(D \setminus N)}$:

$$\mathcal{R}_{(D \setminus N)} := \{(\{u(t)\}_{t \geq 0}, \{y(t)\}_{t \geq 0}) \text{ s.t. } D(q)y(t) = N(q)u(t) \text{ for all } t \geq 0\}.$$

As is common in control engineering, we can now write the mapping from the input signals to the output signals based on this relation. This mapping, which one could formally denote with $\mathcal{R}_{(D \setminus N)}(\{u(t)\}_{t \geq 0})$ is commonly denoted with the rational expression $G(q) := D(q)^{-1}N(q)$, that is

$$y(t) = G(q)u(t)$$

and, when applicable, referred to as the transfer operator of a linear time invariant model.

A more general modelling framework reaching beyond state evolutions over Euclidean spaces given in (2.1) is offered by the notion of transitions systems introduced next. Introduced first by Keller in [Keller, 1976], they are now a standard class of models for representing hardware and software systems. They have been introduced in the control systems community by the work of inter alia [Tabuada, 2009].

Definition 2.1 (Transition system) The tuple $\text{TS} = (\mathcal{X}, \mathcal{X}_0, \mathcal{A}, \rightarrow, \mathcal{Z}, \mathcal{H})$ defines a transition system for which

- \mathcal{X} is a (possibly infinite) set of states;
- \mathcal{X}_0 is a (possibly infinite) set of initial states;
- \mathcal{A} is a (possibly infinite) set of actions, with elements $u \in \mathcal{A}$;
- $\rightarrow \subseteq \mathcal{X} \times \mathcal{A} \times \mathcal{X}$ is a transition relation;
- \mathcal{Z} is a (possibly infinite) set of observations;
- $\mathcal{H} : \mathcal{X} \rightarrow \mathcal{Z}$ is a map assigning to each $x \in \mathcal{X}$ an observation $\mathcal{H}(x) \in \mathcal{Z}$.

A metric transition system is a transition system endowed with a metric over the observation space \mathcal{Z} .

As before, we say that $\{x(t)\}_{t \geq 0} \in \mathcal{X}^{\mathbb{N}}$ is a trajectory of the transition system if for all $t \geq 0$ there exists an action or input $u(t) \in \mathcal{A}$ such that $(x(t), u(t), x(t+1)) \in \rightarrow$ and if it is initialised, that is, if $x(0) \in \mathcal{X}_0$. Similarly $\{y(t)\}_{t \geq 0}$ defines a trace of the transition system, if there exists a trajectory $\{x(t)\}_{t \geq 0}$ and if $\{y(t)\}_{t \geq 0} = \{\mathcal{H}(x(t))\}_{t \geq 0}$.

Models in the format of (2.1) can also be written as a transition system if they are strictly proper, that is, if $D = 0$. The transition relation is then defined as $\rightarrow := \{(x, u, x^+) \in \mathbb{X} \times \mathbb{R}^m \times \mathbb{X} : x^+ = Ax + Bu\}$. Similarly, the output map is defined as $\mathcal{H}(x) = Cx$.

Considering the definition of transition systems, note that for every state-action pair, there can either be no possible next state, a unique next state, or a set of possible next states. The last type of transitions are referred to as being non-deterministic within computer science. A system is blocking if there are state-action pairs with no successors (that is next states). Within this thesis we say that a transition system is deterministic if every state-action pair yields a unique transition.

We will use this modelling framework, both as a way to represent the dynamics of physical systems, and to represent the controllers. As such, we will introduce the required notions of composition and control next. But first let us consider the properties of interest for control. We say that the transition system has a behaviour associated to it. More precisely, the behaviour *generated* by the system is the set of all *trajectories*. Whereas the output behaviour is the set of all *traces* of the system. It is over the output behaviour, denoted with \mathcal{B} , of the transition system that we define the properties of interest.

2.3.a Linear-time properties

Linear-time properties are a subset of formal properties and are based on a linear ordering of events in time. Linear time properties have the benefit that they can be

described through operations on *languages*. We will briefly introduce them in this subsection, an extensive introduction can be found in [Baier and Katoen, 2008]. Further, the history of temporal logic properties is recapped in [Vardi, 2009].

Definition 2.2 (Languages) Σ , the alphabet, is a finite set with elements $\sigma \in \Sigma$, the letters of the alphabet. A word is an infinite sequence of letters described by a mapping

$$\pi : \mathbb{N} \rightarrow \Sigma.$$

Any set of words is a language \mathcal{L} and is a subset of the set of all words $\Sigma^{\mathbb{N}}$.

For a given set of signals taking values in some abstract signal space \mathbb{W} , we can now define the concept of language as follows.

Definition 2.3 The labelling map Lab labels the elements in the signal space \mathbb{W} with letters of the alphabet Σ :

$$\text{Lab} : \mathbb{W} \rightarrow \Sigma.$$

Given a signal $\omega \in \mathbb{W}^{\mathbb{N}}$, the composition of the labelling map with the signal $\pi = \text{Lab} \circ \omega$ is a word. This is also denoted as $\text{Lab}(\omega)$. The set of words $\text{Lab}(\Omega)$ generated by a set of trajectories $\Omega \subset \mathbb{W}^{\mathbb{N}}$ is defined as

$$\text{Lab}(\Omega) := \{\text{Lab}(\omega) : \omega \in \Omega\} \subseteq \Sigma^{\mathbb{N}}.$$

We say that every language $\mathcal{L}_\psi \subset \Sigma^{\mathbb{N}}$ defines an Linear-time (LT) specification ψ . For a single trajectory $\omega \in \mathbb{W}^{\mathbb{N}}$ and a linear-time specification $\psi \in \text{LT}$, with LT the abstract set of Linear-time specifications, the satisfaction relation is defined over the set $\models \subseteq (\mathbb{W}^{\mathbb{N}} \times \text{LT})$ as

$$\omega \models \psi \Leftrightarrow \text{Lab}(\omega) \in \mathcal{L}_\psi.$$

We recall that \mathcal{B} is the collection of all traces generated by a transitions system TS. Then the verification of a specification over a transition system TS, composed with the labelling map Lab , can be expressed with the satisfaction relation based on the *language generated* by the system, i.e., $\text{Lab}(\mathcal{B})$.

Definition 2.4 (TS $\models \psi$) Let TS be a given transition system and let Lab be its labelling map. We say that the satisfaction relation admits the specification ψ for TS, denoted $\text{TS} \models \psi$, if

$$\text{Lab}(\mathcal{B}) \subset \mathcal{L}_\psi.$$

thus if the language generated by TS is a subset of the accepting language \mathcal{L}_ψ of the LT property ψ .

We now consider the use of logical propositions to define the language based properties that are of interest to us. For this we introduce *Linear-time Temporal Logic* (LTL) which has been introduced in the present form in computer science by [Pnueli, 1977]. LTL is a common specification language used in inter alia [Tabuada and Pappas, 2003, Baier and Katoen, 2008] and it contains temporal logic formulas, which are built up recursively as

- true and $p \in \text{AP}$ are LTL formulas, with p elements of the finite set of Atomic Propositions (AP);
- if ψ is an LTL formula, then $\neg\psi$ is an LTL formula;
- if ψ_1 and ψ_2 are LTL formulas, then $\psi_1 \wedge \psi_2$ and $\psi_1 \vee \psi_2$ are LTL formulas;
- if ψ_1 and ψ_2 are LTL formulas, then $\bigcirc\psi_1$ and $\psi_1 \text{ U } \psi_2$ are LTL formulas.

This recursive composition defines all the possible LTL formulas and is called the syntax, with short notation

$$\psi ::= \text{true} \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \bigcirc\psi \mid \psi \text{ U } \psi.$$

The formula $\psi_1 \vee \psi_2$ is equal to $\neg(\neg\psi_1 \wedge \neg\psi_2)$. Thus the smaller syntax fragment

$$\psi ::= \text{true} \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc\psi \mid \psi \text{ U } \psi$$

is sufficient to describe all LTL formulas.

Let $\pi = \pi(0), \pi(1), \pi(2), \dots \in \Sigma^{\mathbb{N}}$ be a infinite word and let $\pi_t = \pi(t), \pi(t+1), \pi(t+2), \dots$ as a subsequence (postfix) of π , then the satisfaction relation between π and ψ ; $\pi \models \psi$ or equivalently $\pi_0 \models \psi$ is defined recursively over ψ as given in the following definition.

Definition 2.5 Let $\pi_t = \pi(t), \pi(t+1), \pi(t+2), \dots$ be an infinite word starting at $t \in \mathbb{N}$ with letters $\pi(\cdot) \in \Sigma$. Let the alphabet $\Sigma = 2^{\text{AP}}$ be the power set of the set of atomic propositions $p \in \text{AP}$. Then the satisfaction relation \models is defined recursively over π_t as

(true)	$\pi_t \models \text{true}$	$\Leftrightarrow \text{true}$
(atomic proposition)	$\pi_t \models p$	$\Leftrightarrow p \in \pi(t)$
(negation)	$\pi_t \models \neg\psi$	$\Leftrightarrow \pi_t \not\models \psi$
(conjunction)	$\pi_t \models \psi_1 \wedge \psi_2$	$\Leftrightarrow \pi_t \models \psi_1 \text{ and } \pi_t \models \psi_2$
(disjunction)	$\pi_t \models \psi_1 \vee \psi_2$	$\Leftrightarrow \pi_t \models \psi_1 \text{ or } \pi_t \models \psi_2$
(next)	$\pi_t \models \bigcirc\psi$	$\Leftrightarrow \pi_{t+1} \models \psi$
(until)	$\pi_t \models \psi_1 \text{ U } \psi_2$	$\Leftrightarrow \exists i \in \mathbb{N} : \pi_{t+i} \models \psi_2,$ $\text{and } \forall j \in \mathbb{N} : 0 \leq j < i, \pi_{t+j} \models \psi_1.$

Of special interest to us are specifications defined over a finite time horizon N . These are specifications that can be represented by a collection of finite words. A finite word π' over is defined by a finite sequence of letters, i.e., $\pi' = \pi'(0), \pi'(1), \dots, \pi'(N)$, and defines a *basic language*, this is a language (cf. Definition 2.2) defined as $\{\pi'\} \times \prod_{i=n+1}^{\infty} \Sigma$. Thus any specification over bounded time is a collection of basic languages. Bounded Linear-time Temporal Logic (BLTL) is a subset of LTL with syntax fragment

$$\psi ::= \text{true} \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc\psi.$$

2.4 Probabilistic framework

Whatever way uncertainty is approached, probability is the only sound way to think about it.

Dennis Lindley

2.4.a Measurable space, probability, and random variables

A measurable space is a pair $(\mathbb{X}, \mathcal{F})$ with sample space \mathbb{X} and σ -algebra \mathcal{F} defined over \mathbb{X} , which is equipped with a topology. As a specific instance of \mathcal{F} consider the Borel measurable space $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. In this work, we restrict our attention to Polish spaces and generally consider the Borel σ -field [Bogachev, 2007]. Recall that a Polish space is a separable completely metrisable topological space. In other words, the space admits a topological isomorphism to a complete metric space which is dense with respect to a countable subset. A simple example of such a space is the real line \mathbb{R} .

A probability measure \mathbb{P} for $(\mathbb{X}, \mathcal{F})$ is a non-negative map, $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ such that $\mathbb{P}(\mathbb{X}) = 1$ and such that for all countable collections $\{A_i\}_{i=1}^{\infty}$ of pairwise disjoint sets in \mathcal{F} , it holds that $\mathbb{P}(\bigcup_i A_i) = \sum_i \mathbb{P}(A_i)$. Together with the measurable space, such a probability measure \mathbb{P} defines the probability space, which is denoted as $(\mathbb{X}, \mathcal{F}, \mathbb{P})$ and has realisations $x \sim \mathbb{P}$. Let us further denote the set of all probability measures for a given measurable pair $(\mathbb{X}, \mathcal{F})$ as $\mathcal{P}(\mathbb{X}, \mathcal{F})$. For a probability space¹ $(\mathbb{X}, \mathcal{F}_{\mathbb{X}}, \mathbb{P})$ and a measurable space $(\mathbb{Y}, \mathcal{F}_{\mathbb{Y}})$, a $(\mathbb{Y}, \mathcal{F}_{\mathbb{Y}})$ -valued *random variable* is a function $\mathbf{y} : \mathbb{X} \rightarrow \mathbb{Y}$ that is $(\mathcal{F}_{\mathbb{X}}, \mathcal{F}_{\mathbb{Y}})$ -measurable, and which induces the probability measure $\mathbf{y}_* \mathbb{P}$ in $\mathcal{P}(\mathbb{Y}, \mathcal{F}_{\mathbb{Y}})$.

Consider the measurable space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and let \mathbf{z} be a $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ -valued random variable with induced distribution $\mathbb{P}_{\mathbf{z}}$ defined on a probability space $(\mathbb{X}, \mathcal{F}, \mathbb{P})$. For brevity, we will refer to \mathbf{z} as a random variable with distribution $\mathbb{P}_{\mathbf{z}}$ and realisations $z \sim \mathbb{P}_{\mathbf{z}}$. A random variable has an expected value defined as

$$\mathbb{E}[\mathbf{z}] = \int_{\mathbb{R}} z \mathbb{P}_{\mathbf{z}}(dz)$$

and a second order moment defined as $\mathbb{E}[\mathbf{z}^2]$. Similarly the variance of \mathbf{z} , or equivalently the second central moment is defined as $\mathbb{E}[(\mathbf{z} - \mathbb{E}[\mathbf{z}])^2]$. The covariance between two random variables \mathbf{z}_1 and \mathbf{z}_2 is defined as

$$\text{cov}(\mathbf{z}_1, \mathbf{z}_2) = \mathbb{E}[(\mathbf{z}_1 - \mathbb{E}[\mathbf{z}_1])(\mathbf{z}_2 - \mathbb{E}[\mathbf{z}_2])].$$

As standard, we say that two random variables are uncorrelated if their covariance is zero. For a vector-valued random variable \mathbf{x} taking values in \mathbb{R}^n we define the (co-)variance matrix as $\text{var}(\mathbf{x}) = \mathbb{E}[(\mathbf{x} - \mathbb{E}[\mathbf{x}])(\mathbf{x} - \mathbb{E}[\mathbf{x}])^T]$. Let \mathbf{z}_1 be a

¹The index \mathbb{X} in $\mathcal{F}_{\mathbb{X}}$ distinguishes the given σ -algebra on \mathbb{X} from that on \mathbb{Y} , which is denoted as $\mathcal{F}_{\mathbb{Y}}$. Whenever possible this index will be dropped.

$(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ -valued random variable and let \mathbf{z}_2 be a $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ -valued random variable, then \mathbf{z}_1 and \mathbf{z}_2 are independent if for all $Z_1, Z_2 \in \mathcal{B}(\mathbb{R})$:

$$\mathbb{P}(\mathbf{z}_1 \in Z_1 \wedge \mathbf{z}_2 \in Z_2) = \mathbb{P}(\mathbf{z}_1 \in Z_1)\mathbb{P}(\mathbf{z}_2 \in Z_2).$$

A timed series of random variables $\mathbf{x}(t)$ with $t \in \mathbb{N}$ defines a stochastic process. A stochastic process $\mathbf{x}(t)$ is called stationary if $\mathbf{x}(t)$ has the same probability distribution as $\mathbf{x}(t + \tau)$ for all $\tau \in \mathbb{N}$. A white noise process is a stationary process consisting of uncorrelated random variables with zero mean and finite variance; that is, over time the signal has independent and identically distributed random variables.

When it is of interest to analyse properties of signals averaged over time, we can work with signals that are not fully stochastic and stationary, but also include a deterministic part, as if they are stationary stochastic signals. For this we employ the notion of quasi-stationary signals as given by Ljung [1999].

Definition 2.6 (Quasi-stationary) *A signal $\{\mathbf{y}(t)\}$ is said to be quasi-stationary if it is subject to*

$$\begin{aligned} \mathbb{E}[\mathbf{y}(t)] &= m_y(t), \quad |m_y(t)| \leq C \quad \text{for all } t; \\ \mathbb{E}[\mathbf{y}(t)\mathbf{y}(r)] &= R_y(t, r), \quad R_y(t, r) \leq C; \\ \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N R_y(t, t - \tau) &= R_y(\tau), \quad \text{for all } \tau \in \mathbb{N}. \end{aligned}$$

We denote with $\bar{\mathbb{E}}$ the generalised expectation operation defined as

$$\bar{\mathbb{E}}[\mathbf{y}(t)] = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E}[\mathbf{y}(k)]. \quad (2.2)$$

Now the correlation function of a quasi stationary signal is defined as $R_y(\tau) := \bar{\mathbb{E}}[\mathbf{y}(t)\mathbf{y}(t - \tau)]$, then the power spectral density of a signal is defined as

$$\Phi_y(\omega) := \sum_{\tau=-\infty}^{\infty} R_y(\tau) e^{i\omega\tau}.$$

The power of a signal is defined as $\mathcal{P}_y := \bar{\mathbb{E}}[\mathbf{y}^2(t)]$ and note that

$$\bar{\mathbb{E}}[\mathbf{y}^2(t)] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_y(\omega) d\omega.$$

2.4.b Markov decision processes and properties

Consider a linear time invariant model as in (2.1) disturbed by Gaussian noise, then the state updates are given as

$$x(t+1) = Ax(t) + Bu(t) + B_w w(t), \quad (2.3)$$

with $w(t)$ a white noise sequence with a standard Gaussian distribution, that is, at each time instance $w(t)$ is a realisation of a zero mean and unit variance

Gaussian distribution. This is denoted as $w(t) \sim \mathcal{N}(0, I)$. Since $w(t)$ is a white noise signal, this means that at any time instant given the state $x(t)$ the next state is conditionally independent from the past states and the past noise inputs $w(t-1), w(t-2), \dots$. This means that the model of the state evolutions satisfies the Markov property.

Models like (2.3) satisfying the Markov property and driven by an external control input or action can be collected within the notion of Markov decision processes (MDP) [Bertsekas and Shreve, 1996, Meyn and Tweedie, 1993, Hernández-Lerma and Lasserre, 1996] defined as follows. The tuple $\mathbf{M} = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U})$ defines a discrete-time MDP over an uncountable state space \mathbb{X} ; and is characterised by \mathbb{T} , a conditional stochastic kernel that assigns to each point $x \in \mathbb{X}$ and control $u \in \mathbb{U}$ a probability measure $\mathbb{T}(\cdot | x, u)$ over $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. More precisely, for any set $A \in \mathcal{B}(\mathbb{X})$, $\mathbb{P}_{x,u}(x(t+1) \in A) = \int_A \mathbb{T}(dy | x(t) = x, u)$, where $\mathbb{P}_{x,u}$ denotes the conditional probability $\mathbb{P}(\cdot | x, u)$. In general, we will require the stochastic kernel $\mathbb{T}(\cdot | x, u)$ to be measurable mapping. The initial probability distribution is $\pi_{x(0)} : \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$. Markov decision processes will be formally introduced in Chapter 3 in Definition 3.1 and repeated later in Chapter 7. As standard, we can select actions or controls $u \in \mathbb{U}$ based on a Markov policy. Within stochastic optimal control, this policy is often defined as a mapping from state to actions. More precisely, a Markov policy μ over the horizon $[0, N]$ is a sequence $\mu = (\mu_0, \mu_1, \dots, \mu_{N-1})$ of universally measurable maps

$$\mu_t : \mathbb{X} \rightarrow \mathbb{U}, \quad t \in \mathbb{N}_{N-1} := \{0, 1, \dots, N-1\},$$

from the state space \mathbb{X} to the action space \mathbb{U} . In a more general setting, necessary for more complicated control objectives, policies can also be stochastic.

Definition 2.7 (Markov policy) *For an MDP $\mathbf{M} = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U})$, a Markov policy μ is a sequence $\mu = (\mu_1, \mu_2, \mu_3, \dots)$ of universally measurable maps $\mu_t : \mathbb{X} \rightarrow \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ $t = 0, 1, 2, \dots$ from the state space \mathbb{X} to the set of controls.*

Note that this definition of a stochastic Markov policy also trivially includes the deterministic Markov policies. The stochastic notion of a Markov policy will be used in Chapter 3 Definition 3.3.

We consider the evolution of the Markov decision process for $t = 0, 1, \dots, N$ with $N \in \mathbb{N}$. For a given initial state and Markov policy μ and a given initial state $x_0 \in \mathbb{X}$, an execution or realisation of the process characterises a state trajectory given as $\{x(t) | t \in \mathbb{N}_N\}$. At each time instant t , given $x(t)$ and the policy μ_t , the next state is a realisation of the controlled Borel-measurable stochastic kernel $\mathbb{T}(\cdot | x(t), \mu_t(x(t)))$.

The execution $\{x(t) | t \in \mathbb{N}_N\}$ initialised by $x_0 \in \mathbb{X}$ drawn from $\pi_{x(0)}$ and controlled with Markov policy μ is a stochastic process defined on the canonical sample space $\Omega := \mathbb{X}^{N+1}$ endowed with its product topology $\mathcal{B}(\Omega)$. This stochastic process has a probability measure \mathbb{P} uniquely defined by the transition kernel \mathbb{T} , policy μ , and initial distribution $\pi_{x(0)}$ [Bertsekas and Shreve, 1996, Prop. 7.45].

Note that once controlled, the MDP is a Markov process since all non-determinism is resolved.

Similar to the deterministic framework for modelling systems, we are interested in temporal properties defined over the behaviour of the MDPs. For this, we first define simpler properties such as probabilistic invariance (or equivalently safety), reachability, and reach-avoid properties. The probabilistic invariance property is defined for a given time horizon N and safe set $A \in \mathcal{B}(\mathbb{X})$ as the probability that the realisation $\{x(t)|t \in \mathbb{N}_N\}$ of a controlled MDP stays in A

$$\mathbb{P}(\forall t \in \mathbb{N}_N : x(t) \in A).$$

Similarly the probabilistic reachability property for a given target set $K \in \mathcal{B}(\mathbb{X})$ is defined as the probability that the realised trajectory $\{x(t)|t = 0, 1, \dots, N\}$ reaches K that is

$$\mathbb{P}(\exists j \in \mathbb{N}_N : x(j) \in K).$$

The probability of reaching K while staying in A and before the end of the time horizon N is defined as a probabilistic reach-avoid property, quantified as

$$\mathbb{P}(\exists j \in \mathbb{N}_N : x(j) \in K \wedge \forall t \in \mathbb{N}_{j-1} : x(t) \in A),$$

where \wedge denotes the logical “and” operator. The complement of A is then the “unsafe” set that is avoided in the probabilistic reach-avoid property. Note that beyond the set of simple safety or reachability properties, LTL specifications can often be rewritten as stochastic reach-avoid properties over extended Markov decision processes [Lahijanian et al., 2009].

Based on the above, a verification problem formulated on a Markov process is the verification whether the probability of one of the above properties is above, below or within a bounded region of a required probability p . A synthesis problem, considers the computation of a policy for an MDP such that the probability in a property of interest is maximised or above a required threshold. For continuous, or more generally uncountably-infinite, state-spaces the computation of a policy, even when performed approximately, will induce heavy computational load. In Chapter 3 we circumvent the direct computation of policies and their probabilistic quantification. We show that a policy computed for an abstraction of the MDP, which is of reduced order and allows for simpler computations, can be refined to the original MDP with guarantees.

Beyond the case of probabilistic state-transitions, one could also have the problem that the state cannot always be observed exactly. In other words we could have that the state of the model (2.3) can only be measured via noisy sensor measurements $y(t)$ modelled as

$$y(t) = Cx(t) + Du(t) + D_w w(t), \quad (2.4)$$

where w is a white noise sequence and generally we take $D = 0$. For the more general case of partially observable Markov decision processes it is known that the verification and controller synthesis is very difficult. Instead, we will tackle controller synthesis for partially observable LTI systems modelled as (2.3-2.4) in Chapter 4.

Part I

Control

There should be no such thing as boring mathematics.

Edsger Dijkstra

3

Approximate similarity relations and controller refinement for general Markov decision processes

In this chapter we introduce new approximate similarity relations that are shown to be key for policy (or control) synthesis over general Markov decision processes. The models of interest are discrete-time Markov decision processes, endowed with uncountably-infinite state spaces and metric output (or observation) spaces. The new relations, underpinned by the use of metrics, allow in particular for a useful trade-off between deviations over probability distributions on states, and distances between model outputs. We show that the new probabilistic similarity relations, inspired by a notion of simulation developed for finite-state models, can be effectively employed over general Markov decision processes for verification purposes, and specifically for control refinement from abstract models.

3.1 Introduction

The formal verification of computer systems allows for the quantification of their properties and for their correct functioning. Whilst the bulk of verification methods has classically focused on finite- or countable-state models, with the ever more ubiquitous embedding of digital components into physical systems richer models

are needed and correct functioning can only be expressed over the combined behaviour of both the digital computer and the surrounding physical system. It is in particular of interest to synthesise the part of the computer software that controls or interacts with the physical system automatically, with low likelihood of malfunctioning. Furthermore, when computers interact with physical systems such as biological processes, power networks, and smart-grids, stochastic models are key. Consider, as an example, a power network for which we would like to quantify the likelihood of blackouts and to synthesise strategies to minimise this.

Systems with uncertainty and non-determinism can be naturally modelled as Markov decision processes (MDP). In this work, we focus on general Markov decision processes (gMDP) that have uncountable state spaces as well as metric output spaces. The characterisation of properties over such processes cannot in general be attained analytically [Abate et al., 2008], so an alternative is to approximate these models by simpler processes that are prone to be mathematically analysed or algorithmically verified [Esmail Zadeh Soudjani and Abate, 2013], such as finite-state MDP [Esmail Zadeh Soudjani et al., 2015]. Clearly, it is then key to provide formal guarantees on this approximation step, such that solutions of the verification or synthesis problem for a property on the simpler process can be extended to the original model. Our verification problems include the synthesis of a policy (or a control strategy) that maximises the likelihood of the specification of interest.

In this work we develop a new notion of approximate similarity relation, aimed to attain a computationally efficient controller synthesis over Markov decision processes with metric output spaces. We show that it is possible to obtain a control strategy for a gMDP as a refinement of a strategy synthesised for an abstract model, at the expense of accuracy defined on a similarity relation between them, which quantifies bounded deviations in transition probabilities and output distances. In summary, we provide results allowing us to quantitatively relate the outcome of verification problems performed over the simpler (abstract) model to the original (concrete) model, and further to refine control strategies synthesised over the abstract model to strategies for the original model.

The use of similarity relations on probabilistic models with countable states has been broadly investigated, either via exact notions of probabilistic simulation and bisimulation relations [Larsen and Skou, 1989, Segala, 1995, Segala and Lynch, 1995, D’Argenio et al., 2002], or (more recently) via approximate notions [Desharnais et al., 2008] and [D’Innocenzo et al., 2012]. On the other hand, similar notions over *general, uncountable-state spaces* have been only recently studied: available relations either hinge on stability requirements on model outputs [Julius and Pappas, 2009, Zamani et al., 2014] (established via martingale theory or contractivity analysis), or alternatively enforce structural abstractions of a model [Desharnais et al., 2004] by exploiting continuity conditions on its probability laws [Abate, 2013, Abate et al., 2014a].

In this work, we want to quantify properties with a certified precision *both* in the deviation of the probability laws for finite-time events (as in the classical notion of probabilistic bisimulation) and of the output trajectories (as studied for dynamical models). Additionally, we impose no strict requirements on the dynamics

of the given gMDP and its abstraction. To these ends, we consider the type of exact probabilistic simulation and bisimulation relations originally introduced in [Larsen and Skou, 1989, Jonsson and Larsen, 1991]. We first extend the lifting based definition of these exact probabilistic simulation and bisimulation relations for finite-state probabilistic automata and stochastic games [Segala, 1995, Segala and Lynch, 1995, Zhang and Pang, 2010] to gMDP (Section 3.3). We then generalise these notions to allow for errors on the probability laws *and* deviations over the output space (Section 3.4). Two case studies in the area of smart buildings (Section 3.5) are used to evaluate these new approximate probabilistic simulation relations. Unlike cognate recent work [Abate, 2013, Julius and Pappas, 2009], we are interested in similarity relations that allow refining over the concrete model a control strategy synthesised on the abstract one. We zoom in on relations that, quite like the alternating notions in [Alur et al., 1998, Tabuada, 2009] for non-probabilistic models and in [Zhang and Pang, 2010] for stochastic ones, quantitatively bound the difference in the controllable behaviour of pairs of models (namely a gMDP and its abstraction). In Section 3.6 we show how over a class of Markov processes (without controls), this newly developed approximate similarity relation practically generalises notions of probabilistic (bi-)simulations of Labeled Markov processes Desharnais et al. [2002, based on zigzag-morphisms], Desharnais et al. [2003, based on equivalence relations], and their approximate versions Desharnais et al. [2004, 2008], D’Innocenzo et al. [2012, based on binary relations].

In the following section, the problem statement is given in more details. Subsequently, we first provide a new characterisation of similarity relations for controlled Markov processes with general state spaces and give its relevant controller refinement in Section 3.3. Secondly, the new notion of approximate simulation relations is given in Section 3.4, where we additionally also provide properties of these similarity relations. This is followed with two clarifying case studies in Section 3.5. In the penultimate section of this chapter we tackle connections to literature and measurability issues.

3.2 Verification of general Markov decision processes: problem setup

3.2.a gMDP models - syntax and semantics

General Markov decision processes are related to control Markov processes [Abate, 2013] and Markov decision processes [Bertsekas and Shreve, 1996, Meyn and Tweedie, 1993, Hernández-Lerma and Lasserre, 1996], and formalised as follows.

Definition 3.1 (Markov decision process (MDP)) *The tuple $\mathbf{M} = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U})$ defines a discrete-time MDP over an uncountable state space \mathbb{X} , and is characterised by \mathbb{T} , a conditional stochastic kernel that assigns to each point $x \in \mathbb{X}$ and control $u \in \mathbb{U}$ a probability measure $\mathbb{T}(\cdot \mid x, u)$ over $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. The initial probability distribution is $\pi_{x(0)} : \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$.*

At every state the state transition depends non-deterministically on the choice of $u \in \mathbb{U}$. When chosen according to a distribution $\mu_u : \mathcal{B}(\mathbb{U}) \rightarrow [0, 1]$, we refer to the stochastic control input as μ_u . Moreover the transition kernel is denoted as $\mathbb{T}(\cdot | x, \mu_u) = \int_{\mathbb{U}} \mathbb{T}(\cdot | x, u) \mu_u(du) \in \mathcal{P}(\mathbb{X}, \mathcal{B}(\mathbb{X}))$. Given a string of inputs $u(0), u(1), \dots, u(N)$, over a finite time horizon $\{0, 1, \dots, N\}$, and an initial condition x_0 (sampled from distribution $\pi_{x(0)}$), the state at the $(t + 1)$ -st time instant, $x(t + 1)$, is obtained as a realisation of the controlled Borel-measurable stochastic kernel $\mathbb{T}(\cdot | x(t), u(t))$ – these semantics induce paths (or executions) of the MDP.

Definition 3.2 (General Markov decision process (gMDP))

$\mathbf{M} = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U}, h, \mathbb{Y})$ is a discrete-time gMDP consisting of an MDP combined with output space \mathbb{Y} and a measurable output mapping $h : \mathbb{X} \rightarrow \mathbb{Y}$. A metric $\mathbf{d}_{\mathbb{Y}}$ decorates the output space \mathbb{Y} .

The gMDP semantics are directly inherited from those of the MDP. Further, output traces of gMDP are obtained as mappings of MDP paths, namely $\{y(t)\}_{0:N} := y(0), y(1), \dots, y(N)$, where $y(t) = h(x(t))$. Denote the class of all gMDP with the metric output space \mathbb{Y} as $\mathcal{M}_{\mathbb{Y}}$. Note that gMDP can be regarded as a superclass of the known labelled Markov processes (LMP) [Desharnais et al., 2004] as elucidated in [Abate et al., 2014a].

Example 3.1 Consider the stochastic process

$$\mathbf{M} : x(t + 1) = f(x(t), u(t)) + e(t), \quad y(t) = h(x(t)) \in \mathbb{Y},$$

with variables $x(t), u(t), e(t)$, taking values in \mathbb{R}^n , representing the state, control input¹, and noise terms respectively. The process is initialised as $x(0) \sim \pi_{x(0)}$, and driven by $e(t)$, a white noise sequence with zero-mean normal distributions and covariance matrix Σ_e . This stochastic process, defined as a dynamical model, is a gMDP characterised by a tuple $(\mathbb{R}^n, \pi_{x(0)}, \mathbb{T}, \mathbb{R}^n, h, \mathbb{Y})$, where the conditional transition kernel is defined as $\mathbb{T}(\cdot | x, u) = \mathcal{N}(f(x(t), u(t)), \Sigma_e)$, a normal probability distribution with mean $f(x(t), u(t))$ and covariance matrix Σ_e .

A policy is a selection of control inputs based on the past history of states and controls. We allow controls to be selected via universally measurable maps [Bertsekas and Shreve, 1996] from the state to the control space, so that time-bounded properties such as safety can be maximised [Abate et al., 2008]. When the selected controls are only dependent on the current states, and thus conditionally independent of history (or memoryless), the policy is referred to as Markov.

Definition 3.3 (Markov policy) For a gMDP $\mathbf{M} = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U}, h, \mathbb{Y})$, a Markov policy μ is a sequence $\mu = (\mu_0, \mu_1, \mu_2, \mu_3, \dots)$ of universally measurable maps $\mu_t = \mathbb{X} \rightarrow \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ $t = 0, 1, 2, \dots$, from the state space \mathbb{X} to the set of controls.

¹ In other domains one also refers to the control variables as actions (Machine Learning, Stochastic Games) or as external non-determinism (Computer science).

Recall that a function $f : \mathbb{K}_1 \rightarrow \mathbb{K}_2$ is universally measurable if the inverse image of every Borel set is measurable with respect to every complete probability measure on \mathbb{K}_1 that measures all Borel subsets of \mathbb{K}_1 .

The execution $\{x(t), t \in [0, N]\}$ initialised by $x_0 \in \mathbb{X}$ and controlled with Markov policy μ is a stochastic process defined on the canonical sample space $\Omega := \mathbb{X}^{N+1}$ endowed with its product topology $\mathcal{B}(\Omega)$. This stochastic process has a probability measure \mathbb{P} uniquely defined by the transition kernel \mathbb{T} , policy μ , and initial distribution $\pi_{x(0)}$ [Bertsekas and Shreve, 1996, Prop. 7.45].

Of interest are time-dependent properties such as those expressed as specifications in a temporal logic of choice. This leads to problems where one maximises the probability that a sequence of labelled sets is reached within a time limit and in the right order. One can intuitively realise that in general the optimal policy leading to the maximal probability is not a Markov (memoryless) policy, as introduced in Def. 3.3. We introduce the notion of a *control strategy*, and define it as a broader, memory-dependent version of the Markov policy above. This strategy is formulated as a Markov process that takes as an input the state of the to-be-controlled gMDP.

Definition 3.4 (Control strategy) *A control strategy $\mathbf{C} = (\mathbb{X}_{\mathbf{C}}, x_{\mathbf{C}0}, \mathbb{X}, \mathbb{T}_{\mathbf{C}}^t, h_{\mathbf{C}}^t)$ for a gMDP \mathbf{M} with state space \mathbb{X} and control space \mathbb{U} over the time horizon $t = 0, 1, 2, \dots, N$ is an inhomogenous Markov process with state space $\mathbb{X}_{\mathbf{C}}$; an initial state $x_{\mathbf{C}0}$; inputs $x \in \mathbb{X}$; time-dependent, universally measurable kernels $\mathbb{T}_{\mathbf{C}}^t$, $t = 0, 1, \dots, N$; and with universally measurable output maps $h_{\mathbf{C}}^t : \mathbb{X}_{\mathbf{C}} \rightarrow \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$, $t = 0, 1, \dots, N$, with elements $\mu \in \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$.*

Unlike a Markov policy, the control strategy is in general dependent on the history, as it has an internal state that can be used to remember relevant past events. As elucidated in Algorithm 1, note that the first control $u(0)$ is selected by drawing $x_{\mathbf{C}}(1)$ according to $\mathbb{T}_{\mathbf{C}}^0(\cdot | x_{\mathbf{C}}(0), x(0))$, where $x_{\mathbf{C}}(0) = x_{\mathbf{C}0}$, and selecting $u(0)$ from measure $\mu_{\mathbf{C}}^0 = h_{\mathbf{C}}^0(x_{\mathbf{C}}(1))$.² The control strategy applied to \mathbf{M} can be both stochastic (as a realisation of $\mathbb{T}_{\mathbf{C}}^0(\cdot | x_{\mathbf{C}}(0), x(0))$), a function of the initial state $x(0)$, and of time.

The execution $\{(x(t), x_{\mathbf{C}}(t)), t \in [0, N]\}$ of a gMDP \mathbf{M} controlled with strategy \mathbf{C} is defined on the canonical sample space $\Omega := (\mathbb{X} \times \mathbb{X}_{\mathbf{C}})^{N+1}$ endowed with its product topology $\mathcal{B}(\Omega)$. This stochastic process is associated to a unique probability measure $\mathbb{P}_{\mathbf{C} \times \mathbf{M}}$, since the stochastic kernels $\mathbb{T}_{\mathbf{C}}^t$ for $t \in [0, N]$ and \mathbb{T} are Borel measurable and composed via universally measurable policies [Bertsekas and Shreve, 1996, Prop. 7.45].

²Note that the stochastic transitions for the control strategy and the gMDP are selected in an alternating fashion. The output map of the strategy is indexed based on the time instant at which the resulting policy will be applied to the gMDP.

Algorithm 1 Execution of the controlled model $\mathbf{C} \times \mathbf{M}$

```

set  $t := 0$  and  $x_{\mathbf{C}}(0) := x_{\mathbf{C}0}$ 
draw  $x(0) \sim \pi_{x(0)}$  {from  $\mathbf{M}$ }
while  $t < N$  do
  draw  $x_{\mathbf{C}}(t+1) \sim \mathbb{T}_{\mathbf{C}}^t(\cdot | x_{\mathbf{C}}(t), x(t))$  {from  $\mathbf{C}$ }
  set  $\mu_t := h_{\mathbf{C}}^t(x_{\mathbf{C}}(t+1))$ , draw  $u(t)$  from  $\mu_t$ 
  draw  $x(t+1) \sim \mathbb{T}(\cdot | x(t), u(t))$  {from  $\mathbf{M}$ }
  set  $t := t + 1$ 
end while

```

3.2.b gMDP verification and strategy refinement: problem statement

We qualitatively introduce the main problem that we want to solve in this work: How can one provide a general framework to synthesise control policies over a formal abstraction $\tilde{\mathbf{M}}$ of a concrete complex model \mathbf{M} , with the understanding that $\tilde{\mathbf{M}}$ is much simpler to be manipulated (analytically or computationally) than \mathbf{M} is? We approach this problem by defining a simulation relation under which a control strategy $\tilde{\mathbf{C}}$ for the abstract Markov process $\tilde{\mathbf{M}}$ implies the existence of a strategy \mathbf{C} for \mathbf{M} , so that we can quantify differences in the stochastic transition kernels and in the output trajectories for the two controlled models. This allows us to derive bounds on the probability of satisfaction of a specification for $\mathbf{C} \times \mathbf{M}$ from the satisfaction probability of modified specifications for $\tilde{\mathbf{C}} \times \tilde{\mathbf{M}}$. We will show that with this setup we can deal with finite-horizon temporal properties, including safety verification as a relevant instance.

The results in this paper are to be used in parallel with optimisation, both for selecting the control refinement and for synthesising a policy on the abstract model. It has been shown in [Bertsekas and Shreve, 1996] that stochastic optimal control even for a system on a “basic” space can lead to measurability issues: in order to avoid these issues we follow [Bertsekas and Shreve, 1996, Desharnais et al., 2008] and the developed theory for Polish spaces and Borel (or universally) measurable notions. Throughout the paper we will give as clarifying examples Markov processes evolving, as in Example 3.1, over Euclidean spaces that are a special instances of Polish spaces. This allows us to elucidate the theory.

3.3 Exact (bi-)simulation relations based on lifting**3.3.a Introduction**

In this section, we define probabilistic simulation and bisimulation relations that are, respectively, a preorder and an equivalence relation on $\mathcal{M}_{\mathbb{Y}}$. Before introducing these relations, we first extend Segala’s notion [Segala, 1995, Segala and Lynch, 1995] of *lifting* to uncountable state spaces, which allows us to equate the

transition kernels of two given gMDPs. Thereafter, we leverage liftings to define (bi-)simulation relations over $\mathcal{M}_{\mathbb{Y}}$, which characterise the similarity in the controllable behaviours of the two gMDPs. Subsequently we show that these similarity relations also imply controller refinement, i.e., within the similarity relation a control strategy for a given gMDP can be refined to a controller for another gMDP. In the next section, we show that this exact notion of similarity allows a more general notion of approximate probabilistic simulation. The new notions of similarity relations extend the known exact notions in [Larsen and Skou, 1989], and the approximate notions of [Desharnais et al., 2008, D’Innocenzo et al., 2012]. Additionally, we will show that these results can be naturally extended to allow for both differences in probability and deviations in the outputs of the two gMDPs.

We work with pairs of gMDPs put in a relationship, denoting them with numerical indices (M_1, M_2) , with the intention to apply the developed notions to an abstraction \tilde{M} of a concrete model M , respectively.

3.3.b Lifting for general Markov decision processes

Consider two gMDP $M_1, M_2 \in \mathcal{M}_{\mathbb{Y}}$ mapping to a common output space \mathbb{Y} with metric $d_{\mathbb{Y}}$. For $M_1 = (\mathbb{X}_1, \pi_{x(0)_1}, \mathbb{T}_1, \mathbb{U}_1, h_1, \mathbb{Y})$ and $M_2 = (\mathbb{X}_2, \pi_{x(0)_2}, \mathbb{T}_2, \mathbb{U}_2, h_2, \mathbb{Y})$ at given state-action pairs $x_1 \in \mathbb{X}_1, u_1 \in \mathbb{U}_1$ and $x_2 \in \mathbb{X}_2, u_2 \in \mathbb{U}_2$, respectively, we want to relate the corresponding transition kernels, namely the probability measures $\mathbb{T}_1(\cdot \mid x_1, u_1) \in \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\mathbb{T}_2(\cdot \mid x_2, u_2) \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$.

Similar to the coupling of measures in $\mathcal{P}(\mathbb{X}, \mathcal{F})$ [Abate et al., 2014b, Lindvall, 2002], consider the *coupling* of two arbitrary probability spaces $(\mathbb{X}_1, \mathcal{F}_1, \mathbb{P}_1)$ and $(\mathbb{X}_2, \mathcal{F}_2, \mathbb{P}_2)$ (cf. [Skala, 1993, Strassen, 1965]). A probability measure \mathbb{P}_c defined on $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{F})$ *couples* the two spaces if the projections p_1, p_2 , with $x_1 = p_1(x_1, x_2)$ and $x_2 = p_2(x_1, x_2)$, define respectively an $(\mathbb{X}_1, \mathcal{F}_1)$ - and an $(\mathbb{X}_2, \mathcal{F}_2)$ -valued random variables, such that $\mathbb{P}_1 = p_{1*}\mathbb{P}_c$ and $\mathbb{P}_2 = p_{2*}\mathbb{P}_c$. For *finite- or countable-state* stochastic processes a related concept has been referred to as *lifting* in [Segala, 1995, Segala and Lynch, 1995]: the transition probabilities are coupled using a weight function in a way that respects a given relation over the combined state spaces. Rather than using weight functions over a countable or finite domain [Segala, 1995], we introduce lifting as a coupling of measures over Polish space and their corresponding Borel measurable σ -fields.

Since we assume that the state spaces are Polish and have a corresponding Borel σ -field for the given probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1), \mathbb{P}_1)$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2), \mathbb{P}_2)$ with $\mathbb{P}_1 := \mathbb{T}_1(\cdot \mid x_1, u_1)$ and $\mathbb{P}_2 := \mathbb{T}_2(\cdot \mid x_2, u_2)$, the natural choice for the σ -algebra becomes $\mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2) = \mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2)$ ³ and the question of finding a coupling can be reduced to finding a probability measure in $\mathcal{P}(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2))$.

Definition 3.5 (Lifting for general state spaces) *Let $\mathbb{X}_1, \mathbb{X}_2$ be two sets with associated measurable spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ and let the Borel measurable set $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ be a relation. We denote by $\tilde{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$*

³ $\mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2)$ denotes the product σ -algebra of $\mathcal{B}(\mathbb{X}_1)$ and $\mathcal{B}(\mathbb{X}_2)$.

the corresponding lifted relation, so that $\Delta \bar{\mathcal{R}} \Theta$ holds if there exists a probability space $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ (equivalently, a lifting \mathbb{W}) satisfying

1. for all $X_1 \in \mathcal{B}(\mathbb{X}_1)$: $\mathbb{W}(X_1 \times \mathbb{X}_2) = \Delta(X_1)$;
2. for all $X_2 \in \mathcal{B}(\mathbb{X}_2)$: $\mathbb{W}(\mathbb{X}_1 \times X_2) = \Theta(X_2)$;
3. for the probability space $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ it holds that $x_1 \mathcal{R} x_2$ with probability 1, or equivalently that $\mathbb{W}(\mathcal{R}) = 1$.

With reference to the connection with the notion of coupling, an equivalent definition of lifting is obtained by replacing 1. and 2. by the condition that for $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ the projections p_1, p_2 , with $x_1 = p_1(x_1, x_2)$ and $x_2 = p_2(x_1, x_2)$, we can define $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ -valued random variables $\Delta = p_{1*} \mathbb{W}$ and $\Theta = p_{2*} \mathbb{W}$. An example is portrayed in Figure 3.1 containing two models M_1, M_2 and a relation (denoted by equally labelled/coloured pairs of states), where the transition kernels for a pair of states is lifted with respect to the relation.

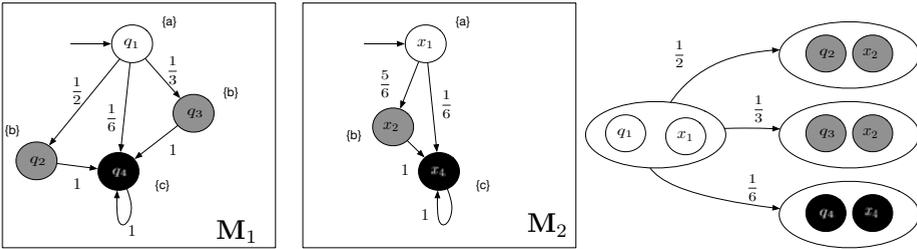


Figure 3.1: Finite-state Markov processes M_1 and M_2 (left & middle) with $S = \{q_1, q_2, q_3, q_4\}$ and $T = \{x_1, x_2, x_4\}$ the respective state spaces. The states are labelled with three different colours. Lifting probabilities of the transition kernels for (q_1, x_1) are given on the edges of the rightmost figure.

Remark 3.2 Notice that the extension of the notion of lifting to general spaces has required the use of measures, rather than weight functions over a countable or finite domain, as in [Segala, 1995]. We have required that the σ -algebra $\mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$ contains not only sets of the form $X_1 \times \mathbb{X}_2$ and $\mathbb{X}_1 \times X_2$, but also specifically the sets that characterise the relation \mathcal{R} . Since the spaces \mathbb{X}_1 and \mathbb{X}_2 have been assumed to be Polish, it holds that every open (closed) set in $\mathbb{X}_1 \times \mathbb{X}_2$ belongs to $\mathcal{B}(\mathbb{X}_1) \otimes \mathcal{B}(\mathbb{X}_2) = \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$ [Bogachev, 2007, Lemma 6.4.2]. As an instance consider the diagonal relation $\mathcal{R}_{diag} := \{(x, x) : x \in \mathbb{X}\}$ over $\mathbb{X} \times \mathbb{X}$, of importance for examples introduced later. This is a Borel measurable set [Bogachev, 2007, Theorem 6.5.7].

3.3.c Exact probabilistic (bi-)simulation relations via lifting

Similar to the alternating notions for probabilistic game structures in [Zhang and Pang, 2010], we provide a simulation that relates any input chosen for the first

process with one for the second process. As such, we allow for more elaborate handling of the inputs than in the probabilistic simulation relations discussed in [Desharnais et al., 2008, D’Innocenzo et al., 2012], and further pave the way towards the inclusion of output maps. We extend the notions in [Segala, 1995, Zhang and Pang, 2010, Jonsson and Larsen, 1991] by allowing for more general Polish spaces. Further, we introduce the notion of *interface function* in order to connect the controllable behaviour of two gMDP:

$$\mathcal{U}_v : \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2 \rightarrow \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2)),$$

where we require that \mathcal{U}_v is a Borel measurable function. This means that \mathcal{U}_v induces a Borel measurable stochastic kernel, again denoted by \mathcal{U}_v , over \mathbb{U}_2 given $(u_1, x_1, x_2) \in \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$. The notion of interface function is known in the context of correct-by-design controller synthesis and of hierarchical controller refinement [Girard and Pappas, 2009, Tabuada, 2009]. For the objective of hierarchical controller refinement, an interface function implements (or refines) any control action synthesised over the abstract model to an action for the concrete model. In order to establish an exact simulation relation between abstract and concrete models, we can attempt to refine the control actions from one model to the other by choosing an interface function that matches their stochastic behaviours. On the other hand in the next section, the interface function will be used to establish approximate simulation relations: for this goal, the optimal selection of the interface function is the one that optimises the accuracy of the relation. This is topic of ongoing research.

In this work we extend standard interface functions for deterministic systems by allowing randomised actions $\mu_2 \in \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2))$. The lifting of the transition kernels for the chosen interface generates a stochastic kernel $\mathbb{W}_{\mathbb{T}}$ conditional on the values of signals in \mathbb{U}_1 and in $\mathbb{X}_1 \times \mathbb{X}_2$. Let us trivially extend the interface function to $\mathcal{U}_v(\mu_1, x_1, x_2) := \int_{\mathbb{U}_1} \mathcal{U}_v(u_1, x_1, x_2) \mu_1(du_1)$.

Definition 3.6 (Probabilistic simulation) *Consider two gMDP $\mathbf{M}_i, i = 1, 2$, $\mathbf{M}_i = (\mathbb{X}_i, \pi_{x(0)_i}, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$. The gMDP \mathbf{M}_1 is stochastically simulated by \mathbf{M}_2 if there exists an interface function \mathcal{U}_v and a relation $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ with $\mathcal{R} \in \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$, for which there exists a Borel measurable stochastic kernel $\mathbb{W}_{\mathbb{T}}(\cdot | u_1, x_1, x_2)$ on $\mathbb{X}_1 \times \mathbb{X}_2$ given $\mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$, such that*

1. $\forall (x_1, x_2) \in \mathcal{R}, h_1(x_1) = h_2(x_2)$;
2. $\forall (x_1, x_2) \in \mathcal{R}, \forall u_1 \in \mathbb{U}_1, \mathbb{T}_1(\cdot | x_1, u_1) \bar{\mathcal{R}} \mathbb{T}_2(\cdot | x_2, \mathcal{U}_v(u_1, x_1, x_2))$, with lifted probability measure $\mathbb{W}_{\mathbb{T}}(\cdot | u_1, x_1, x_2)$;
3. $\pi_{x(0)_1} \bar{\mathcal{R}} \pi_{x(0)_2}$.

The relationship between the two models is denoted as $\mathbf{M}_1 \preceq \mathbf{M}_2$.

The Borel measurability for both \mathcal{U}_v (see above) and $\mathbb{W}_{\mathbb{T}}$ (as in this definition), which is technically needed for the well posedness of the controller refinement, can be relaxed to universal measurability, as will be discussed in the Appendix.

Definition 3.7 (Probabilistic bisimulation) *Under the same conditions as above, \mathbf{M}_1 is a probabilistic bisimulation of \mathbf{M}_2 if there exists a relation $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ such that $\mathbf{M}_1 \preceq \mathbf{M}_2$ w.r.t. \mathcal{R} and $\mathbf{M}_2 \preceq \mathbf{M}_1$ w.r.t. the inverse relation $\mathcal{R}^{-1} \subseteq \mathbb{X}_2 \times \mathbb{X}_1$. \mathbf{M}_1 and \mathbf{M}_2 are said to be probabilistically bisimilar, which is denoted $\mathbf{M}_1 \approx \mathbf{M}_2$.*

For every gMDP \mathbf{M} : $\mathbf{M} \preceq \mathbf{M}$ and $\mathbf{M} \approx \mathbf{M}$. This can be seen by considering the diagonal relation $\mathcal{R}_{diag} = \{(x_1, x_2) \in \mathbb{X} \times \mathbb{X} \mid x_1 = x_2\}$ and selecting equal inputs for the associated interfaces. The resulting equal transition kernels

$$\mathbb{T}(\cdot|x, u)\bar{\mathcal{R}}_{diag}\mathbb{T}(\cdot|x, u)$$

are lifted by the measure $\mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_2|u, x_1, x_2) = \delta_{x'_1}(dx'_2)\mathbb{T}(dx'_1|x_1, u)$ where $\delta_{x'_1}$ denotes the Dirac distribution located at x'_1 .

Example 3.3 (Lifting for diagonal relations)

a. Consider the gMDP (\mathbf{M}_1) introduced in Ex. 3.1 and a slight variation of it (\mathbf{M}_2), given as stochastic dynamic processes,

$$\begin{aligned} \mathbf{M}_1 : x(t+1) &= f(x(t), u(t)) + e(t), & y(t) &= h(x(t)), \\ \mathbf{M}_2 : x(t+1) &= f(x(t), u(t)) + \tilde{e}(t) + \tilde{u}(t), & y(t) &= h(x(t)), \end{aligned}$$

with variables $x(t), x(t+1), u(t), \tilde{u}(t), e(t), \tilde{e}(t)$ taking values in \mathbb{R}^n , and with dynamics initialised with the same probability distribution at $t = 0$ and driven by white noise sequences $e(t), \tilde{e}(t)$, both with zero mean normal distributions and with variance $\Sigma_e, \Sigma_{\tilde{e}}$, respectively. Notice that if $\Sigma_e - \Sigma_{\tilde{e}}$ is positive definite then $\mathbf{M}_1 \preceq \mathbf{M}_2$. To see this, select the control input pair $(u_2, \tilde{u}_2) \in \mathbb{U}_2$ as $u_2 = u_1$, and \tilde{u}_2 according to the zero-mean normal distribution with variance $\Sigma_e - \Sigma_{\tilde{e}}$, then the associated interface is $\mathcal{U}_v(\cdot|u_1, x_1, x_2) = \delta_{u_1}(du_2)\mathcal{N}(d\tilde{u}_2|0, \Sigma_e - \Sigma_{\tilde{e}})$. For this interface the stochastic dynamics of the two processes are equal, and can be lifted with \mathcal{R}_{diag} .

b. Similar as above, consider two gMDP modelled as Gaussian processes

$$\begin{aligned} \mathbf{M}_1 : x(t+1) &= (A + BK)x(t) + Bu(t) + e(t), & y(t) &= h(x(t)), \\ \mathbf{M}_2 : x(t+1) &= Ax(t) + Bu(t) + e(t), & y(t) &= h(x(t)), \end{aligned}$$

with variables $x(t), x(t+1), e(t)$ taking values in \mathbb{R}^n and $u(t) \in \mathbb{R}^m$, matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $K \in \mathbb{R}^{m \times n}$. Then $\mathbf{M}_1 \preceq \mathbf{M}_2$, since in \mathcal{R}_{diag} for every action u_1 chosen for \mathbf{M}_1 , the choice of interface $u_2 = u_1 + Kx_2$ for \mathbf{M}_2 results in the same transition kernel for the second model.

Remark 3.4 Over $\mathcal{M}_{\mathbb{V}}$, the class of gMDP with a shared output space, the relation \preceq is a preorder, since it is reflexive (see Example 3.3) and transitive (see later Cor. 3.16). Moreover the relation \approx is an equivalence relation as it is also symmetric (see Cor. 3.7).

3.3.d Controller refinement via probabilistic simulation relations

The ideas underlying the controller refinement are first discussed, after which it is shown that the refined controller induces a strategy as per Def. 3.4. Finally the equivalence of properties defined over the controlled gMDPs is shown.

Consider two gMDP $\mathbf{M}_i = (\mathbb{X}_i, \pi_{x(0)_i}, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$ $i = 1, 2$ with $\mathbf{M}_1 \preceq \mathbf{M}_2$. Given the entities \mathcal{U}_v and $\mathbb{W}_{\mathbb{T}}$ associated to $\mathbf{M}_1 \preceq \mathbf{M}_2$, the distribution of the next state x'_2 of \mathbf{M}_2 is given as $\mathbb{T}_2(\cdot | x_2, \mathcal{U}_v(u_1, x_1, x_2))$, and is equivalently defined via the lifted measure as the marginal of $\mathbb{W}_{\mathbb{T}}(\cdot | u_1, x_1, x_2)$ on \mathbb{X}_2 . Therefore, the distribution of the combined next state (x'_1, x'_2) , defined as $\mathbb{W}_{\mathbb{T}}(\cdot | u_1, x_1, x_2)$, can be expressed as

$$\mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_2 | u_1, x_1, x_2) = \mathbb{W}_{\mathbb{T}}(dx'_1 | x'_2, u_1, x_1, x_2) \mathbb{T}_2(dx'_2 | x_2, \mathcal{U}_v(u_1, x_1, x_2)),$$

where $\mathbb{W}_{\mathbb{T}}(dx'_1 | x'_2, u_1, x_1, x_2)$ is referred to as the conditional probability given x'_2 (c.f. [Borkar, 2012, Corollary 3.1.2]).⁴ Similarly, the conditional measure for the initialisation $\mathbb{W}_{\pi_{x(0)}}$ is denoted as

$$\mathbb{W}_{\pi_{x(0)}}(dx_1(0) \times dx_2(0)) = \mathbb{W}_{\pi_{x(0)}}(dx_1(0) | x_2(0)) \pi_{x(0)_2}(dx_2(0)).$$

Now suppose that we have a control strategy for \mathbf{M}_1 , referred to as \mathbf{C}_1 , and we want to construct the refined control strategy \mathbf{C}_2 for \mathbf{M}_2 , which is such that events defined over the output space have equal probability. This refinement procedure follows directly from the interface and the conditional probability distributions, and is described in Algorithm 2. This execution algorithm is separated into the refined control strategy \mathbf{C}_2 and its gMDP \mathbf{M}_2 . \mathbf{C}_2 is composed of \mathbf{C}_1 , the stochastic kernel $\mathbb{W}_{\mathbb{T}}$, and the interface \mathcal{U}_v , and it remembers the previous state of \mathbf{M}_2 (cf. line 8 in Algorithm 2).

Algorithm 2 Refinement of control strategy \mathbf{C}_1 as \mathbf{C}_2

- 1: set $t := 0$
 - 2: draw $x_2(0)$ from $\pi_{x(0)_2}$,
 - 3: draw $x_1(0)$ from $\mathbb{W}_{\pi_{x(0)}}(\cdot | x_2(0))$.
 - 4: **loop**
 - 5: given $x_1(t)$, select $u_1(t)$ according \mathbf{C}_1 ,
 - 6: set $\mu_{2t} := \mathcal{U}_v(u_1(t), x_1(t), x_2(t))$,
 - 7: draw $x_2(t+1)$ from $\mathbb{T}_2(\cdot | x_2(t), \mu_{2t})$,
 - 8: draw $x_1(t+1)$ from $\mathbb{W}_{\mathbb{T}}(\cdot | x_2(t+1), u_1(t), x_1(t), x_2(t))$,
 - 9: set $t := t + 1$.
 - 10: **end loop**
-

Theorem 3.8 (Refined control strategy) *Let gMDP \mathbf{M}_1 and \mathbf{M}_2 be related as $\mathbf{M}_1 \preceq \mathbf{M}_2$, and consider the control strategy $\mathbf{C}_1 = (\mathbb{X}_{\mathbf{C}_1}, x_{\mathbf{C}_1 0}, \mathbb{X}_1, \mathbb{T}_{\mathbf{C}_1}^t, h_{\mathbf{C}_1}^t)$ for \mathbf{M}_1 as given. Then there exists at least one refined control strategy $\mathbf{C}_2 = (\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}_{\mathbf{C}_2}^t, h_{\mathbf{C}_2}^t)$, as defined in Def. 3.4, with*

- state space $\mathbb{X}_{\mathbf{C}_2} := \mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2$, with elements $x_{\mathbf{C}_2} = (x_{\mathbf{C}_1}, x_1, x_2)$;
- initial state $x_{\mathbf{C}_2 0} := (x_{\mathbf{C}_1 0}, 0, 0)$;

⁴ Beyond Borel measurability, this also holds when the kernels are universally measurable, as corresponding universally measurable regular conditional probability measures are obtained [Edalat, 1999].

- input variable $x_2 \in \mathbb{X}_2$, namely the state variable of \mathbf{M}_2 ;
- time-dependent stochastic kernels $\mathbb{T}_{\mathbf{C}_2}^t$, defined as

$$\mathbb{T}_{\mathbf{C}_2}^0(dx_{\mathbf{C}_2}|x_{\mathbf{C}_2 0}, x_2(0)) := \mathbb{T}_{\mathbf{C}_1}^0(dx_{\mathbf{C}_1}|x_{\mathbf{C}_1 0}, x_1) \mathbb{W}_{\pi_{x(0)}}(dx_1|x_2) \delta_{x_2(0)}(dx_2) \text{ and}$$

$$\mathbb{T}_{\mathbf{C}_2}^t(dx'_{\mathbf{C}_2}|x_{\mathbf{C}_2}(t), x_2(t)) := \mathbb{T}_{\mathbf{C}_1}^t(dx'_{\mathbf{C}_1}|x_{\mathbf{C}_1}, x'_1)$$

$$\mathbb{W}_{\mathbb{T}}(dx'_1|x'_2, h_{\mathbf{C}_1}^t(x_{\mathbf{C}_1}), x_2, x_1) \delta_{x_2(t)}(dx'_2) \text{ for } t \in [1, N];$$
- measurable output maps $h_{\mathbf{C}_2}^t(x_{\mathbf{C}_1}, \tilde{x}_1, x_2) := \mathcal{U}_v(h_{\mathbf{C}_1}^t(x_{\mathbf{C}_1}), x_1, x_2)$.

Both the time-dependent stochastic kernels $\mathbb{T}_{\mathbf{C}_2}^t$ and the output maps $h_{\mathbf{C}_2}^t$, for $t \in [0, N]$, are universally measurable, since Borel measurable maps are universally measurable and the latter are closed under composition [Bertsekas and Shreve, 1996, Chapter 7].

Since, by the above construction of \mathbf{C}_2 , the output spaces of the controlled systems $\mathbf{C}_1 \times \mathbf{M}_1$ and $\mathbf{C}_2 \times \mathbf{M}_2$ have equal distribution, it follows that measurable events have equal probability, as stated next and proved in the Appendix.

Theorem 3.9 *If $\mathbf{M}_1 \preceq \mathbf{M}_2$, then for all control strategies \mathbf{C}_1 there exists a control strategy \mathbf{C}_2 such that, for all measurable events $A \in \mathcal{B}(\mathbb{Y}^{N+1})$,*

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{y_1(t)\}_{0:N} \in A) = \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{y_2(t)\}_{0:N} \in A).$$

3.4 New ϵ, δ -approximate (bi-)simulation relations via lifting

3.4.a Motivation and δ -lifting

The requirement on an exact simulation relation between two models is evidently restrictive. Consider the following example, where two Markov processes have a bounded output deviation.

Example 3.5 (Models with a shared noise source) *Consider an output space $\mathbb{Y} := \mathbb{R}^d$, with a metric $\mathbf{d}_{\mathbb{Y}}(x, y) := \|x - y\|$ (the Euclidean norm), and two gMDP expressed as noisy dynamic processes:*

$$\begin{aligned} \mathbf{M}_1 : x_1(t+1) &= f(x_1(t), u_1(t)) + e_1(t), & y_1(t) &= h(x_1(t)), \\ \mathbf{M}_2 : x_2(t+1) &= f(x_2(t), u_2(t)) + e_2(t), & y_2(t) &= h(x_2(t)), \end{aligned}$$

where f and h are both globally Lipschitz, satisfying $\|f(x_1, u) - f(x_2, u)\| \leq L\|x_1 - x_2\|$ for $0 < L < 1$, and in addition $\|h(x_1) - h(x_2)\| \leq H\|x_1 - x_2\|$ for an $0 < H$ valid for all $x_1, x_2 \in \mathbb{R}^n$ and for all u . Suppose that the probability distributions of the random variable e_1 and of e_2 depend on a shared noise source ω , with $\omega \in \Omega$ and distribution \mathbb{P}_{ω} , and are such that $e_1(t) = g_1(\omega(t))$ and $e_2(t) = g_2(\omega(t))$. Assume now that there exists a value $c \in \mathbb{R}$, such that $\mathbb{P}_{\omega}[\|g_1(\omega) - g_2(\omega)\| < c] = 1$. Then for every pair of states $x_1(t)$

and $x_2(t)$ of \mathbf{M}_1 and \mathbf{M}_2 respectively, the difference between state transitions is bounded as $\|x_1(t+1) - x_2(t+1)\| \leq L\|x_1(t) - x_2(t)\| + c$ with probability 1. By induction it can be shown that if $\|x_1(0) - x_2(0)\| \leq \frac{c}{1-L}$, then for all $t \geq 0$, $\|x_1(t) - x_2(t)\| \leq \frac{c}{1-L}$, and $\|y_1(t) - y_2(t)\| \leq \frac{cH}{1-L}$.

Even though the difference in the output of the two models is bounded by the quantity $\frac{cH}{1-L}$ with probability 1, it is impossible to provide an approximation error using either the method in [Julius and Pappas, 2009] (hinging on stochastic stability assumptions), nor using (approximate) relations as in [Desharnais et al., 2008, D’Innocenzo et al., 2012]: with the former approach, for the same input sequence $u(t)$ the output trajectories of \mathbf{M}_1 and \mathbf{M}_2 have bounded difference, but do not converge to each other; with the latter approach, the relation defined via a normed difference cannot satisfy the required notion of transitivity.

As mentioned before and highlighted in the previous Ex. 3.5, we are interested in introducing a new approximate version of the notion of probabilistic simulation relation, which allows for both δ -differences in the stochastic transition kernels, and ϵ -differences in the output trajectories. For the former prerequisite, we relax the requirements on the lifting in Def. 3.5; subsequently, we define the resulting approximate (bi-)simulation relation according to the latter prerequisite on the outputs.

Definition 3.10 (δ -lifting for general state spaces) Let $\mathbb{X}_1, \mathbb{X}_2$ be two sets with associated measurable spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$, $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$, and let $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ be a relation for which $\mathcal{R} \in \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2)$. We denote by $\bar{\mathcal{R}}_\delta \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ the corresponding lifted relation (acting on $\Delta\bar{\mathcal{R}}_\delta\Theta$), if there exists a probability space $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ satisfying

1. for all $X_1 \in \mathcal{B}(\mathbb{X}_1)$: $\mathbb{W}(X_1 \times \mathbb{X}_2) = \Delta(X_1)$;
2. for all $X_2 \in \mathcal{B}(\mathbb{X}_2)$: $\mathbb{W}(\mathbb{X}_1 \times X_2) = \Theta(X_2)$;
3. for the probability space $(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$ it holds that $x_1 \mathcal{R} x_2$ with probability at least $1 - \delta$, or equivalently that $\mathbb{W}(\mathcal{R}) \geq 1 - \delta$.

We leverage Definition 3.10 to introduce a new approximate similarity relation that encompasses both approximation requirements, obtaining the following ϵ, δ -approximate probabilistic simulation.

Definition 3.11 (ϵ, δ -approximate probabilistic simulation) Consider two gMDP $\mathbf{M}_i = (\mathbb{X}_i, \pi_{x(0)_i}, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$, $i = 1, 2$, over a shared metric output space $(\mathbb{Y}, \mathbf{d}_\mathbb{Y})$. \mathbf{M}_1 is ϵ, δ -stochastically simulated by \mathbf{M}_2 if there exists an interface function \mathcal{U}_v and a relation $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$, for which there exists a Borel measurable stochastic kernel $\mathbb{W}_\mathbb{T}(\cdot | u_1, x_1, x_2)$ on $\mathbb{X}_1 \times \mathbb{X}_2$ given $\mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2$, such that:

1. $\forall (x_1, x_2) \in \mathcal{R}$, $\mathbf{d}_\mathbb{Y}(h_1(x_1), h_2(x_2)) \leq \epsilon$;
2. $\forall (x_1, x_2) \in \mathcal{R}$, $\forall u_1 \in \mathbb{U}_1$: $\mathbb{T}_1(\cdot | x_1, u_1) \bar{\mathcal{R}}_\delta \mathbb{T}_2(\cdot | x_2, \mathcal{U}_v(u_1, x_1, x_2))$, with lifted probability measure $\mathbb{W}_\mathbb{T}(\cdot | u_1, x_1, x_2)$;

$$3. \pi_{x(0)_1} \bar{\mathcal{R}}_\delta \pi_{x(0)_2}.$$

The simulation relation is denoted as $\mathbf{M}_1 \preceq_\varepsilon^\delta \mathbf{M}_2$.

Definition 3.12 (ε, δ -approximate probabilistic bisimulation) Under the same conditions as before \mathbf{M}_1 is an ε, δ -probabilistic bisimulation of \mathbf{M}_2 if there exists a relation $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$ such that $\mathbf{M}_1 \preceq_\varepsilon^\delta \mathbf{M}_2$ w.r.t. \mathcal{R} and $\mathbf{M}_2 \preceq_\varepsilon^\delta \mathbf{M}_1$ w.r.t. $\mathcal{R}^{-1} \subset \mathbb{X}_2 \times \mathbb{X}_1$. \mathbf{M}_1 and \mathbf{M}_2 are said to be ε, δ -probabilistically bisimilar, denoted as $\mathbf{M}_1 \approx_\varepsilon^\delta \mathbf{M}_2$.

In this section we have provided similarity relations quantifying the difference between two Markov processes. The end use of the introduced similarity relations is to quantify the probability of events of a gMDP via its abstraction and to refine controllers: this is achieved in the next section.

3.4.b Controller refinement via approximate simulation relations

Consider two gMDP \mathbf{M}_1 and \mathbf{M}_2 , for which \mathbf{M}_1 is the abstraction of the concrete model \mathbf{M}_2 . The following result is an approximate version of Theorem 3.9, and presents the main result of this paper, namely the approximate equivalence of properties defined over the gMDP \mathbf{M}_1 and \mathbf{M}_2 .

Theorem 3.13 If $\mathbf{M}_1 \preceq_\varepsilon^\delta \mathbf{M}_2$, then for all control strategies \mathbf{C}_1 there exists a control strategy \mathbf{C}_2 such that, for all measurable events $A \subset \mathbb{Y}^{N+1}$

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{y_1(t)\}_{0:N} \in A_{-\varepsilon}) - \gamma \leq \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{y_2(t)\}_{0:N} \in A) \leq \mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{y_1(t)\}_{0:N} \in A_\varepsilon) + \gamma,$$

with constant $1 - \gamma := (1 - \delta)^{N+1}$, and with the ε -expansion of A defined as

$$A_\varepsilon := \left\{ \{y_\varepsilon(t)\}_{0:N} \mid \exists \{y(t)\}_{0:N} \in A : \max_{t \in [0, N]} \mathbf{d}_{\mathbb{Y}}(y_\varepsilon(t), y(t)) \leq \varepsilon \right\}$$

and similarly the ε -contraction defined as $A_{-\varepsilon} := \left\{ \{y(t)\}_{0:N} \mid \{ \{y(t)\}_{0:N} \}_\varepsilon \subset A \right\}$ where $\{ \{y(t)\}_{0:N} \}_\varepsilon$ is the point-wise ε -expansion of $\{y(t)\}_{0:N}$.

While the details of the proof can be found in the Appendix, its key aspect is the existence of a refined control strategy \mathbf{C}_2 , which we detail next. Given a control strategy \mathbf{C}_1 over the time horizon $t \in \{0, \dots, N\}$, there is a control strategy \mathbf{C}_2 that refines \mathbf{C}_1 over \mathbf{M}_2 . The control strategy is conceptually given in Algorithm 3. Whilst the state (x_1, x_2) of \mathbf{C}_2 is in \mathcal{R} , the control refinement from \mathbf{C}_1 follows in the same way (cf. Alg.3 line 4-9) as for the exact case of Section 3.3.d. Hence, similar to the control refinement for exact probabilistic simulations, the *basic ingredients* of \mathbf{C}_2 are the states x_1 and x_2 , whose stochastic transition to the pair (x'_1, x'_2) is governed firstly by a point distribution $\delta_{x_2(t)}(dx'_2)$ based on the measured state $x_2(t)$ of \mathbf{M}_2 ; and, subsequently, by the lifted probability measure $\mathbb{W}_{\mathbb{T}}(dx'_1 \mid x'_2, u_1, x_2, x_1)$, conditioned on x'_2 .

On the other hand, whenever the state (x_1, x_2) leaves \mathcal{R} the control chosen by strategy \mathbf{C}_1 cannot be refined to \mathbf{M}_2 : instead, an alternative control strategy \mathbf{C}_{rec}

has to be used to control the residual trajectory of M_2 . The choice is of no importance to the result in Theorem 3.13. This stage of the execution (cf. Alg. 3 line 11-15) referred to as *recovery* makes the choice of the overall control strategy C_2 non-unique. In practice we will only synthesise the control strategy over a finite-time.

Algorithm 3 Refinement of C_1 as C_2

```

1: set  $t := 0$  {Start}
2: draw  $x_2(0)$  from  $\pi_{x(0)_2}$ 
3: draw  $x_1(0)$  from  $\mathbb{W}_{\pi_{x(0)}}(\cdot \mid x_2(0))$ 
4: while  $(x_1(t), x_2(t)) \in \mathcal{R}$  do {Refine}
5:   given  $x_1(t)$ , select  $u_1(t)$  from  $C_1$ ,
6:   set input  $\mu_{2t} := \mathcal{U}_v(u_1(t), x_1(t), x_2(t))$ ,
7:   draw  $x_2(t+1)$  from  $\mathbb{T}_2(\cdot \mid x_2(t), \mu_{2t})$ ,
8:   draw  $x_1(t+1)$  from  $\mathbb{W}_{\mathbb{T}}(\cdot \mid x_2(t+1), u_1(t), x_1(t), x_2(t))$ ,
9:   set  $t := t + 1$ 
10: end while
11: loop {Recover}
12:   given  $x_2(t)$ , select  $\mu_t$  (from  $C_{rec}$ ),
13:   draw  $x_2(t+1)$  from  $\mathbb{T}_2(\cdot \mid x_2(t), \mu_t)$ ,
14:   set  $t := t + 1$ 
15: end loop

```

By splitting the execution in Algorithm 3 into a control strategy and a gMDP M_2 , we can again obtain the refined control strategy.

Theorem 3.14 (Refined control strategy) *Let gMDP M_1 and M_2 , with $M_1 \preceq_\epsilon^\delta M_2$, and control strategy $C_1 = (\mathbb{X}_{C_1}, x_{C_1 0}, \mathbb{X}_1, \mathbb{T}_{C_1}^t, h_{C_1}^t)$ for M_1 be given. Then for any given recovery control strategy C_{rec} , a refined control strategy, denoted as*

$$C_2 = (\mathbb{X}_{C_2}, x_{C_2 0}, \mathbb{X}_2, \mathbb{T}_{C_2}^t, h_{C_2}^t),$$

can be obtained as an inhomogenous Markov process with two discrete modes of operation, {refinement} and {recovery}, based on Algorithm 3.

The details of the tuple $(\mathbb{X}_{C_2}, x_{C_2 0}, \mathbb{X}_2, \mathbb{T}_{C_2}^t, h_{C_2}^t)$ are given in the Appendix, together with the proof of the theorem. They follow from Algorithm 3, in a similar way as Theorem 3.8 follows from Algorithm 2.

3.4.c Examples and properties

Example 3.6 (Models with a shared noise source – continued from above)

Based on the relation $\mathcal{R} := \{(x_1, x_2) : \|x_1 - x_2\| \leq \frac{c}{1-L}\}$ it can be shown that $M_1 \approx_\epsilon^0 M_2$ with $\epsilon = \frac{Hc}{1-L}$, since, firstly, it holds that $\mathbf{d}_Y(h(x_1) - h(x_2)) \leq \epsilon$ for all $(x_1, x_2) \in \mathcal{R}$ with $\mathbf{d}_Y = \|\cdot\|$. Additionally, for all $(x_1, x_2) \in \mathcal{R}$ and for any input u_1

the selection $u_2 = u_1$ is such that $\mathbb{T}_1(\cdot|x_1, u_1)\bar{\mathcal{R}}_0\mathbb{T}_2(\cdot|x_2, u_1)$, note that $\bar{\mathcal{R}}_0$ is equal to $\bar{\mathcal{R}}$ (the lifted relation from \mathcal{R}). The lifted stochastic kernel is $\mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_2|u_1, x_1, x_2) := \int_{\omega} \delta_{f(x_1, u_1) + g_1(\omega)}(dx'_1) \delta_{f(x_2, u) + g_2(\omega)}(dx'_2) \mathbb{P}_{\omega}(d\omega)$, this stochastic kernel is Borel measurable if $f(x_1, u_1) + g_1(\omega)$ and $f(x_2, u) + g_2(\omega)$ are assumed Borel measurable maps. Note that the employed identity interface is also Borel measurable.

Example 3.7 (Relationship to model with truncated noise) Consider the stochastic dynamical process $\mathbf{M}_1 : x(t+1) = f(x(t), u(t)) + e(t)$ with output mapping $y(t) = h(x(t))$, operating over the Euclidean state space \mathbb{R}^n , and driven by a white noise sequence $e(t) \in \mathbb{R}^n$ with distribution \mathbb{P}_e . The output space $y \in \mathbb{Y} \subseteq \mathbb{R}^d$ is endowed with the Euclidean norm $\mathbf{d}_{\mathbb{Y}} = \|\cdot\|$. Select a domain $D \subset \mathbb{R}^n$ so that, at any given time instant t , $e(t) \in D$ with probability $1 - \delta$. Then define a truncated white noise sequence $\tilde{e}(t)$, with distribution $\mathbb{P}_e(\cdot|D)$. The resulting model \mathbf{M}_2 driven by $\tilde{e}(t)$ is $\mathbf{M}_2 : x(t+1) = f(x(t), u(t)) + \tilde{e}(t)$, with the same output mapping $y(t) = h(x(t))$. We show that \mathbf{M}_2 is a $0, \delta$ -approximate probabilistic bisimulation of \mathbf{M}_1 , i.e. $\mathbf{M}_1 \approx_0^{\delta} \mathbf{M}_2$. Select $\mathcal{R} := \{(x_1, x_2) \text{ for } x_1, x_2 \in \mathbb{R}^n | x_1 = x_2\}$, and choose as interface the identity one, i.e., $\mathcal{U}_v(u_1, x_1, x_2) = u_1$. A viable lifting measure is

$$\begin{aligned} \mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_2|u_1, x_1, x_2) &:= \int_{e \in D} \delta_{x'_1}(dx'_1) \delta_{t_1(e)}(dx'_2) \mathbb{P}_e(de) \\ &\quad + \int_{e \in \mathbb{R}^n \setminus D} \delta_{t_1(e)}(dx'_1) \mathbb{P}_e(de) \int_{\tilde{e}} \delta_{t_2(\tilde{e})}(dx'_2) \mathbb{P}_e(d\tilde{e}|D) \end{aligned} \quad (3.1)$$

with $t_1(e) = f(x_1, u_1) + e$ and $t_2(\tilde{e}) = f(x_2, u_1) + \tilde{e}$.

Example 3.8 (Relationship between noiseless and truncated-noise models)

Consider the model with truncated noise \mathbf{M}_2 as defined in Ex. 3.7. In what sense is \mathbf{M}_2 approximated by its noiseless version \mathbf{M}_3 , namely $\mathbf{M}_3 : x(t+1) = f(x(t), u(t))$, $y(t) = h(x(t))$? Under requirements on the Lipschitz continuity $\|f(x_1, u) - f(x_2, u)\| \leq L\|x_1 - x_2\|$, $0 < L < 1$, $\|h(x_1) - h(x_2)\| \leq H\|x_1 - x_2\|$, and on the boundedness of D and of $c = \max_{d \in D} \|d\|$, Ex. 3.5 can be leveraged by concluding that $\mathbf{M}_2 \approx_{\varepsilon}^0 \mathbf{M}_3$, with $\varepsilon = \frac{Hc}{1-L}$.⁵

In Examples 3.7 and 3.8 we have that \mathbf{M}_1 is approximated by \mathbf{M}_2 , which is subsequently approximated by \mathbf{M}_3 . The following theorem and corollary attain a quantitative answer on the question whether \mathbf{M}_1 is approximated by \mathbf{M}_3 .

Theorem 3.15 (Transitivity of $\preceq_{\varepsilon}^{\delta}$) Consider three gMDP \mathbf{M}_i , $i = 1, 2, 3$, defined by tuples $(\mathbb{X}_i, \pi_{x(0)_i}, \mathbb{T}_i, \mathbb{U}_i, h_i, \mathbb{Y})$. If

- \mathbf{M}_1 is ε_a, δ_a -stochastically simulated by \mathbf{M}_2 , and
- \mathbf{M}_2 is ε_b, δ_b -stochastically simulated by \mathbf{M}_3 ,

then \mathbf{M}_1 is $(\varepsilon_a + \varepsilon_b), (\delta_a + \delta_b)$ -stochastically simulated by \mathbf{M}_3 . Equivalently, if

$$\mathbf{M}_1 \preceq_{\varepsilon_a}^{\delta_a} \mathbf{M}_2 \text{ and } \mathbf{M}_2 \preceq_{\varepsilon_b}^{\delta_b} \mathbf{M}_3, \text{ then } \mathbf{M}_1 \preceq_{\varepsilon_a + \varepsilon_b}^{\delta_a + \delta_b} \mathbf{M}_3.$$

⁵ Alternatively, if \mathbf{M}_2 with non-deterministic input $\tilde{e} \in D$ is an ε_a -alternating bisimulation [Tabuada, 2009] of \mathbf{M}_3 then $\mathbf{M}_2 \approx_{\varepsilon_a}^0 \mathbf{M}_3$.

Next, as a corollary of this theorem, we derive properties of the notion of approximate bisimulation, and discuss the transitivity of the (exact) notions of simulation and of bisimulation relation. The latter implies that the simulation relation (cf. Def.3.6) is a preorder, and that the bisimulation relation (cf. Def.3.7) is an equivalence relation over the category of gMDP $\mathcal{M}_{\mathbb{Y}}$.

Corollary 3.16 (Transitivity properties) *Following Theorem 3.15,*

- if $\mathbf{M}_1 \approx_{\epsilon_a}^{\delta_a} \mathbf{M}_2$ and $\mathbf{M}_2 \approx_{\epsilon_b}^{\delta_b} \mathbf{M}_3$, then $\mathbf{M}_1 \approx_{\epsilon_a + \epsilon_b}^{\delta_a + \delta_b} \mathbf{M}_3$, and
- if $\mathbf{M}_1 \preceq \mathbf{M}_2$ and $\mathbf{M}_2 \preceq \mathbf{M}_3$, then $\mathbf{M}_1 \preceq \mathbf{M}_3$, and
- if $\mathbf{M}_1 \approx \mathbf{M}_2$ and $\mathbf{M}_2 \approx \mathbf{M}_3$, then $\mathbf{M}_1 \approx \mathbf{M}_3$.

Here notice that for $\mathcal{R}_{13} := \{(x_1, x_3) | \exists x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}\}$ we show that if $\Delta_1 \bar{\mathcal{R}}_{12\delta_a} \Delta_2$ and $\Delta_2 \bar{\mathcal{R}}_{23\delta_b} \Delta_3$, then $\Delta_1 \bar{\mathcal{R}}_{13(\delta_a + \delta_b)} \Delta_3$, where the used lifting measure $\mathbb{W}_{\mathbb{T}}$ is a function of the respective liftings $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$, i.e. for all $x_1, x_3 \in \mathcal{R}_{13} \exists x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}$, $\mathbb{W}_{\mathbb{T}}$ is given as

$$\begin{aligned} \mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_3 | u_1, x_1, x_2) = \\ \int_{\mathbb{X}_2} \mathbb{W}_{23}(dx'_3 | x'_2, \mathcal{U}_{v12}(u_1, x_1, x_2), x_2, x_3) \mathbb{W}_{12}(dx'_1 \times dx'_2 | u_1, x_1, x_2). \end{aligned}$$

Furthermore, the interface \mathcal{U}_{v13} is the composition of \mathcal{U}_{v12} and \mathcal{U}_{v23} . The proof of Theorem 3.15 and Corollary 3.16 can be found in the Appendix.

Example 3.9 (Combination of Examples 3.7 and 3.8 via Corollary 3.16)

For the models in Examples 3.7 and 3.8 we can conclude that $\mathbf{M}_1 \approx_{\epsilon}^{\delta} \mathbf{M}_3$. This means that a stochastic system as in \mathbf{M}_1 in Ex. 3.7 can be approximated via its deterministic counterpart, and that the approximation error can be expressed via the probability (i.e. amount of truncation cf. Ex. 3.7) and the output error (i.e. Ex. 3.8). This allows for explicit trading off between output deviation and deviation in probability.

3.5 Case studies

3.5.a Introduction: energy management in smart buildings

We are interested in developing advanced solutions for the energy management of smart buildings. In this work we first describe a simple example with a three-dimensional model of the thermal dynamics in an office building: we consider a simple building that is divided in two connected zones, each with a radiator affecting the heat exchange in that zone by controlling the water temperature in a boiler. With this case study we aim at elucidating the theory of the previous sections. In the third subsection we work with a more realistic model of an office building: this 5-dimensional model shows how the given approximate similarity relations can be used for the design of controllers that verifiably satisfy properties expressed as quantitative specifications.

3.5.b First case study

A model of the temperature dynamics in an office building with two zones to heat [Haesaert et al., 2015a, Holub and Macek, 2013] assumes that the temperature fluctuations in the two zones, as well as the ambient temperature dynamics, can be modelled as a Gaussian process

$$\mathbf{M} : x(t+1) = Ax(t) + Bu(t) + Fe(t), \quad y(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} x(t), \quad (3.2)$$

with stable dynamics characterised by matrices

$$A = \begin{bmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \\ 0 & 0 & 0.9900 \end{bmatrix}, \quad B = \begin{bmatrix} 0.0650 & 0 \\ 0 & 0.60 \\ 0 & 0 \end{bmatrix}, \quad F = \begin{bmatrix} 0.05 & -0.02 & 0 \\ -0.02 & 0.05 & 0 \\ 0 & 0 & 0.1 \end{bmatrix},$$

where $x_{1,2}(t)$ are the temperatures in zone 1 and 2, respectively; $x_3(t)$ is the deviation of the ambient temperature from its mean; and $u(t) \in \mathbb{R}^2$ is the control input. The disturbance $e(t)$ is a white noise sequence with standard Gaussian distributions, for all $t \in \mathbb{R}^+$. The state variables are initiated as $x(0) = [16 \ 14 \ -5]^T$. This stochastic process can be written as a gMDP, as detailed in Example 3.1. As the model abstraction, we select the controllable and deterministic dynamics of the mean of the state variables, and consequently omit the ambient temperature and the additive noise term:

$$\tilde{\mathbf{M}} : \begin{cases} \tilde{x}(t+1) &= \tilde{A}\tilde{x}(t) + \tilde{B}\tilde{u}(t) \in \mathbb{R}^2, \text{ with } \tilde{A} := \begin{bmatrix} 0.8725 & 0.0625 \\ 0.0625 & 0.8775 \end{bmatrix}, \\ \tilde{y}(t) &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tilde{x}(t), \quad \tilde{B} := \begin{bmatrix} 0.0650 & 0 \\ 0 & 0.60 \end{bmatrix}. \end{cases} \quad (3.3)$$

We then obtain that, as intuitive, $\tilde{\mathbf{M}} \preceq_{\varepsilon}^{\delta} \mathbf{M}$. In order to compute specific values of ε and δ , we select the relation $\mathcal{R} := \{(\tilde{x}, x) \in \mathbb{R}^2 \times \mathbb{R}^3 \mid \sqrt{(\tilde{x}_1 - x_1)^2 + (\tilde{x}_2 - x_2)^2} \leq \varepsilon\}$ and the interface function $\mathcal{U}_v(\tilde{u}, \tilde{x}, x) = \tilde{u} + \tilde{B}^{-1}(\tilde{A}\tilde{x} - \bar{A}x)$, with

$$\bar{A} = \begin{bmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \end{bmatrix}.$$

The structure of the interface is arbitrary: in the specific instance the interface is selected to optimally correct the difference in room temperatures at the next time step.

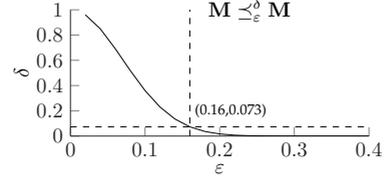
A stochastic kernel $\mathbb{W}_{\mathbb{T}}$ for the lifting is

$$\mathbb{W}_{\mathbb{T}}(d\tilde{x}' \times dx' \mid \tilde{u}, \tilde{x}, x) = \int_e \delta_{\tilde{f}}(d\tilde{x}') \delta_{f(e)}(dx') \mathcal{N}(de \mid 0, I),$$

with $\tilde{f} = \tilde{A}\tilde{x} + \tilde{B}\tilde{u}$ and $f(e) = Ax + BU_v(\tilde{u}, \tilde{x}, x) + Fe$. The lower bound on $\mathbb{W}_{\mathbb{T}}(\mathcal{R} \mid \tilde{u}, \tilde{x}, x) \leq 1 - \delta$ has been computed and traded off against the output deviation, as in Figure 3.2.

We are interested in the goal, expressed for the model \mathbf{M} , of increasing the likelihood of trajectories reaching the target set $T = [20.5, 21]^2$ and staying there thereafter. For the abstract model we have developed a strategy, as in [Haesaert et al., 2015a], satisfying by construction the property expressed in LTL-like nota-

Figure 3.2: Trade-off between the output error ε and the probability error δ for the δ, ε -approximate probabilistic simulation $\tilde{M} \stackrel{\delta, \varepsilon}{\sim} M$. We have selected the pair $(\varepsilon, \delta) = (0.16, 0.073)$ as an ideal trade-off.



tion with the formula $\varphi = \diamond \square T$ and shrunken to $\varphi_{-\varepsilon}$ (as per Theorem 3.13). This strategy is synthesised as a correct-by-construction controller using PESSOA [Mazo Jr et al., 2010], where the discrete-time dynamics in (3.3) are further discretised over state and action spaces: we have selected a state quantisation of 0.05 over the range $[15, 25]^2$ for the two state variables, and an input quantisation of 0.05 over the set $[10, 30]^2$. It can be observed that the controller regulates the abstract model \tilde{M} to eventually remain within the target region, as shown in Figure 3.3. We now want to verify that indeed, when refined to the concrete stochastic model, this strategy implies the reaching and staying in the safe set up to some probabilistic error. The refined strategy is obtained from this control strategy as discussed in Section 3.4.b, and recovers from exits out of the relation \mathcal{R} by resetting the abstract states in the relation.

In a simulation study reported in Figure 3.3, we have executed the refined control strategy over a time horizon of 200 steps. Observe that for the execution displayed in the top/left plot the behaviour of the controlled concrete model M remains close to that of \tilde{M} . Only at 4 incidents (circled) does the output error exceed the level $\varepsilon = 0.16$. This reflects our expectations, since at any point in time the probability that the output error exceeds the level $\varepsilon = 0.16$ over the following X time steps is provably less than $1 - (1 - \delta)^X \approx X\delta = 0.073X$, as per Theorem 3.13, which leads to an upper bound of 15 occurrences. Within this case study, whenever the state of the abstract and concrete model leave the relation \mathcal{R} , then the recovery strategy consists of resetting the state of the abstract model and continuing with the refined control strategy. Thanks to the use of the ε -contraction $\varphi_{-\varepsilon}$ of the concrete specification φ , model M will still abide by φ with a high confidence.

3.5.c Second case study

We consider a realistic model for an office building, with the dynamics obtained from [Bacher and Madsen, 2011]. With a time sampling of 5 minutes, the following model describes stochastic temperature fluctuations around a known mean value:

$$\begin{aligned}
 \mathbf{M}_{\text{office}} : \quad & \begin{cases} x_b(t+1) &= \Xi x_b(t) + \Gamma q(t) + B_p w_p(t) + B_s \Phi_s(t) + B_a T_a(t) \\ y(t) &= [0 \ 1 \ 0 \ 0] x_b(t), \end{cases} \\
 \Xi &= \begin{bmatrix} 0.4487 & 0.216 & 0.2164 & 0.1186 \\ 0.216 & 0.1778 & 0.3719 & 0.2334 \\ 0.09639 & 0.1657 & 0.6569 & 0.08082 \\ 0.005234 & 0.0103 & 0.008007 & 0.9708 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 2.65\text{e-}5 \\ 7.45\text{e-}5 \\ 2.06\text{e-}4 \\ 0.07\text{e-}5 \end{bmatrix}, \quad B_p = \begin{bmatrix} 1.0939\text{e-}4 \\ 2.16\text{e-}4 \\ 7.45\text{e-}5 \\ 3.92\text{e-}6 \end{bmatrix} \\
 B_s &= \begin{bmatrix} 6.60\text{e-}4 \\ 1.31\text{e-}3 \\ 4.49\text{e-}4 \\ 2.36\text{e-}5 \end{bmatrix}, \quad B_a = \begin{bmatrix} 2.96\text{e-}4 \\ 8.79\text{e-}4 \\ 1.93\text{e-}4 \\ 5.67\text{e-}3 \end{bmatrix}.
 \end{aligned}$$

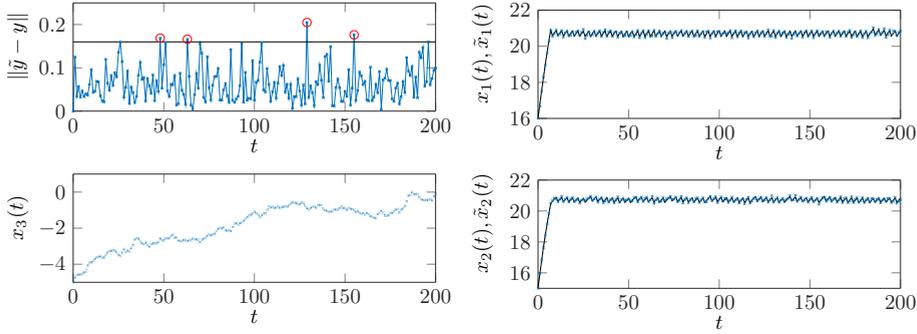


Figure 3.3: Refined control for deterministic model applied to M . The figure (top left) evaluates the accuracy of the approximation, and gives with red circles the instances in which the relation is left. The plot (bottom left) shows the ambient temperature. The plots on the right display the temperature inside the two rooms. The small blue crosses give the actual temperature in the rooms (x_1, x_2) whereas the deterministic simulation of $(\tilde{x}_1, \tilde{x}_2)$ is drawn in black and mostly covered by the crosses.

The output $y(t)$ models the temperature deviation of the internal air. The 4-dimensional state of the model, obtained from a frequency-based identification procedure, represents the fluctuation of internal temperatures in the building, including the building envelope and the interior [Bacher and Madsen, 2011, TiTeThTs model], where the influence of mean values dynamics have been eliminated from the model. The objective of this model is to capture the influence of stochastic effects acting upon the system and control them via the heater with input $q(t)$. The model represents the stochastic disturbances on the building temperature. We foresee three major sources of stochastic disturbance to the system, as explained next.

The first, $w_p(t)$ is the randomness of the heat generated by people in the building. An average person generates 100 Watt [W] under normal circumstances. We presume that the occupancy of the office adds a random element to this average number, which we capture as an independently and identically distributed random signal with Gaussian distribution and a standard deviation equal to 20 % per person: when there are $n_p := 10$ people in the office this standard deviation becomes $\sqrt{n_p} \times 20$ [W].

The second source of stochastic disturbance is the ambient temperature, for which we model the stochastic deviation $T_a(t)$ from accurate weather forecasts. As this deviation is correlated over time, this is modelled as a first-order coloured noise, with a time constant of 20 minutes. The choice of the time constant gives a measure of correlation in time [Therrien, 1992], so we use it to choose the time over which there is a significant correlation between successive values of $T_a(t)$. Additionally, we choose it such that the stationary variance is equal to 1, i.e., $\mathbf{E}[T_a(t)^2] = 1$. The resulting weather model is a first-order (1-dimensional) model $T_a(t+1) = 0.7788T_a + 0.6273w_w(t)$, which is driven by a white noise source with standard

Gaussian distribution, namely $w_w(t) \sim \mathcal{N}(0, I)$.

The third and final source of disturbance $\Phi_s(t)$ is the energy flow from solar radiation. Though measurable, this disturbance cannot exactly be predicted and has a high impact on the temperature inside the office. The impact depends on the effective window area of the building, which has been estimated as 6.03 [m²] in [Bacher and Madsen, 2011]. Based on the measured solar radiation in [Bacher and Madsen, 2011], we model this disturbance as a white noise source with standard deviation of 0.1 [kW/m²].

Including the weather model for T_a , which requires encompassing the noise signal $w_w(t)$, leads to the following 5-dimensional model for the temperature fluctuations in the office building:

$$\mathbf{M} = (A, B, B_w, C) : \begin{cases} x(t+1) &= Ax(t) + B_w w(t) + Bu(t) \\ y(t) &= [0 \ 1 \ 0 \ 0 \ 0] x(t) \end{cases}$$

$$A = \begin{bmatrix} 0.4487 & 0.216 & 0.2164 & 0.1186 & 2.96\text{e-}4 \\ 0.216 & 0.1778 & 0.3719 & 0.2334 & 8.789\text{e-}4 \\ 0.09639 & 0.1657 & 0.6569 & 0.08082 & 1.928\text{e-}4 \\ 0.005234 & 0.0103 & 8.007\text{e-}3 & 0.9708 & 0.005667 \\ 0 & 0 & 0 & 0 & 0.7788 \end{bmatrix},$$

$$B = \begin{bmatrix} 0.1326 \\ 0.3725 \\ 1.029 \\ 4.309\text{e-}3 \\ 0 \end{bmatrix}, \quad B_w = \begin{bmatrix} 0.006918 & 0.06596 & 0 \\ 0.01372 & 0.1308 & 0 \\ 0.004712 & 0.04492 & 0 \\ 2.485\text{e-}4 & 0.002369 & 0 \\ 0 & 0 & 0.6273 \end{bmatrix}.$$

In order to avoid numerical ill-conditioning issues, both the heat input $q(t)$ (expressed in kW) and the corresponding matrix Γ have been replaced by scaled versions, namely the input signal $u(t)$ and the input matrix B . At full throttle the heating input $q(t) = 5$ [kW] corresponds to the scaled input $u(t) = 1$. Similarly the three noise sources discussed above have been normalised together with the respective system matrices, so that $w(t)$ is the new driving noise, as a white-noise sequence with a standard Gaussian distribution, encompassing the unpredicted heat caused by people, solar radiation, and weather fluctuations.

We are interested in controlling the obtained stochastic system \mathbf{M} to verify a quantitative property over its output signal, which is the inner air temperature. More precisely, we want to maximise the probability that the deviation of the inner air temperature stays within a 0.5 degrees difference from the nominal temperature, over an horizon of 30 minutes. This property can be encoded as a PCTL specification for the discrete time model as follows: $\mathbb{P}_{\geq p} (\square^6 [|y| < 0.5])$, where p is one of the parameters to be optimised over.

In order to solve this type of probabilistic safety problems we would normally employ formal abstractions, as implemented in the software tool FAUST² [Esmaeil Zadeh Soudjani et al., 2015]. However, a straightforward use of the tool on the non-autonomous 5-dimensional model does not yield tight guarantees.⁶ Hence, we first obtain several reduced-order models; then, over the input range of interest, we quantify the corresponding ϵ, δ -approximate probabilistic bisimulation relations; finally, we design a controller over the obtained formal abstractions with FAUST², and refine it to the original 5-dimensional model of the office building.

⁶For this 5 dimensional model obtainable guarantees δ exceed 1.

In the refinement step we tune the trade-off between the conservativeness with respect to heating inputs and the accuracy of the approximation.

Model abstraction

We use model order reduction via balanced truncations, as implemented in `Matlab`, to obtain lower-order approximations preserving the dynamics of interest. We seek to obtain either first- or second-order models, from two types of concrete dynamics: firstly, the native dynamics of model $M = (A, B, B_w, C)$, and secondly the dynamics of model $M' = (A + BF, B, B_w, C)$. In the latter case, the state-feedback gain F is chosen⁷ so that it reduces the importance of the controllable modes of the system: $F = [0.48456 \quad 0.39865 \quad 0.85352 \quad 0.56387 \quad 0.0024252]$.

As a result, we obtain four reduced-order models $M_i = (A_i, B_i, B_{wi}, C_i)$ ($i = 1, 2, 3, 4$) of M via balanced truncation⁸:

$$M_i : \begin{cases} x_s(t+1) &= A_i x_s(t) + B_{wi} w(t) + B_i u_s(t) \\ y_s(t) &= C_i x_s(t), \end{cases} \quad (3.4)$$

where the resulting matrices are given in the appendix.

Models M_1 and M_3 are obtained based on $M = (A, B, B_w, C)$, whereas M_2 and M_4 are based on the dynamics of $M' = (A + BF, B, B_w, C)$. As expected the quality of the reduced models depends on the choice of M' or M : in the former case, the part of the dynamics that we cannot compensate with a control is approximated best, whereas for M the most prominent dynamics are approximated best, notwithstanding how well they can be controlled.

Approximate probabilistic simulation relations

The reduced models M_1, M_2, M_3, M_4 are approximations of M and it is expected that, even when using an interface function, the error between these reduced models and M will increase with the input u_s . Therefore we quantify the performance of M_i for $i = 1, 2, 3, 4$ only over a bounded input set $\mathbb{U}_s := \{u_s \in \mathbb{R} \mid u_s^2 \leq c_1\}$. To choose a relevant c_1 suppose we would take constant c_1 of $0.25 = 0.5^2$, then this would be equal to an allowed deviation of 50 percent of the maximal input for the nominal heat input, which is 5[kW] for the original system. As we only want to correct the heating with respect to stochastic fluctuations we take the more realistic value for c_1 of $0.2^2 = 0.04$.

Let us now compute the parameters pair (ϵ, δ) establishing the relationship $M_i \preceq_{\epsilon}^{\delta} M$ between reduced-order and concrete models. Similarly to the work [Girard and Pappas, 2009] on hierarchical control based on model reduction we consider a putative relation between the two state spaces as

$$\mathcal{R} := \{(x, x_s) \mid (x - Px_s)^T M (x - Px_s) \leq \epsilon^2\},$$

⁷The gain term is obtained with the `dare(A, B, C^T C, 0.02)` command in `Matlab`.

⁸This results from the application of the `balred` function in `Matlab`.

with properly-sized matrices M and P , satisfying the Sylvester equation $PA_i = AP + BQ$, for a choice of Q , and $C_i = CP$, and so that $M - C^T C$ is positive semi-definite, namely $M - C^T C \succeq 0$. Introduce the interface $\mathcal{U}_v : \mathbb{U}_s \times \mathbb{X}_s \times \mathbb{X} \rightarrow \mathbb{U}$ as

$$u = Ru_s + Qx_s + K(x - Px_s),$$

and notice that \mathcal{U}_v is a function of both P and Q above, alongside the additional design variables R and K (to be further discussed shortly). The interface function is chosen to reduce the differences in the observed stochastic behaviours of the two systems. It refines any choice of u_s to a control input u , as such it implements any control strategy for \mathbf{M}_i to the original model \mathbf{M} . In this case study we have considered a concrete model that is controllable, linear, time-invariant, and driven by an additive stochastic noise. The chosen interface \mathcal{U}_v , with design variables Q , K , and R , fully parameterises the set of possible interfaces that refine controls synthesised over a reduced model that is deterministic, linear, and time-invariant, as suggested in [Girard and Pappas, 2009].

Let us next focus on the characterisation of the relation $\mathbf{M}_i \preceq_{\epsilon}^{\delta} \mathbf{M}$. Condition 1 in Definition 3.11, namely $\forall(x, x_s) \in \mathcal{R} : d_{\mathbb{Y}}(y(t), y_s(t)) \leq \epsilon$, holds since $\|y - y_s\|^2 = \|Cx - CPx_s\|^2$ and $(x - Px_s)^T C^T C (x - Px_s) \leq (x - Px_s)^T M (x - Px_s)$, and the latter is bounded by ϵ^2 for $(x, x_s) \in \mathcal{R}$.

For condition 2, i.e., $\forall(x, x_s)$ and $\forall u_s \in \mathbb{U}_s : \mathbb{T}_s(\cdot | x_s, u_s) \bar{\mathcal{R}}_{\delta} \mathbb{T}(\cdot | x, \mathcal{U}_v(u_s, x_s, x))$, we construct a lifted probability measure $\mathbb{W}_{\mathbb{T}}(\cdot | u_s, x_s, x)$ based on the shared input noise $w(t)$. From this lifting measure, the original transition kernels can easily be recovered by marginalising over \mathbb{X}_s and over \mathbb{X} , respectively, as $\mathbb{T}(\cdot | x, u) = \mathcal{N}(\cdot | Ax + B\mathcal{U}_v(u_s, x_s, x), B_w B_w^T)$, and $\mathbb{T}_s(\cdot | x_s, u_s) = \mathcal{N}(\cdot | A_s x_s + B_s u_s, B_{wi} B_{wi}^T)$. The last condition requires that, with probability at least $1 - \delta$, the pair $(x', x'_s) \in \mathcal{R}$ is distributed as $(x', x'_s) \sim \mathbb{W}_{\mathbb{T}}(\cdot | u_s, x_s, x)$. This condition can be encoded as: $\forall w^T w \leq c_w, \forall(x, x_s) \in \mathcal{R}, \forall u_s \in \mathbb{U}_s$ it holds that: $(x' - x'_s) \in \mathcal{R}$. Note that the latter can be written as $(x' - Px'_s)^T M (x' - Px'_s) \leq \epsilon^2$, where

$$x' - Px'_s = (A + BK)(x - Px_s) + (B_w - PB_{wi})w + (BR - PB_s)u_s. \quad (3.5)$$

The conditions above can be expressed as a single matrix inequality via the S -procedure [Boyd and Vandenberghe, 2004]. We know that $w \sim \mathcal{N}(0, I)$, $w^T w$ has a Chi-square distribution with 2 degrees of freedom. Thus for a required level of $1 - \delta$, we select c_w as $c_w = \chi_2^{-1}(1 - \delta)$ and solve the resulting constraints with respect to ϵ for given values of K, P, Q and R , for each of the reduced models \mathbf{M}_i using CVX [Grant and Boyd, 2014]. Note that χ_2^{-1} is the chi-square inverse cumulative distribution function with 2 degrees of freedom. The gains K and R are selected together with M by alternately optimising their choice. The chosen P and Q follow from the Sylvester equation, for which additional freedom is used to minimise the influence of w and u_s in (3.5).

Table 3.1 provides a number of ϵ, δ values, derived from the approximate probabilistic simulation relation, for each of the models \mathbf{M}_i . Notice that for increasing values of δ, ϵ decreases to a positive lower bound: this lower bound is a function of the size of the set \mathbb{U}_s . Based on these outcomes, we have decided to proceed with \mathbf{M}_2 .

Table 3.1: ε, δ -simulation relation trade-off for the reduced-order models. The table gives for each model and δ the computed ε .

δ	1	$10^{-\frac{1}{3}}$	$10^{-\frac{2}{3}}$	10^{-1}	$10^{-\frac{4}{3}}$	$10^{-\frac{5}{3}}$	10^{-2}	$10^{-\frac{7}{3}}$	$10^{-\frac{8}{3}}$	10^{-3}
M_1	0.1233	0.4803	0.6247	0.7347	0.827	0.9082	0.9816	1.049	1.112	1.171
M_2	0.01445	0.1037	0.132	0.1534	0.1714	0.1871	<u>0.2014</u>	0.2145	0.2267	0.2381
M_3	0.05206	0.7612	0.997	1.175	1.325	1.456	1.575	1.684	1.785	1.881
M_4	0.1839	0.3029	0.3358	0.3604	0.3809	0.3988	0.415	0.4298	0.4435	0.4564

Control synthesis over abstract model M_2 : use of FAUST²

For a given choice of ε, δ we follow Theorem 3.13 and modify the given PCTL property $\psi := \mathbb{P}_{\geq p}(\Box^6[|y| < 0.5])$ to obtain $\psi_{\varepsilon, \delta} := \mathbb{P}_{\geq p+\gamma}(\Box^6[|y| < 0.5 - \varepsilon])$. Here γ gives the accumulation of the error in the probability over the time horizon of interest: for this case we have $1 - \gamma := (1 - \delta)^6$, which is $\gamma \approx 6\delta$. We then apply FAUST² to obtain a grid-based approximation of the safety probability over the six time steps of the formula (which adds up to 30 minutes in the model), with an accuracy of 0.1. More precisely, we first quantise the input space (this on its own generates an exact simulation), then we apply FAUST² [Esmail Zadeh Soudjani et al., 2015] over the obtained continuous space, finite action model. For this work we have optimised the algorithms in FAUST² to use less memory for models with Gaussian noise: by first decoupling the noise by means of a simple state transform, the storage of the discretised probability transitions can be done in a structured and more efficient manner. This leads to perform the computations with 2.6×10^7 grid points to attain the desired accuracy of 0.1 (more precisely 0.0983) with a 2,6 GHz Intel Core i5 with 16 GB memory within less than 20 minutes. We finally obtain that the modified safety property is satisfied with probability of at least $0.8412 - 0.0983 = 0.7429$ for the reduced order model M_2 initialised at zero.

Control refinement: simulation results

We refine the policy obtained from FAUST² for the reduced-order model M_2 to the original model M . Recall that we expect this refined policy to have a quantifiable safety, expressed via the property ψ , which is a requirement that the inner air temperature remains within the bound $y_s \in [-0.5, 0.5]$ of the nominal temperature during the next 30 minutes. The safety probability for the concrete model M initialised at the origin is lower bounded by the computed probability $p = (0.7429 - \gamma) = (0.7429 - 0.0585) = 0.6844$ (this is according to Theorem 3.13).

We empirically validate this result as follows. We first initialise the system and the state of the reduced-order model (in the controller) at the origin. Then we perform 10^5 Monte-Carlo simulations and observe that executions of the reduced-order model remain in the modified safe set 85.81 percent of the time, whereas they exit it 14.19 percent of the time. For the same noise sequences, the controlled 5-dimensional model, where the control is refined based on the interface introduced before, stays in the *original* safe set 99.9 percent of the time, and exits it in 0.10 percent of the times. The concrete model is further seen to stay within the

modified safe set 86.05 percent of the times, which is much closer to the computed probability for the reduced-order model. Notice that these empirical outcomes are expected to be higher than indicated in the error bounds, as these bounds are conservative especially when considering states starting in the middle of the relation.

Similarly, starting at the edge of the modified safe set $y_s \in [0.2986, -0.2986]$ of the reduced-order model, we have considered the initialisation as follows $x_s(0) = [-0.4229 \ -0.2987]^T$ and $x(0) = Px_s(0)$, where P has been discussed above. For this initial state 0.7289 is the lower bound on the safety probability for the reduced-order model, and $p = 0.6704$ for the full-order model. With 10^5 empirical Monte-Carlo runs, we obtain that the reduced-order model stays in the modified safe set 84.30 percent of the time, whereas the concrete model with the refined control policy stays in the safe set in 99.87 percent of the runs. Similar results were obtained upon initialising at other points on the edges of the (modified) safe set, or on the edge of the relation.

3.6 Connections to literature and measurability issues

In this section we establish quantitative connections between the notion of approximate similarity that we have introduced for gMDPs and known and established concepts that have been discussed in the literature for processes that are special cases of gMDPs.

As measurability issues are key in this discussion we would like to first point out that the results in this paper can be extended to analytical spaces with universally measurable kernels. When we allow the gMDPs to have universally measurable kernels, we need to show the existence of a conditional probability measure $\mathbb{W}_T(dx'_1|x'_2, u_1, x_1, x_2)$: for this we refer to [Edalat, 1999] which discusses the existence of universally measurable regular conditional probabilities.

3.6.a Early results for Markov chains with finite state spaces

From the perspective of testing, the concept of probabilistic bisimulation has been first introduced in [Larsen and Skou, 1989], based on a relational notion, and later used to define equivalence between Labelled Markov processes (LMPs) [Desharnais et al., 2002]. LMPs are different from gMDPs in that transition are not governed by actions but by observable labels, and the acceptance of a label (and the consequent transition) defines the behaviour of such a process. LMPs are defined over a finite state space \mathbb{S} , a set of labels L , and stochastic transition kernels $\mathbb{T}_l : \mathbb{S} \times \mathbb{S} \rightarrow [0, 1]$ that are finitely indexed by $l \in L$. There is a strong relationship between LMPs and standard MDPs with labels [Abate et al., 2014a], despite their different semantics.

Definition 3.17 (Probabilistic bisimulation (relational notion))

Let $T = (\mathbb{S}, \mathbb{P}_{l \in L}, L)$ be a labelled Markov chain, with L the finite set of labels. Then an

equivalence relation \equiv_p on \mathbb{S} is probabilistic bisimulation if whenever $s \equiv_p t$, the following holds:

$$\forall l \in L : \forall A \in \mathbb{S} / \equiv_p, \sum_{s' \in A} \mathbb{T}_l(s|s') = \sum_{s' \in A} \mathbb{T}_l(t|s').$$

Two states s and t are said to be probabilistically bisimilar ($s \sim_{SL} t$) if the pair (s, t) is contained in a probabilistic bisimulation relation.

An extension of this definition is used to compare two separate processes by combining their state spaces (as a disjoint union) and defining the probabilistic bisimulation on the obtained extended state space [Desharnais et al., 2002]. (More details on this operation is given in the following subsection for continuous state-space models.)

For countable-state probabilistic processes combining probability and non-determinism, [Segala, 1995, Segala and Lynch, 1995] has discussed probabilistic simulations based on a lifting notion – this has inspired the extension (over more general models) that is elaborated in this work. Over finite- or countable-state sets, [Segala, 1995, Lemma 8.2.2] has shown that lifting coincides with \mathcal{R}_{eq} -equivalence of the corresponding probability distributions.

3.6.b Exact bisimulation relations for models with continuous state spaces

The early notion of bisimulation between labelled Markov chains [Larsen and Skou, 1989] has been extended to processes (again denoted as LMPs) defined over analytical state spaces in [Desharnais et al., 2002], by employing zigzag morphisms. This work combines and extends earlier results on zigzag-based bisimulations [Blute et al., 1997, Desharnais et al., 1998, Edalat, 1999], provides the fundamental measure theoretical results to support bisimulations over continuous spaces, and shows their logical characterisation and their transitivity property. Alternative but equivalent to the zigzag definition, the follow-up work in [Desharnais et al., 2003] discusses an extension of the relational notion in [Larsen and Skou, 1989], based on the concept of measurable \mathcal{R}_{eq} -closed sets.

Suppose that we have a LMP $\mathbf{S} = (\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T}_l, L)$, with a finite label set $l \in L$ and with \mathbb{X} being a Polish space. Note that, unlike in the discrete-space case, this process is defined together with a Borel σ -algebra $\mathcal{B}(\mathbb{X})$. Then based on [Desharnais et al., 2003] an equivalence relation, denoted \mathcal{R}_{eq} , defines a bisimulation iff for any $x_1 \mathcal{R}_{eq} x_2$ and for any measurable \mathcal{R}_{eq} -closed set B (or equivalently for every measurable set $B \subset \mathbb{X} / \mathcal{R}_{eq}$) it holds that

$$\mathbb{T}_l(B|x_1) = \mathbb{T}_l(B|x_2), \forall l \in L.$$

As an extension, a bisimulation between two different LMPs $\mathbf{S}_i = (\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{T}_{l,i}, L)$, $i = 1, 2$ can be constructed by working on the disjoint union of their state spaces. More precisely, an equivalence relation \mathcal{R}_{eq} over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ defines a bisimulation iff

for every $x_1 \mathcal{R}_{eq} x_2$ (where $x_1 \in \mathbb{X}_1$ and $x_2 \in \mathbb{X}_2$) and for every \mathcal{R}_{eq} -closed set B , it holds that

$$\mathbb{T}_{l,1}(B \cap \mathbb{X}_1 | x_1) = \mathbb{T}_{l,2}(B \cap \mathbb{X}_2 | x_2), \quad \forall l \in L.$$

An example of an equivalence relation over the disjoint union between two heterogeneous spaces, along with the induced quotient space, is given in Figure 3.4a. The discussed notion of equivalence between LMPs crucially depends on the equivalence of the *probability spaces* $(\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{P}_i)$ with probability measures $\mathbb{P}_i := \mathbb{T}_{l,i}(\cdot | x_i)$, given for a fixed l and state x_i . For an equivalence relation \mathcal{R}_{eq} over $\mathbb{X}_1 \sqcup \mathbb{X}_2$, the probability spaces are equivalent if for every measurable \mathcal{R}_{eq} -closed set B it holds that

$$\mathbb{P}_1(B \cap \mathbb{X}_1) = \mathbb{P}_2(B \cap \mathbb{X}_2),$$

which is denoted as $\mathbb{P}_1 \equiv_{\mathcal{R}_{eq}} \mathbb{P}_2$.

This type of equivalence between *probability spaces* has also been used for bisimulation relations between control Markov processes [Abate, 2013], a simpler instance of the gMDP framework discussed in this work. As such, it is a natural extension of the notion in [Desharnais et al., 2002, 2003] from LMPs to control Markov processes.

An equivalence relation defined over the disjoint union of \mathbb{X}_1 , and \mathbb{X}_2 , i.e., $\mathcal{R}_{eq} \subset (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, can also be expressed as a relation over their Cartesian product, namely $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$. As an example, we provide in Figure 3.4b the relation over the Cartesian product of two spaces, corresponding to the equivalence relation defined in Figure 3.4a over their disjoint union. This connection raises the question of whether probability spaces related via \mathcal{R}_{eq} are also in a lifted relation. When working with finite or countable sets, we know that this connection holds [Segala, 1995]. On the other hand, for continuous or uncountable spaces this depends on the absence of measure-theoretical issues, and will be studied in depth to answer when the following claim holds.

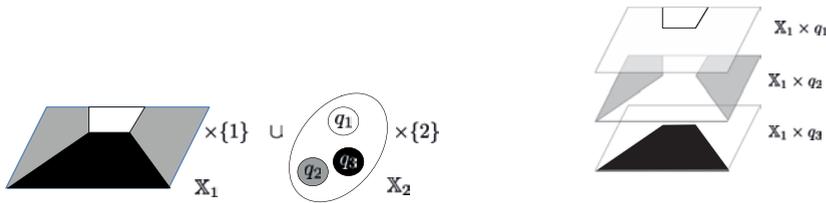
Claim 1 Consider two measure spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ and an equivalence relation \mathcal{R}_{eq} that induces a relation over $\mathbb{X}_1 \times \mathbb{X}_2$ as $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$. Then,

- for any two probability measures $\Delta \in \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\Theta \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$, we have

$$\Delta \bar{\mathcal{R}} \Theta \text{ if and only if } \Delta \equiv_{\mathcal{R}_{eq}} \Theta.$$

- for any two universally measurable transition kernels \mathbb{T}_1 and \mathbb{T}_2 , there exists a universally measurable kernel $\mathbb{W}_{\mathbb{T}}$ that lifts the transition kernels for \mathcal{R} as required in Def. 3.6.

In order to prove this claim and to construct the lifted measure based on an equivalence relation, we exploit the notion of zigzag morphism [Desharnais et al., 2002, Edalat, 1999] and its properties. More precisely, consider a tuple $(\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T})$, with \mathbb{X} a Polish space and $\mathbb{T} : \mathbb{X} \times \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$ a transition probability function.



(a) An equivalence relation \mathcal{R}_{eq} over the disjoint union $\mathbb{X}_1 \sqcup \mathbb{X}_2$, where two elements from each set are in the relation if they share the same colour.

(b) Relation \mathcal{R} over the Cartesian product of $\mathbb{X}_1 \subset \mathbb{R}^2$ and $\mathbb{X}_2 = \{q_1, q_2, q_3\}$, induced by the equivalence relation \mathcal{R}_{eq} . Elements of the relation are coloured.

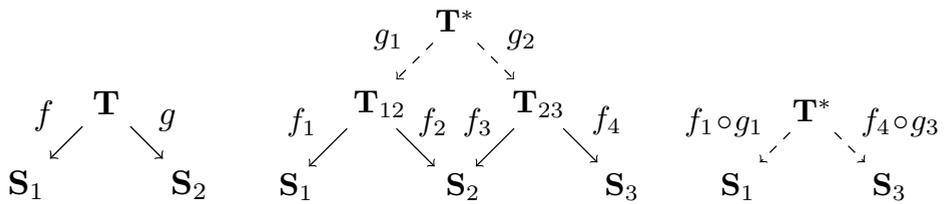
Figure 3.4: Example of an equivalence relation over the disjoint union of two heterogeneous spaces, and the corresponding relation over their Cartesian product.

Definition 3.18 (Morphism) A function $f : (\mathbb{X}, \mathcal{B}(\mathbb{X}), \mathbb{T}) \rightarrow (\mathbb{X}', \mathcal{B}(\mathbb{X}'), \mathbb{T}')$ is a morphism if it is a continuous surjective map $f : \mathbb{X} \rightarrow \mathbb{X}'$, such that for all $s \in \mathbb{X}$ and for all $B \in \mathcal{B}(\mathbb{X})$,

$$\mathbb{T}(f^{-1}(B)|s) = \mathbb{T}'(B|f(s)),$$

i.e., it is preserving transition probabilities.

Consider two labelled Markov processes $S_i = (\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \{k_{l,i} | l \in L\})$ with a shared finite set of labels L , then a morphism f is a *zigzag morphism* if it preserves the two transition probability functions for all $l \in L$. Two LMPs S_1 and S_2 are *probabilistically bisimilar* if there is a generalised span of zigzag morphisms between them [Desharnais et al., 2002]; namely, if there exists a labelled Markov process T (with universally measurable transition kernels) and zigzag morphisms f and g from T to S_1 and S_2 , respectively (see Figure 3.5 (a)). In order to prove



(a) Generalised span of zigzag morphisms

(b) Construct T^* as a semi-pullback of co-span $T_{12} \rightarrow S_2 \leftarrow T_{22}$.

(c) Transitive bisimulation based on semi-pullback

Figure 3.5: Probabilistic bisimulation between S_1 and S_2 established by zigzag morphism. Transitivity of probabilistic bisimulations S_1 and S_2 and S_2 and S_3 follows as a semi-pullback.

that this notion of probabilistic bisimulation is transitive, [Edalat, 1999] has shown that

- the category of Markov processes with universally measurable transition

probability functions \mathbb{T} on Polish spaces and with surjective and continuous transition probability preserving maps has *semi-pullbacks* [Edalat, 1999, Corollary 5.3];

- the category of probability measures \mathbb{P} on Polish spaces and measure-preserving surjective maps has *semi-pullbacks* [Edalat, 1999, Corollary 5.4].

By adding a labelling to the transition probability function \mathbb{T} , one can trivially show the existence of semi-pullbacks on an LMP. Moreover, the transitivity of probabilistic bisimulations follows based on semi-pullbacks: if \mathbf{S}_1 is probabilistically bisimilar to \mathbf{S}_2 , which is also bisimilar to \mathbf{S}_3 , then \mathbf{S}_1 and \mathbf{S}_3 are bisimilar, as in Figure 3.5b.

Let us go back to Claim 1. Firstly recall that, as depicted in Figure 3.4a, an equivalence relation \mathcal{R}_{eq} over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ induces a quotient space, denoted by $\mathcal{Q} := (\mathbb{X}_1 \sqcup \mathbb{X}_2)/\mathcal{R}_{eq}$, and partitions the unionised state space by disjoint sets, namely $\bigcup_{q \in \mathcal{Q}} q = \mathbb{X}_1 \sqcup \mathbb{X}_2$, and $q_1 \cap q_2 = \emptyset$ for $q_1 \neq q_2$, $q_1, q_2 \in \mathcal{Q}$. Thus starting from the Markov processes $\mathbf{S}_1 = (\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1), \mathbb{T}_1)$ and $\mathbf{S}_2 = (\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2), \mathbb{T}_2)$, we show that the claim holds under either of the following two conditions.

Condition 1 (Polish quotient space) *The equivalence relation of interest \mathcal{R}_{eq} induces a quotient space $(\mathcal{Q}, \mathcal{F})$ that is Polish and the maps from \mathbb{X}_1 and \mathbb{X}_2 to the quotient space $f_1 : \mathbb{X}_1 \rightarrow \mathcal{Q}$ and $f_2 : \mathbb{X}_2 \rightarrow \mathcal{Q}$ are measurable and surjective.*

Condition 2 (Analytic Borel quotient space) *The equivalence relation of interest \mathcal{R}_{eq} induces a quotient space that is analytical as in [Desharnais et al., 2002, Edalat, 1999] and the maps from \mathbb{X}_1 and \mathbb{X}_2 to the quotient space $f_1 : \mathbb{X}_1 \rightarrow \mathcal{Q}$ and $f_2 : \mathbb{X}_2 \rightarrow \mathcal{Q}$ are measurable and surjective.*

Notice that condition 1 implies condition 2, and further note that f_1 and f_2 are constructed based on the injection ι_1 and ι_2 , i.e., $\iota_i : \mathbb{X}_i \rightarrow \mathbb{X}_1 \sqcup \mathbb{X}_2$ for $i = 1, 2$, composed with $q : \mathbb{X}_1 \sqcup \mathbb{X}_2 \rightarrow \mathcal{Q}$.

Then we can construct the quotient Markov process as the tuple $\mathbf{S} := (\mathcal{Q}, \mathcal{F}, \mathbb{T})$ such that $(\mathcal{Q}, \mathcal{F})$ is a Borel measurable space with $\mathcal{Q} = (\mathbb{S}_1 \sqcup \mathbb{S}_2)/\mathcal{R}_{eq}$, and \mathcal{F} is defined as $\mathcal{F} := \{E \subset \mathcal{Q} : q^{-1}(E) \in \mathcal{B}(\mathbb{S}_1 \sqcup \mathbb{S}_2)\}$. The stochastic transition kernel \mathbb{T} is constructed as in [Desharnais et al., 2002, Proof of Proposition 9.4]. For any $B \in \mathcal{F}$ it holds that

$$\mathbb{T}(B|t) = \mathbb{T}_1(f_1^{-1}(B)|s) \quad \text{with } s \in f_1^{-1}(t) \quad (3.6)$$

and $\mathbb{T}(B|\cdot)$ is Borel measurable.

Then f_1 and f_2 are zigzag morphisms from, respectively, \mathbf{S}_1 and \mathbf{S}_2 to \mathbf{S} , and they form a co-span. Based on [Edalat, 1999] we now know that there exists a Markov process $\mathbf{W} := ((\mathbb{X}_1 \times \mathbb{X}_2), \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2), \mathbb{W})$, which is a semi-pullback, and where \mathbb{W} lifts the relation over $\mathbb{X}_1 \times \mathbb{X}_2$ and defines a universally measurable stochastic kernel. If \mathbf{S}_1 , \mathbf{S}_2 and \mathbf{S} have analytical Borel spaces (this includes Polish spaces)

and universally measurable transition kernels, then $\mathbb{W} : \mathcal{R} \times \mathcal{B}(\times)$ is defined as

$$\mathbb{W}(dx'_1 \times dx'_2 \mid (x_1, x_2)) = \int_{q' \in Q} \mathbb{T}_1(dx'_1 \mid x_1, q') \mathbb{T}_2(dx'_2 \mid x_2, q') \mathbb{T}(dq' \mid f_1(x_1)), \quad (3.7)$$

where $\mathbb{T}_i(dx'_i \mid x_i, q')$ for $i = 1, 2$ are universally measurable regular conditional probability distributions, such that for measurable subsets $X_i \subset \mathbb{X}_i$ and $Q \subset \mathcal{Q}$ it holds that

$$\mathbb{T}_i(X_i \cap f_1^{-1}(Q) \mid x_i) = \int_Q \mathbb{T}_i(dx'_i \mid x_i, q') \mathbb{T}(dq' \mid f_1(x_1)).$$

The details of this reasoning follow from [Edalat, 1999] together with the existence proof for the regular conditional probability distributions.

Remark 3.10 (Measurability assumptions) *The measurability assumption above is a nontrivial but natural assumption, since, as proven for LMPs, any equivalence relation on $\mathbb{X}_1 \sqcup \mathbb{X}_2$ based on logics induces a quotient LMP that has an analytical Borel space and measurable canonical maps [Desharnais et al., 2002, Proposition 9.4].*

3.6.c Approximate probabilistic bisimulation relations

A relaxation of exact equivalence relations in a probabilistic context has been first introduced for (finite-state) labelled Markov chains in [Desharnais et al., 2004], and later employed in [D’Innocenzo et al., 2012].

Definition 3.19 *A relation $\mathcal{R} \subseteq S \times S$ is an (probabilistic) ϵ -simulation if whenever sRt , then for all labels $l \in L$, and sets in the event space $X \in \Sigma$, it holds that*

$$\mathbb{T}_l(\mathcal{R}(X) \mid t) \geq \mathbb{T}_l(X \mid s) - \epsilon.$$

Note that the relation is not required to be an equivalence relation, hence it does not induce a partitioning of the state space. For continuous-space systems, [Abate, 2013] has discussed an approximate (bi-)simulation notion derived from the finite-state definition. This definition relates to an approximate equivalence of the probability spaces $(\mathbb{X}_i, \mathcal{B}(\mathbb{X}_i), \mathbb{P}_i)$ $i = 1, 2$ as follows. For an equivalence relation \mathcal{R}_{eq} over $\mathbb{X}_1 \sqcup \mathbb{X}_2$ the probability spaces are approximately equivalent if for every measurable \mathcal{R}_{eq} -closed set B it holds that

$$|\mathbb{P}_1(B \cap \mathbb{X}_1) - \mathbb{P}_2(B \cap \mathbb{X}_2)| \leq \delta,$$

which is denoted as $\mathbb{P}_1 \equiv_{\mathcal{R}_{eq}}^{\delta} \mathbb{P}_2$.

Theorem 3.20 *Consider two measure spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ and an equivalence relation \mathcal{R}_{eq} satisfying condition 1. Then for any two probability measures $\Delta \in$*

$\mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $\Theta \in \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ we have that

$$\Delta \equiv_{\mathcal{R}_{eq}}^{\delta} \Theta \text{ if and only if } \Delta \bar{\mathcal{R}}_{\delta} \Theta,$$

with as standard $\mathcal{R} := \{(x_1, x_2) \in \mathbb{X}_1 \times \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{eq}\}$.

Proof: 1. $\Delta \bar{\mathcal{R}}_{\delta} \Theta \implies \Delta \equiv_{\mathcal{R}_{eq}}^{\delta} \Theta$

If $\Delta \bar{\mathcal{R}}_{\delta} \Theta$ then for each $C \subset (\mathbb{X}_1 \sqcup \mathbb{X}_2) / \mathcal{R}_{eq}$ with subsets $\tilde{S} = \mathbb{X}_1 \cap C \in \mathcal{B}(\mathbb{X}_1)$ and $\tilde{T} = \mathbb{X}_2 \cap C \in \mathcal{B}(\mathbb{X}_2)$, then $|\Delta(\tilde{S}) - \Theta(\tilde{T})| \leq \delta$ because $\mathbb{W}(\tilde{S} \times (\mathbb{X}_2 \setminus \tilde{T})) \leq \delta$ and $\mathbb{W}((\mathbb{X}_1 \setminus \tilde{S}) \times \tilde{T}) \leq \delta$. This can be shown as follows

$$\Delta(\tilde{S}) \leq \Delta(\tilde{S}) + \mathbb{W}((\mathbb{X}_1 \setminus \tilde{S}) \times \tilde{T}) = \Theta(\tilde{T}) + \mathbb{W}(\tilde{S} \times (\mathbb{X}_2 \setminus \tilde{T})) \leq \Theta(\tilde{T}) + \delta$$

and repeating the reasoning starting from $\Theta(\tilde{T})$ we get $\Theta(\tilde{T}) \leq \Delta(\tilde{S}) + \delta$, and hence $|\Delta(\tilde{S}) - \Theta(\tilde{T})| \leq \delta$.

2. $\Delta \equiv_{\mathcal{R}_{eq}}^{\delta} \Theta \implies \Delta \bar{\mathcal{R}}_{\delta} \Theta$

Under Condition 1 we have that the quotient space has the Borel measure space $(\mathcal{Q}, \mathcal{F})$ where \mathcal{Q} is Polish. Additionally we have measurable mappings $f_i : \mathbb{X}_i \rightarrow \mathcal{Q}$. We denote the induced probability measures $f_{1*}\Delta \in \mathcal{P}(\mathcal{Q}, \mathcal{F})$ and $f_{2*}\Theta \in \mathcal{P}(\mathcal{Q}, \mathcal{F})$. Denote a measure that lifts these over the diagonal relation as $\mathbb{W}_{\mathcal{Q}} \in \mathcal{P}(\mathcal{Q}^2, \mathcal{F}^2)$. This is equivalent to maximal coupling of $f_{1*}\Delta$ and $f_{2*}\Theta$. Specifically for Polish spaces we take the γ -coupling given as $\mathbb{W}_{\mathcal{Q}} := \gamma(f_{1*}\Delta, f_{2*}\Theta) \in \mathcal{P}(\mathcal{Q}^2, \mathcal{F}^2)$ [Abate et al., 2014b] based on [Lindvall, 2002, Section 1.5] and given as follows

Definition 3.21 Let Z be a Borel space and let $\nu, \tilde{\nu} \in (Z)$ be two probability measures on it. The γ -coupling of $(\nu, \tilde{\nu})$ is a measure $\gamma \in (Z^2)$ given by

$$\gamma(\nu, \tilde{\nu}) := \Psi_Z(\nu \wedge \tilde{\nu}) + \mathbf{1}_{[0,1)}(\|\nu \wedge \tilde{\nu}\|) \cdot \frac{(\nu - \tilde{\nu})^+ \otimes (\nu - \tilde{\nu})^-}{1 - \|\nu - \tilde{\nu}\|}$$

where $\Psi_Z : Z \rightarrow Z^2$ is the diagonal map on Z given by $\Psi_Z : z \mapsto (z, z)$.

The lifted measure over $\mathbb{W} \in \mathcal{P}(\mathbb{X}_1 \times \mathbb{X}_2, \mathcal{B}(\mathbb{X}_1 \times \mathbb{X}_2))$ is given as

$$\mathbb{W} := \int_{\mathcal{Q} \times \mathcal{Q}} \Delta(dx_1 | q_1) \Theta(dx_2 | q_2) \mathbb{W}_{\mathcal{Q}}(dq_1 \times dq_2).$$

□

3.7 Conclusions

In this work we have discussed new and general approximate similarity relations for general control Markov processes, and shown that they can be effectively employed for abstraction-based verification goals as well as for controller synthesis and refinement over quantitative specifications. The new relations in particular

allow for a useful trade-off between the deviations in probability distribution on states and the deviations between model outputs. We have extended results on control refinement for deterministic LTI systems to construct interface functions effectively. For this and other model classes within the set of gMDPs the algorithmic construction of appropriate interface functions together with the optimal quantification of the ε, δ -approximate similarity relation is topic of further research. Alongside practical applications of the developed notions, current efforts focus on further generalisation of Theorem 3.13 to specific quantitative properties expressed via temporal logics. We are moreover interested in further expanding our understanding of the properties of similarity relations.

List of symbols

General

$\mathbf{d}_{\mathbb{Y}}$	metric on the space \mathbb{Y}
$\mathcal{B}(\mathbb{Y})$	is the Borel σ -algebra on space \mathbb{Y}
$\mathcal{N}(0, \Sigma)$	Zero-mean Gaussian distribution with co-variance matrix Σ .

Markov processes

\mathbf{M}	Markov decision process (MDPs) or general Markov decision process (gMDP) cf. Definitions 3.1 and 3.2
$\pi_{x(0)}$	Initial probability distribution
\mathbb{X}	State space (restricted to Polish spaces) of MDP or gMDP
\mathbb{U}	Set of possible actions (restricted to Polish space) of MDP or gMDP
\mathbb{T}	Borel measurable, conditional stochastic kernel of an MDP Defines probability distribution of next state $x(t+1) \in \mathbb{X}$ conditional on $(x(t), u(t)) \in \mathbb{X} \times \mathbb{U}$.
\mathbb{Y}	Output space
h	Measurable output mapping; $h : \mathbb{X} \rightarrow \mathbb{Y}$
$\mathcal{M}_{\mathbb{Y}}$	class of all gMDPs with output space \mathbb{Y}

Control

μ_u	Stochastic control input $\mu_u : \mathcal{B}(\mathbb{U}) \rightarrow [0, 1]$
μ	Markov policy $\mu := (\mu_0, \mu_1, \mu_2, \dots)$ with $\mu_i : \mathbb{X} \rightarrow \mathcal{P}(\mathbb{U}, \mathcal{B}(\mathbb{U}))$ cf. Definition 3.3
\mathbf{C}	Control strategy, modelled as a time-inhomogeneous Markov decision process cf. Definition 3.4
\mathcal{U}_v	Interface function introduced in Section 3.3.c $\mathcal{U}_v : \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2 \rightarrow \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2))$
$\mathbf{C} \times \mathbf{M}$	Feedback composition of \mathbf{C} with \mathbf{M}
$\mathbb{P}_{\mathbf{C} \times \mathbf{M}}$	Probability measure associated to the composed stochastic process

Relations and preorders

\mathcal{R}	Relation over $\mathcal{R} \subseteq \mathbb{X}_1 \times \mathbb{X}_2$.
$\bar{\mathcal{R}}$	Relation over $\bar{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ obtained via lifting from \mathcal{R} , as per Def. 3.5.
$\bar{\mathcal{R}}_\delta$	Relation over $\bar{\mathcal{R}} \subseteq \mathcal{P}(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1)) \times \mathcal{P}(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ obtained via the approximate lifting with a deviation in probability bounded with δ obtained from \mathcal{R} , as per Def. 3.10.
$\equiv_{\mathcal{R}_{eq}}$	Relation between two probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ based on the equivalence relation $\mathcal{R}_{eq} \subseteq (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, à la [Desharnais et al., 2003], as reviewed in Section 3.6.
$\equiv_{\bar{\mathcal{R}}_{eq}}^\delta$	Approximate relation between two probability spaces $(\mathbb{X}_1, \mathcal{B}(\mathbb{X}_1))$ and $(\mathbb{X}_2, \mathcal{B}(\mathbb{X}_2))$ based on the equivalence relation $\mathcal{R}_{eq} \subseteq (\mathbb{X}_1 \sqcup \mathbb{X}_2) \times (\mathbb{X}_1 \sqcup \mathbb{X}_2)$, à la [Abate, 2013], as reviewed in Section 3.6.
\sqsubseteq	Probabilistic simulation relation, see Def. 3.6.
\approx	Probabilistic bisimulation relation, see Def. 3.7.
$\sqsubseteq_\varepsilon^\delta$	ε, δ -approximate probabilistic simulation relation, see Def. 3.11.

List of definitions

- A *metric* on \mathbb{Y} is a function $d_{\mathbb{Y}} : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}^+$ satisfying the following conditions $\forall y_1, y_2, y_3 \in \mathbb{Y}$:
 1. $d_{\mathbb{Y}}(y_1, y_2) = 0 \Leftrightarrow y_1 = y_2$;
 2. $d_{\mathbb{Y}}(y_1, y_2) = d_{\mathbb{Y}}(y_2, y_1)$;
 3. $d_{\mathbb{Y}}(y_1, y_3) \leq d_{\mathbb{Y}}(y_1, y_2) + d_{\mathbb{Y}}(y_2, y_3)$.
- A topological space is called *separable* if it contains a countable, dense subset.
- A *Polish space* is separable and completely metrisable topological space.
- A *Borel set* is any set that can be generated from open sets based on countable unions, intersections and complements.
- A σ -algebra on the space \mathbb{Y} $\mathcal{F}_{\mathbb{Y}}$ is a collection of subsets of \mathbb{Y} that includes the empty subset, is closed under complement, countable unions and intersections.
- A *Borel σ -algebra* over a topological space \mathbb{Y} is the smallest class of sets that is closed under countable unions and intersections and contains every closed (open) set, see [Bertsekas and Shreve, 1996, Prop. 7.11].
- A *complete measure space* is a measure space in which every subset of every null set is measurable (with measure zero). As an example the completion of the Borel measure gives the complete Lebesgue measure.

- A subset of A of a Polish space \mathbb{Y} is *universally measurable* if it is measurable w.r.t. every complete probability measure on \mathbb{Y} that measures all Borel subsets of \mathbb{Y}
- $\mathcal{F}_{\mathbb{Y}_1} \otimes \mathcal{F}_{\mathbb{Y}_2}$ is the σ -algebra on $\mathbb{Y}_1 \times \mathbb{Y}_2$ generated by sets $F_1 \times F_2$ with $F_1, F_2 \in \mathcal{F}_{\mathbb{Y}_1}, \mathcal{F}_{\mathbb{Y}_2}$

3.A Details on case study and use of FAUST²

The model reduction procedure via balanced truncation⁹ yields four reduced-order models $\mathbf{M}_i = (A_i, B_i, B_{wi}, C_i)$ $i = 1, 2, 3, 4$:

$$\mathbf{M}_i : \begin{cases} x_s(t+1) &= A_i x_s(t) + B_{wi} w(t) + B_i u_s(t) \\ y_s(t) &= C_i x_s(t), \end{cases}$$

which are characterised by the following constant matrices

$$\begin{aligned} \mathbf{M}_1 : A_1 &= \begin{bmatrix} 0 & -0.8572 \\ 1 & 1.857 \end{bmatrix}, B_1 = \begin{bmatrix} -0.5343 \\ 0.5523 \end{bmatrix}, B_{w1} = \begin{bmatrix} -5.916\text{e-}3 & -0.0564 & 8.62\text{e-}3 \\ 6.138\text{e-}3 & 0.05852 & -6.739\text{e-}3 \end{bmatrix}, C_1 = [0 \ 1]; \\ \mathbf{M}_2 : A_2 &= \begin{bmatrix} 0 & -0.05267 \\ 0.125 & -0.1081 \end{bmatrix}, B_2 = \begin{bmatrix} 0.8917 \\ 0.3725 \end{bmatrix}, B_{w2} = \begin{bmatrix} 0.01925 & 0.1835 & 0.002356 \\ 0.01372 & 0.1308 & 3.229\text{e-}5 \end{bmatrix}, C_2 = [0 \ 1]; \\ \mathbf{M}_3 : A_3 &= [0.9951], B_3 = [0.1194], B_{w3} = [0.001497 \ 0.01427 \ 0.01467], C_3 = [1]; \\ \mathbf{M}_4 : A_4 &= [0.1203], B_4 = [0.3829], B_{w4} = [0.01257 \ 0.1198 \ 0.0002907], C_4 = [1]. \end{aligned}$$

Models \mathbf{M}_1 and \mathbf{M}_3 are obtained from $\mathbf{M} = (A, B, B_w, C)$, whereas \mathbf{M}_2 and \mathbf{M}_4 are based on the dynamics of $\mathbf{M}' = (A + BF, B, B_w, C)$. We have synthesised F to be $[0.4846 \ 0.3986 \ 0.8535 \ 0.5639 \ 0.002425]$. As expected the reduced models depend on the choice of \mathbf{M}' or \mathbf{M} : in the former case, the part of the dynamics that we cannot compensate with a control is approximated best, whereas for \mathbf{M} the most prominent dynamics are approximated best.

Approximate probabilistic simulation relation

We quantify the performance of \mathbf{M}_i for $i = 1, 2, 3, 4$ only over a bounded input set $\mathbb{U}_s := \{u_s \in \mathbb{R} \mid u_s^2 \leq c_1\}$.

Subsequently solving the Sylvester equations for Q, P and R , tuning a stabilising interface gain K , and then using the S -procedure as described in [Boyd and Vandenberghe, 2004] to compute ε, δ and M , we finally obtain the following matrices for the reduced-order models. For \mathbf{M}_1 we take $R := 1.403$, and we obtain

$$\begin{aligned} Q &:= [-0.08954 \ -0.07712], & K &:= [-0.5717 \ -0.4705 \ -0.9859 \ -0.6213 \ -0.002364], \\ P &:= \begin{bmatrix} -1.061 & 0.09045 \\ 0 & 1 \\ -2.295 & -0.9696 \\ 9.064 & 8.775 \\ 0 & 0 \end{bmatrix}, & M &:= \begin{bmatrix} 0.4797 & 0.1476 & 0.3298 & 0.1397 & -0.001306 \\ 0.1476 & 1.104 & 0.1592 & 0.06704 & -0.00359 \\ 0.3298 & 0.1592 & 0.2862 & 0.1207 & -0.001327 \\ 0.1397 & 0.06704 & 0.1207 & 0.1744 & 0.003174 \\ -0.001306 & -0.00359 & -0.001327 & 0.003174 & 0.003676 \end{bmatrix}. \end{aligned}$$

⁹This is obtained from the application of the `balred` function in `Matlab`.

Note that the latter is optimised for $\delta = 10^{-2}$.

For \mathbf{M}_2 we take $R := 1.004$, and obtain

$$Q := [-1.857 \ 1.406], \quad K := [-0.3553 \ -0.2931 \ -0.65 \ -0.4739 \ -0.002547],$$

$$P := \begin{bmatrix} -0.6186 & 0.2348 \\ 0 & 1 \\ 2.562 & -2.314 \\ -0.009378 & 0.001329 \\ 0 & 0 \end{bmatrix}, \quad M := \begin{bmatrix} 0.2416 & 0.06342 & 0.3159 & 0.1299 & 0.00106 \\ 0.06342 & 1.772 & 0.07267 & 0.02663 & 0.0007664 \\ 0.3159 & 0.07267 & 0.4191 & 0.1728 & 0.001395 \\ 0.1299 & 0.02663 & 0.1728 & 0.08168 & 0.000351 \\ 0.00106 & 0.0007664 & 0.001395 & 0.000351 & 0.0001456 \end{bmatrix}.$$

Again M is chosen based on the S procedure to optimise ε for $\delta = 10^{-2}$. For \mathbf{M}_3 , take $R := 0.3074$ and obtain

$$Q := -0.0008755, \quad K := [-0.5796 \ -0.477 \ -0.9978 \ -0.6265 \ -0.00236],$$

$$P := \begin{bmatrix} 1.004 \\ 1 \\ 1.006 \\ 0.9713 \\ 0 \end{bmatrix}, \quad M := \begin{bmatrix} 8.584 & -4.974 & 4.929 & 2.078 & 0.1158 \\ -4.974 & 3.944 & -3.106 & -1.31 & -0.05919 \\ 4.929 & -3.106 & 3.917 & 1.653 & 0.06135 \\ 2.078 & -1.31 & 1.653 & 0.7024 & 0.02595 \\ 0.1158 & -0.05919 & 0.06135 & 0.02595 & 0.01179 \end{bmatrix}.$$

Note that M is chosen based on the S -procedure to optimise ε for $\delta = 10^{-2}$.

For \mathbf{M}_4 , we take $R := 0.8996$ and

$$Q := -0.6961, \quad K := [-0.5307 \ -0.4366 \ -0.9241 \ -0.5946 \ -0.002391],$$

$$P := \begin{bmatrix} -1.191 \\ 1 \\ 1.242 \\ -0.01296 \\ 0 \end{bmatrix}, \quad M := \begin{bmatrix} 0.03949 & -0.01465 & 0.06076 & 0.02542 & 1.999e-05 \\ -0.01465 & 1.788 & 0.1162 & 0.05143 & -0.0005164 \\ 0.06076 & 0.1162 & 0.128 & 0.05469 & -2.765e-05 \\ 0.02542 & 0.05143 & 0.05469 & 0.04108 & -0.0004062 \\ 1.999e-05 & -0.0005164 & -2.765e-05 & -0.0004062 & 0.0003725 \end{bmatrix}.$$

δ	1	$10^{-\frac{1}{3}}$	$10^{-\frac{2}{3}}$	10^{-1}	$10^{-\frac{4}{3}}$	$10^{-\frac{5}{3}}$	10^{-2}	$10^{-\frac{7}{3}}$	$10^{-\frac{8}{3}}$	10^{-3}
\mathbf{M}_1	0.1233	0.4803	0.6247	0.7347	0.827	0.9082	0.9816	1.049	1.112	1.171
\mathbf{M}_2	0.01445	0.1037	0.132	0.1534	0.1714	0.1871	0.2014	0.2145	0.2267	0.2381
\mathbf{M}_3	0.05206	0.7612	0.997	1.175	1.325	1.456	1.575	1.684	1.785	1.881
\mathbf{M}_4	0.1839	0.3029	0.3358	0.3604	0.3809	0.3988	0.415	0.4298	0.4435	0.4564

Table 3.6: Trade-off for parameters ε, δ in the simulation relation.

3.A.a FAUST² computations on a 2-dimensional model

For a given x, u pair the probability distribution of the next state is distributed with the following stochastic density kernel $t_x(\bar{x} \mid x, u) \sim \mathcal{N}(\cdot; A_i x + B_i u, \Sigma)$, where $\Sigma := B_{w_2} B_{w_2}^T$.

We resort to the algorithms implemented in [Esmaeil Zadeh Soudjani et al., 2015] to maximise the probability of a stochastic event. We set up a stochastic dynamic programming scheme, leading to a final value function providing the probability of the property as

$$V_0(x) = \mathbb{P} [\square^6(|y(t)| \leq 0.5 - \varepsilon)].$$

Define the safe set $\mathcal{A} := \mathbb{R} \times [-0.5 + \varepsilon, 0.5 - \varepsilon] \subset \mathbb{X} = \mathbb{R}^2$, then the property to be maximised can be written as $V_0(x) = \mathbb{P} [\square^6 \mathcal{A}]$.

3.A.a.1 The error computation

Assume there are constants H_1, H_2 , such that

$$\int_{\mathbb{R}^2} |t_x(\bar{x} | x, u) - t_x(\bar{x} | x', u)| d\bar{x} \leq H_1 |x'_1 - x_1| + H_2 |x'_2 - x_2|. \quad (3.8)$$

This gives a linearly increasing error $N(H_1\Delta_1 + H_2\Delta_2)$, where Δ_i is the grid size in the i -th coordinate direction of the state space. Let us compute the two constants next. Starting from

$$t_x(\bar{x} | x, u) = \frac{1}{\sqrt{(2\pi)^2 \det(\sigma)}} \exp \left[-\frac{1}{2} (\bar{x} - A_i x - B_i u)^T \Sigma^{-1} (\bar{x} - A_i x - B_i u) \right],$$

define $m = \begin{bmatrix} m_1 \\ m_2 \end{bmatrix} = A_i x + B_i u$ and $\Sigma^{-1} = \begin{bmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{bmatrix} = L^T L$. Then

$$t_x(\bar{x} | x, u) = \frac{1}{\sqrt{(2\pi)^2 \det(\sigma)}} \exp \left[-\|L\bar{x} - Lm\|^2 \right].$$

Define a change of variables with $v = L\bar{x} \rightarrow dv = |\det(L)|d\bar{x}$. Then the error computation follows from the maximal difference between the probability density distributions [Esmail Zadeh Soudjani et al., 2015] as given in (3.8) and can be rewritten as follows:

$$\int_{\mathbb{R}^2} \left| \frac{1}{\sqrt{(2\pi)^2 \det(\Sigma)}} \left(\exp \left[-\frac{1}{2} \|v - Lm\|^2 \right] - \exp \left[-\frac{1}{2} \|v - Lm'\|^2 \right] \right) \right| \frac{dv}{|\det(L)|}.$$

Note that $\Sigma^{-1} = L^T L$, hence $|\det(L)| = \frac{1}{\sqrt{\det(\Sigma)}}$ and consequently

$$= \int_{\mathbb{R}^2} \frac{1}{2\pi} \left| \left(\exp \left[-\frac{1}{2} \|v - Lm\|^2 \right] - \exp \left[-\frac{1}{2} \|v - Lm'\|^2 \right] \right) \right| dv.$$

Now we can transform a two-dimensional integral into two one-dimensional integrals:

$$\begin{aligned} &\leq \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \left| \left(\exp \left[-\frac{1}{2} \|v_1 - L_1 m_1\|^2 \right] - \exp \left[-\frac{1}{2} \|v_1 - L_1 m'_1\|^2 \right] \right) \right| dv_1 \\ &+ \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \left| \left(\exp \left[-\frac{1}{2} \|v_2 - L_2 m_2\|^2 \right] - \exp \left[-\frac{1}{2} \|v_2 - L_2 m'_2\|^2 \right] \right) \right| dv_2 \\ &\leq \frac{2|L_1 m - L_1 m'|}{\sqrt{2\pi}} + \frac{2|L_2 m - L_2 m'|}{\sqrt{2\pi}} \leq \frac{2}{\sqrt{2\pi}} (|L_1 A_i(x - x')| + |L_2 A_i(x - x')|). \end{aligned}$$

Define $\begin{bmatrix} \bar{a}_{11} & \bar{a}_{12} \\ \bar{a}_{21} & \bar{a}_{22} \end{bmatrix} = L A_i$. Then for (3.8) we have $H_1 = \frac{2}{\sqrt{2\pi}} (|\bar{a}_{11}| + |\bar{a}_{21}|)$, $H_2 = \frac{2}{\sqrt{2\pi}} (|\bar{a}_{12}| + |\bar{a}_{22}|)$.

3.B Proofs of Theorems and Corollaries

3.B.a Control refinement proofs, Theorem 1-4

Let us consider the controller refinement for exact simulation relations first. The execution $\{(x_2(t), x_{C_2}(t)) | t \in [0, N]\}$, is defined on the canonical space $\Omega = (\mathbb{X}_2 \times \mathbb{X}_{C_2})^{N+1}$, and has a unique probability measure $\mathbb{P}_{C_2 \times M_2}$. Therefore in Alg. 1, in order to write the execution of the refined control C_2 and of the gMDP M_2 , we have included the state of M_2 for one transition in the state of the refined control strategy. Therefore, while the execution of Alg. 1 ranges over $\mathbb{X}_{C_1} \times \mathbb{X}_1 \times \mathbb{X}_2$, the execution of the controlled system with C_2 ranges over $\mathbb{X}_{C_2} \times \mathbb{X}_2 = (\mathbb{X}_{C_1} \times \mathbb{X}_1 \times \mathbb{X}_2) \times \mathbb{X}_2$. The marginal of $\mathbb{P}_{C_2 \times M_2}$ on $\mathbb{X}_{C_1} \times \mathbb{X}_1 \times \mathbb{X}_2$ defines the measure for the execution in Alg.1.

Since, by the above construction of C_2 , the output spaces of the closed loop systems $C_1 \times M_1$ and $C_2 \times M_2$ have equal distribution, it follows that measurable events have equal probability, as stated next.

Proof: [of Theorem 3.9] If $\{h_1(x_1(t)) | t \in [0, N]\} \in A$ and $(x_1(t), x_2(t)) \in \mathcal{R} \forall t \in [0, N]$ then $\{h_2(x_2(t)) | t \in [0, N]\} \in A$.

Let us rewrite the stochastic kernel of the combined transition of C_2 and M_2 for $t = 0$ as¹⁰

$$\mathbb{T}_{C_2 \times M_2}^0(dx_{C_2} \times dx_2) = \mathbb{T}_{C_1}^0(dx_{C_1} | x_{C_1,0}, x_1) \mathbb{W}_{\pi_{x(0)}}(dx_1 | x_2) \delta_{x_2(0)}(dx_2) \pi_{x(0)}(dx_2(0)).$$

Marginalised on $\mathbb{X}_{C_1} \times \mathbb{X}_1 \times \mathbb{X}_2$, this becomes (by definition of $\mathbb{W}_{\pi_{x(0)}}$)

$$\begin{aligned} \mathbb{T}_{C_2 \times M_2}^0(dx_{C_1} \times dx_1 \times dx_2) &= \mathbb{T}_{C_1}^0(dx_{C_1} | x_{C_1,0}, x_1) \mathbb{W}_{\pi_{x(0)}}(dx_1 | x_2) \pi_{x(0)}(dx_2) \\ &= \mathbb{T}_{C_1}^0(dx_{C_1} | x_{C_1,0}, x_1) \mathbb{W}_{\pi_{x(0)}}(dx_2 | x_1) \pi_{x(0)}(dx_1). \end{aligned}$$

Further marginalised on $\mathbb{X}_{C_1} \times \mathbb{X}_1$, this becomes

$$\mathbb{T}_{C_2 \times M_2}^0(dx_{C_1} \times dx_1) = \mathbb{T}_{C_1}^0(dx_{C_1} | x_{C_1,0}, x_1) \pi_{x(0)}(dx_1) = \mathbb{T}_{C_1 \times M_1}^0(dx_{C_1} \times dx_1).$$

For $t \in [1, N]$, the stochastic kernel marginalised on $\mathbb{X}_{C_1} \times \mathbb{X}_1 \times \mathbb{X}_2$ is

$$\begin{aligned} \mathbb{T}_{C_2 \times M_2}^t(dx'_{C_1} \times dx'_1 \times dx'_2) &= \mathbb{T}_{C_2}^t(dx'_{C_1} | x_{C_1}, x'_1) \\ &\quad \mathbb{W}_{\mathbb{T}}(dx'_1 | x'_2, h_{C_1}^t(x_{C_1}), x_2, x_1) \mathbb{T}_2(dx'_2 | x_2, h_{C_2}^t(x_{C_2})) \\ &= \mathbb{T}_{C_1}^t(dx'_{C_1} | x_{C_1}, x'_1) \mathbb{W}_{\mathbb{T}}(dx'_1 \times dx'_2 | h_{C_1}^t(x_{C_1}), x_2, x_1) \end{aligned}$$

and can be further marginalised on $\mathbb{X}_{C_1} \times \mathbb{X}_1$ to obtain $\mathbb{T}_{C_1 \times M_1}^t$. Note that since $\mathbb{W}_{\mathbb{T}}(\mathcal{R} | h_{C_1}^t(x_{C_1}), x_2, x_1) = 1$ for $(x_1, x_2) \in \mathcal{R}$ it holds with probability 1 that $(x_1(t), x_2(t)) \in \mathcal{R}$ for $t \in [0, N]$. Therefore we can deduce that

$$\mathbb{P}_{C_1 \times M_1}(\{y_1(t)\}_{0:N} \in A) = \mathbb{P}_{C_2 \times M_2}(\{y_2(t)\}_{0:N} \in A).$$

¹⁰For brevity a part of the argument of the stochastic kernel has been omitted.

□

To prove Theorem 3.14 and 3.13 we leverage their exact versions (Theorem 3.8 and 3.9). We first show the existence of a refined control strategy in case of approximate simulation relation, c.f. Theorem 3.14. Then we leverage these results to prove Theorem 3.13.

Theorem 3.14 states the following. Let gMDP \mathbf{M}_1 and \mathbf{M}_2 , with $\mathbf{M}_1 \preceq_\varepsilon^\delta \mathbf{M}_2$, and control strategy $\mathbf{C}_1 = (\mathbb{X}_{\mathbf{C}_1}, x_{\mathbf{C}_1 0}, \mathbb{X}_1, \mathbb{T}_{\mathbf{C}_1}^t, h_{\mathbf{C}_1}^t)$ for \mathbf{M}_1 be given. Then for every given recovery control strategy \mathbf{C}_{rec} , a refined control strategy

$$\mathbf{C}_2 = (\mathbb{X}_{\mathbf{C}_2}, x_{\mathbf{C}_2 0}, \mathbb{X}_2, \mathbb{T}_{\mathbf{C}_2}^t, h_{\mathbf{C}_2}^t)$$

can be obtained as an *inhomogenous Markov process* with two discrete modes of operation, {refinement} and {recovery}, based on Algorithm 2. More specifically a possible choice of a refined control strategy is build up as follows

- state space $\mathbb{X}_{\mathbf{C}_2} := \{\mathbb{X}_{\mathbf{C}_1} \times \mathbb{X}_1 \times \mathbb{X}_2 \times \{\text{refine}\}\} \cup \mathbb{X}_{\mathbf{C}_{rec}} \times \{\text{recover}\}$ with elements $x_{\mathbf{C}_2} = (x_{\mathbf{C}_1}, x_1, x_2, \text{refine})$ and $x_{\mathbf{C}_2} = (x_{\mathbf{C}_{rec}}, \text{recover})$;
- initial state $x_{\mathbf{C}_2 0} := (x_{\mathbf{C}_1 0}, 0, 0, \text{refinement})$;
- accepting as control inputs $x_2 \in \mathbb{X}_2$;
- time dependent stochastic kernel $\mathbb{T}_{\mathbf{C}_2}^t$, defined for $t = 0$ as

$$\begin{aligned} \mathbb{T}_{\mathbf{C}_2}^0(dx_{\mathbf{C}_2}^{\text{refine}} | x_{\mathbf{C}_2 0}, x_2(0)) &:= \mathbb{T}_{\mathbf{C}_1}^0(dx_{\mathbf{C}_1} | x_{\mathbf{C}_1 0}, x_1) \mathbf{1}_{\mathcal{R}}(x_1, x_2) \\ &\quad \times \mathbb{W}_{\pi_{x(0)}}(dx_1 | x_2) \delta_{x_2(0)}(dx_2) \\ \mathbb{T}_{\mathbf{C}_2}^0(dx_{\mathbf{C}_2}^{\text{recover}} | x_{\mathbf{C}_2 0}, x_2(0)) &:= \mathbb{T}_{\text{init}, \text{rec}}^0(dx_{\mathbf{C}_{rec}} | x_2) \mathbf{1}_{(\mathbb{X}_1 \times \mathbb{X}_2) \setminus \mathcal{R}}(x_1, x_2) \\ &\quad \times \mathbb{W}_{\pi_{x(0)}}(dx_1 | x_2) \delta_{x_2(0)}(dx_2) \end{aligned}$$

and for $t \in [1, N]$ over the {refine} operating mode

$$\begin{aligned} \mathbb{T}_{\mathbf{C}_2}^t(dx_{\mathbf{C}_2}^{\text{refine}'} | x_{\mathbf{C}_2}^{\text{refine}}(t), x_2(t)) &:= \mathbb{T}_{\mathbf{C}_1}^t(dx_{\mathbf{C}_1}' | x_{\mathbf{C}_1}, x_1') \mathbf{1}_{\mathcal{R}}(x_1', x_2') \\ &\quad \times \mathbb{W}_{\mathbb{T}}(dx_1' | x_2', h_{\mathbf{C}_1}^t(x_{\mathbf{C}_1}), x_2, x_1) \delta_{x_2(t)}(dx_2'); \\ \mathbb{T}_{\mathbf{C}_2}^t(dx_{\mathbf{C}_2}^{\text{recover}'} | x_{\mathbf{C}_2}^{\text{refine}}(t), x_2(t)) &:= \mathbb{T}_{\text{init}, \text{rec}}^t(dx_{\mathbf{C}_{rec}}' | x_2') \mathbf{1}_{(\mathbb{X}_1 \times \mathbb{X}_2) \setminus \mathcal{R}}(x_1', x_2') \\ &\quad \times \mathbb{W}_{\mathbb{T}}(dx_1' | x_2', h_{\mathbf{C}_1}^t(x_{\mathbf{C}_1}), x_2, x_1) \delta_{x_2(t)}(dx_2'); \end{aligned}$$

defined based on a stochastic kernel $\mathbb{T}_{\text{init}, \text{rec}}^t$ $t \in [0, N]$ initiates the recovery strategy on the fly and is contained in the choice of recovery strategy. And for $t \in [1, N]$ for the recover operating mode

$$\mathbb{T}_{\mathbf{C}_2}^t(dx_{\mathbf{C}_2}^{\text{recover}'} | x_{\mathbf{C}_2}^{\text{recover}}(t), x_2(t)) := \mathbb{T}_{\mathbf{C}_{rec}}^t(dx_{\mathbf{C}_{rec}}' | x_{\mathbf{C}_{rec}}(t), x_2(t));$$

- universally measurable output map

$$h_{\mathbf{C}_2}^t(x_{\mathbf{C}_2}) := \begin{cases} \mathcal{U}_v(h_{\mathbf{C}_1}^t(x_{\mathbf{C}_1}), x_1, x_2) & \text{for refine,} \\ h_{\mathbf{C}_{rec}}^t(x_{\mathbf{C}_{rec}}) & \text{for recover.} \end{cases}$$

The refined control strategy is composed of the control strategy C_1 , the recovery strategy C_{rec} , the stochastic kernel $\mathbb{W}_{\mathbb{T}}$, and the interface \mathcal{U}_v . Both the time-dependent stochastic kernels $\mathbb{T}_{C_2}^t$ and the output maps $h_{C_2}^t$, for $t \in [0, N]$, can be shown to be universally measurable, since Borel measurable maps (and kernels) are universally measurable and the latter are closed under composition [Bertsekas and Shreve, 1996, Ch.7].

Now we need to use this control strategy to prove Theorem 3.13.

Proof: [of Theorem 3.13] Given C_{rec} consider an auxiliary recover strategy C_{rec}^* such that it has stochastic kernels over $\mathbb{X}_{C_{rec}} \times \mathbb{X}_1 \times \mathbb{X}_{C_1}$:

$$\mathbb{T}_{C_{rec}^*}^t(dx'_{C_{rec}^*} | x_{C_{rec}^*}(t), x_2(t)) = \mathbb{T}_{C_{rec}}^t(dx'_{C_{rec}} | x_{C_{rec}}(t), x_2(t)) \\ \mathbb{T}_{C_1 \times M_1}^t(dx'_{C_1 \times M_1} | x_{C_1 \times M_1}(t))$$

where $\mathbb{T}_{C_1 \times M_1}^t(dx'_{C_1 \times M_1} | x_{C_1 \times M_1}(t))$ is the stochastic kernel over $\mathbb{X}_{C_1 \times M_1} := \mathbb{X}_1 \times \mathbb{X}_{C_1}$. Due to the independence of this kernel the probability distribution $\mathbb{P}_{C_2^* \times M_2}$ of M_2 controlled by C_2^* is, when marginalised on the canonical sample space $(\mathbb{X}_{C_2} \times \mathbb{X}_{M_2})^{N+1}$, equal to $\mathbb{P}_{C_2 \times M_2}$.

Now using the same arguments as in the proof of Theorem 3.9 we know that for all measurable sets $L \subset \mathbb{Y}^{N+1}$

$$\mathbb{P}_{C_1 \times M_1}(\{h_1(x_1(t))\}_{0:N} \in L) = \mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in L).$$

The probability

$$\mathbb{P}_{C_2^* \times M_2}((x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N]) \geq (1 - \delta)^{N+1}.$$

This can be shown by induction starting from $t = 0$, and by showing that at every time step and for every pair of states the probability of staying in \mathcal{R} is at least $1 - \delta$. Now note that if $\{h_1(x_1(t))\} \in A_{-\varepsilon}$ and $(x_1(t), x_2(t)) \in \mathcal{R}$ for $t \in [0, N]$ then $\{y(t)\}_{0:N} \in A$. As a consequence

$$\mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon} \wedge (x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N]) \\ \leq \mathbb{P}_{C_2^* \times M_2}(\{h_2(x_2(t))\}_{0:N} \in A) = \mathbb{P}_{C_2 \times M_2}(\{h_2(x_2(t))\}_{0:N} \in A).$$

Now using the union bounding argument we also have that

$$\mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon}) - (1 - \delta)^{N+1} \\ \leq \mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon} \wedge (x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N])$$

$$1 - \mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon} \wedge (x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N]) \\ \leq (1 - \mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon})) \\ + (1 - \mathbb{P}_{C_2^* \times M_2}((x_1(t), x_2(t)) \in \mathcal{R} \text{ for } t \in [0, N])) \\ \leq (1 - \mathbb{P}_{C_2^* \times M_2}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon})) + (1 - (1 - \delta)^{N+1}).$$

We have deduced that

$$\mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{h_1(x_1(t))\}_{0:N} \in A_{-\varepsilon}) - (1 - (1 - \delta)^{N+1}) \leq \mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N} \in A).$$

If $\{h_2(x_2(t))\}_{0:N} \in A$ and $(\tilde{x}(t), x(t)) \in \mathcal{R}$ then $\{h_1(x_1(t))\}_{0:N} \in A_\varepsilon$. Thus via similar arguments it can be deduced that

$$\mathbb{P}_{\mathbf{C}_2 \times \mathbf{M}_2}(\{h_2(x_2(t))\}_{0:N} \in A) \leq \mathbb{P}_{\mathbf{C}_1 \times \mathbf{M}_1}(\{h_1(x_1(t))\}_{0:N} \in A_\varepsilon) + (1 - (1 - \delta)^{N+1}).$$

□

3.B.b Proof of transitivity statements

Proof: [of Theorem 3.15 and Corollary 3.16] Since $\mathbf{M}_1 \preceq_{\epsilon_a}^{\delta_a} \mathbf{M}_2$ and $\mathbf{M}_2 \preceq_{\epsilon_b}^{\delta_b} \mathbf{M}_3$ there exist

- relations $\mathcal{R}_{12} \subset \mathbb{X}_1 \times \mathbb{X}_2$ and $\mathcal{R}_{23} \subset \mathbb{X}_2 \times \mathbb{X}_3$ that satisfies the required conditions in Def. 3.11.
- Interface $\mathcal{U}_{v12} : \mathbb{U}_1 \times \mathbb{X}_1 \times \mathbb{X}_2 \rightarrow \mathcal{P}(\mathbb{U}_2, \mathcal{B}(\mathbb{U}_2))$, and $\mathcal{U}_{v23} : \mathbb{U}_2 \times \mathbb{X}_2 \times \mathbb{X}_3 \rightarrow \mathcal{P}(\mathbb{U}_3, \mathcal{B}(\mathbb{U}_3))$,
- and corresponding stochastic kernels $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$.

Define the relation $\mathcal{R}_{13} \subset \mathbb{X}_1 \times \mathbb{X}_3$ as $\mathcal{R}_{13} := \{(x_1, x_3) \in \mathbb{X}_1 \times \mathbb{X}_3 \mid \exists x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}\}$. Then $\forall (x_1, x_3) \in \mathcal{R}_{13}$ there exists a $x_2 \in \mathbb{X}_2 : (x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}$. More specifically define a Borel-measurable function $F : \mathbb{X}_1 \times \mathbb{X}_3 \rightarrow \mathbb{X}_2$ such that $\forall (x_1, x_3) \in \mathcal{R}_{13}$ for the mapping $x_2 = F(x_1, x_3)$ it holds that $(x_1, x_2) \in \mathcal{R}_{12}, (x_2, x_3) \in \mathcal{R}_{23}$.

We have $\forall (x_1, x_3) \in \mathcal{R}_{13}$ and $x_2 = F(x_1, x_3) :$

1. $\mathbf{d}(h_1(x_1(t)), h_3(x_3)) \leq \mathbf{d}(h_1(x_1(t)), h_2(x_2(t))) + \mathbf{d}(h_2(x_2(t)), h_3(x_3)) \leq \epsilon_a + \epsilon_b;$
2. $\forall u_1 \in \mathbb{U}_1 : \mathbb{T}_1(\cdot | x_1, u_1) \bar{\mathcal{R}}_{12, \delta_a} \mathbb{T}_2(\cdot | x_2, \mathcal{U}_{v12}(u_1, x_1, x_2))$ and for all $u_2 \in \mathbb{U}_2 : \mathbb{T}_2(\cdot | x_2, u_2) \bar{\mathcal{R}}_{23, \delta_b} \mathbb{T}_3(\cdot | x_3, \mathcal{U}_{v23}(u_2, x_2, x_3))$ and $\mathbb{W}_{\mathbb{T}23} \in \mathcal{P}(\mathbb{X}_2 \times \mathbb{X}_3, \mathcal{B}(\mathbb{X}_2 \times \mathbb{X}_3))$ lifted with $\mathbb{W}_{\mathbb{T}12}(\cdot | u_1, x_1, x_2)$ and $\mathbb{W}_{\mathbb{T}23}(\cdot | u_2, x_2, x_3)$.

Let us derive the stochastic kernel $\mathbb{W}_{\mathbb{T}13}$ by combining $\mathbb{W}_{\mathbb{T}12}$ and $\mathbb{W}_{\mathbb{T}23}$ and marginalising over \mathbb{X}_2

$$\begin{aligned} \mathbb{W}_{\mathbb{T}13}(dx'_1 \times dx'_3 | u_1, x_1, x_2, x_3) &= \int_{\mathbb{X}_2} \mathbb{W}_{\mathbb{T}23}(dx'_3 | x'_2, \mathcal{U}_v(u_1, x_1, x_2), x_2, x_3) \\ &\quad \times \mathbb{W}_{\mathbb{T}12}(dx'_1 \times dx'_2 | u_1, x_1, x_2). \end{aligned}$$

Composed with the mapping F we get a Borel-measurable stochastic kernel $\mathbb{W}_{\mathbb{T}13}(dx'_1 \times dx'_3 | u_1, x_1, x_3) := \mathbb{W}_{\mathbb{T}13}(dx'_1 \times dx'_3 | x_1, F(x_1, x_3), x_3)$. In the sequel we drop the argument of the stochastic kernel. Note that $\mathbb{T}_2(dx_2 | x_2, \mu_{u,2}) = \mathbb{W}_{\mathbb{T}12}(\mathbb{X}_1 \times dx_2) =$

$\mathbb{W}_{\mathbb{T}23}(dx_2 \times \mathbb{X}_3)$. For lifting we have to proof that $\mathbb{W}_{\mathbb{T}13}(\mathcal{R}_{13}) \geq 1 - \delta_a - \delta_b$ or equivalently that $\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times \mathbb{X}_3 \setminus \mathcal{R}_{13}) \leq \delta_a + \delta_b$, namely

$$\begin{aligned}
\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times \mathbb{X}_3 \setminus \mathcal{R}_{13}) &= \int_{\mathbb{X}_1} \int_{\mathbb{X}_2} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&= \int_{\mathcal{R}_{12}} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&\quad + \int_{\mathbb{X}_1} \int_{\mathbb{X}_2 \setminus \mathcal{R}_{12}(x_1)} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
\text{for all } (x_1, x_2) \in \mathcal{R}_{12} : \mathcal{R}_{23}(x_2) &\subseteq \mathcal{R}_{13}(x_1) \\
&\leq \int_{\mathbb{X}_2} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{23}(x_2)} \int_{\mathcal{R}_{12}^{-1}(x_2)} \mathbb{W}_{\mathbb{T}12}(dx_1 | x_2) \mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3) \\
&\quad + \int_{\mathbb{X}_1} \int_{\mathbb{X}_2 \setminus \mathcal{R}_{12}(x_1)} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{13}(x_1)} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&\leq \int_{\mathbb{X}_2} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{23}(x_2)} \int_{\mathbb{X}_1} \mathbb{W}_{\mathbb{T}12}(dx_1 | x_2) \mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3) \\
&\quad + \int_{\mathbb{X}_1} \int_{\mathbb{X}_2 \setminus \mathcal{R}_{12}(x_1)} \int_{\mathbb{X}_3} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&= \int_{\mathbb{X}_2} \int_{\mathbb{X}_3 \setminus \mathcal{R}_{23}(x_2)} \mathbb{W}_{\mathbb{T}23}(dx_2 \times dx_3) + \int_{\mathbb{X}_1} \int_{\mathbb{X}_2 \setminus \mathcal{R}_{12}(x_1)} \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&\leq \delta_a + \delta_b.
\end{aligned}$$

In addition it has to hold that $\mathbb{W}_{\mathbb{T}13}(X_1 \times \mathbb{X}_3) = \mathbb{T}_1(\cdot | x_1, \mu_{u,1})$, namely

$$\begin{aligned}
\mathbb{W}_{\mathbb{T}13}(X_1 \times \mathbb{X}_3) &= \int_{X_1} \int_{\mathbb{X}_3} \int_{\mathbb{X}_2} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&= \int_{X_1} \int_{\mathbb{X}_2} \int_{\mathbb{X}_3} \mathbb{W}_{\mathbb{T}23}(dx_3 | x_2) \mathbb{W}_{\mathbb{T}12}(dx_1 \times dx_2) \\
&= \mathbb{W}_{\mathbb{T}12}(X_1 \times \mathbb{X}_2) = \mathbb{T}_1(\cdot | x_1, \mu_{u,1}).
\end{aligned}$$

The condition $\mathbb{W}_{\mathbb{T}13}(\mathbb{X}_1 \times X_3) = \mathbb{T}_3(\cdot | x_3, \mu_{u,3})$ can be proven via similar arguments. In conclusion $\mathbb{T}_1(\cdot | x_1, \mu_{u,1}) \bar{\mathcal{R}}_{13, \delta_a + \delta_b} \mathbb{T}_3(\cdot | x_3, \mu_{u,3})$. To complete the proof we can show, using the same arguments as before, that if $\pi_{x(0)_1} \bar{\mathcal{R}}_{12, \delta_a} \pi_{x(0)_2}$ and if $\pi_{x(0)_2} \bar{\mathcal{R}}_{23, \delta_b} \pi_{x(0)_3}$ then $\pi_{x(0)_1} \bar{\mathcal{R}}_{13, \delta_a + \delta_b} \pi_{x(0)_3}$. \square

Old words are reborn with new faces.

Terri Guillemets

4

Output-based controller synthesis: a correct-by-design approach for Gaussian LTI systems

For partially observable LTI systems subject to measurement noise and stochastic disturbances on state transitions we give a methodology to obtain correct-by-design controllers. We show that available correct-by-design state-based controllers can be extended to an output-based controller with quantified loss in accuracy. For this we give a design methodology that hinges on the design of a state-observer and interface, and for which the accuracy bound can be given by matrix inequalities. Additionally this expected loss of accuracy can be computed a-priori and can be taken into account in the preceding state-based controller design. In a case study from smart buildings we evaluate the new output-based correct-by-design controller on a physical system with limited sensor information.

4.1 Introduction: correct-by-design control

4.1.a Introduction

Computer-aided tools for the verification and synthesis of controllers for (cyber-) physical systems are demanded by domains dealing with complex new applic-

ations. The development of such tools supporting formal specification and systems evolving over continuous spaces, obtained by [Tabuada and Pappas, 2006, Tabuada, 2009, Mazo Jr et al., 2010], focusses on the correct-by-design automatic synthesis of controllers. Current results cover deterministic systems including linear and stabilisable non-linear systems for a multitude of temporal logics such as linear temporal logic [Ding et al., 2014, Tabuada, 2009], and signal temporal logic [Farahani et al., 2015]. Generally, these correct-by-design controllers are incompatible with systems for which exact knowledge of the dynamics and full state measurements are not available. Nevertheless when considering physical systems, measurements often do not contain the full state and are disturbed by sensor noise. Additionally the exact state evolutions can usually not be defined fully, instead they are modelled subject to some stochastic disturbances.

In this work, we target the development of automatised correct-by-design synthesis. For this we consider properties expressed over linear-time temporal logics (LTL) that have been shown to be impactful for the design and verification of software and hardware systems [Clarke, 2008]. The properties expressed in LTL have a syntax and semantics, that extend upon the basic logical operations with temporal modalities (*next*, *until*, *always*), and allow for complex and realistic task specifications. We start on the premise that when full state-information is available a correct-by-design controller can be designed semi-automatically that provably satisfies any required specification. The objective is to extend this correct-by-design controllers [Tabuada and Pappas, 2006] for the set of linear time invariant (LTI) models to output-based controllers that employ sensor outputs or partial state measurements. For the prerequisite formal properties, these new control architectures should come with quantitative certificates guaranteeing the required functionality. Further, since dynamics of physical systems are often disturbed in a probabilistic sense and associated sensors are noisy, we require that the new output-based controllers show quantifiable robustness with respect to stochastic disturbances on state transitions and output measurements.

Solutions to classical optimal control problems [Franklin et al., 1990] of models with (noisy) output measurements can be distinguished in direct designs based on the input-output behaviour of the system, and in methods exploiting the separation of estimation and control. The former class includes frequency-domain and robust control methods; alternatively, whenever applicable (as in the optimal linear quadratic Gaussian problem) the separation theorem [Witsenhausen, 1971] allows for the distinct design of an observer estimating the state and of a state-feedback controller, yielding a combined output feedback controller.

Within the deterministic framework, the synthesis of state-based controllers over continuous spaces for properties expressed via LTL relies on decidability resulting from the abstraction of the model to a (stochastic) finite state model. More precisely, the synthesis generally starts with abstracting the dynamics of the model evolving over the continuous space to a discrete state model. The latter then represents the behaviour of the physical system symbolically in such a way that the control problem is decidable and can be solved in a certifiable manner. Using an equivalence relation between the symbolic abstraction and the continuous model, the synthesised controller for the symbolic abstraction can then be automatically

refined to the continuous model.

For non-stochastic hybrid systems with partly observable dynamics, the work in [Mickelin et al., 2014] synthesises correct-by-design controllers by combining a robust interpretation of the temporal logic formulas and a super-stable hybrid observer. Alternatively the work in [Ghaemi and Vecchio, 2014] leverages partial orders to deal with imperfect information.

For the control synthesis of models with stochastic transitions, safety and reachability objectives represent fundamental design specifications. For stochastic hybrid systems, [Lesser and Oishi, 2014] solves the probabilistic reachability problem formulated as a dynamics programming problem for partially observable states via sufficient statistics. Similarly controller synthesis with respect to safety is solved via the design of an observer and the quantification of error bounds in [Lesser and Abate, 2015]. Beyond these basic properties, already some efforts have targeted the synthesis of controllers for stochastic finite state models without state observations. This includes the in PRISM implemented synthesis method of [Giro and Rabe, 2012], which studies the synthesis for partially observable models by searching the space of output-feedback controllers via counter-example-guided refinements. A heuristic algorithm in [Chatterjee et al., 2014] finds controllers satisfying LTL properties almost surely over partially-observable Markov decision processes. In contrast, the work of [Zhang et al., 2005] extends PCTL* to target hidden Markov models and proposes a model checking algorithm.

For fully observable Markov processes with general state spaces, verification and controller synthesis problems are reviewed in [Abate, 2013], and generally tackled over a simplified model that can be formally related to the original one. The accuracy of the approximate relation used for the control refinement then defines the accuracy of the refinement. These relations can be quantified either via metrics defined over the marginals of the conditional kernels [Esmail Zadeh Soudjani and Abate, 2011], or via metrics bounding the distance between the output trajectories [Julius and Pappas, 2009]. Instead [Zamani et al., 2014] defines an approximate (bisimulation) relation based on the expected deviation of noisy trajectories.

In this chapter we will use the definition of approximate similarity relations like those defined [Zamani et al., 2014] to quantify the *expected* deviation of noisy trajectories affected by stochastic disturbances. This will give a quantification of the accuracy loss that is very similar to the performance of LQG control design and the H_2 -performance norm. First, we extend correct-by-design controllers for LTI systems to the output feedback case. Secondly, we show that this approach is also viable for systems subject to stochastic disturbances. For both the stochastic case and the non-stochastic case, we give a comprehensive design method, verifiable via matrix inequalities. In the next chapter, we consider in detail the resulting structured matrix inequalities and we seek a reformulation of the design problem to a convex optimisation defined over linear matrix inequalities. This reformulated problem is solvable in polynomial time and can be used with the dual objective of performance and accuracy loss.

4.1.b Models and control

We intend to synthesise a certifiable output-based controller for a physical system. First, consider an Linear Time Invariant (LTI) model not affected by noise on the transitions or measurements

$$\mathbf{M} : \begin{cases} x(t+1) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) \\ z(t) &= C_z x(t), \end{cases} \quad (4.1)$$

where $x(t) \in \mathbb{R}^n$ is the state, initialised by $x(0) \in \mathbb{X}_0 \subset \mathbb{R}^n$, the control input is $u(t) \in \mathbb{R}^m$, and $y(t) \in \mathbb{R}^p$ is the measured output available for control. A, B, C are real matrices of appropriate dimensions. The signals $z(t) \in \mathbb{R}^q$, mapped from the state space via the linear map $C_z x$, are used to define performance and properties. This is unlike [Zhang et al., 2005], which defines specifications over the signals $y(t)$. In contrast to the measured output $y(t)$, the structure of which is physically specified by the sensors attached to the system, the choice of C_z can be adapted to the design requirements, and include $C_z = C$ and $C_z = I$ as special cases.

State evolutions of the model and sensor data are often affected by disturbances; these disturbances can be modelled stochastically. Therefore it is of interest to consider the more complex case where the physical system \mathbf{M} is disturbed by stochastic noise. We work with random variables, denoted bold faced \mathbf{x} , and refer to their realisations as $x \sim \mathbf{x}$. An additive noise source $w(t) \sim \mathbf{w}(t)$, taking values in \mathbb{R}^d , affects both the state transitions and the sensor measurements. The noise source is independent and identically distributed over time, with zero mean and unit variance. This assumption holds for a typical Gaussian process noise with distribution $w(t) \sim \mathcal{N}(0, I_{d \times d})$. The resulting stochastic model is

$$\mathbf{M} : \begin{cases} x(t+1) &= Ax(t) + Bu(t) + B_w w(t) \\ y(t) &= Cx(t) + D_w w(t) \\ z(t) &= C_z x(t), \end{cases} \quad (4.2)$$

where the matrices B_w, D_w , are again real-valued matrices of appropriate dimensions. If $D_w B_w^T = 0$ then the noise affecting the sensor and the state transitions are independent. To simplify the notation we will generally assume that $D_w B_w^T = 0$, but all results presented can be carried over to the case with a correlation between the noise on the measurements on the transitions, that is $D_w B_w^T \neq 0$. The system is initialised as $x(0) \sim \mathbf{x}(0) := \mathcal{N}(x_0, P_0)$.

At every time instant a control input $u(t)$ has to be selected. An allowable control strategy consists of an algorithmic way of choosing $u(t)$ based on the available information at that moment. More precisely, at every time instant t_1 , $u(t_1)$ can be selected based on the available information about initial state $x(0)$, the measurements $y(t)$ and the inputs from $u(t)$ from $t = 0$ until t_1 . The satisfaction of properties of interest defined over $z(t)$ depends on the chosen strategy. We consider the strategy to be embedded into a controller \mathbf{C} , that is a system with an internal state that updates the state based on the output measurements and applies a control action to the physical system accordingly. It is the behaviour of the

resulting controlled (or equivalently closed-loop) system, denoted M_C that we want to verify.

Let us clarify this for the deterministic model (4.1) first. Controlled by C the model generates a set of possible trajectories $x : \mathbb{T} \rightarrow \mathbb{R}^n$, denoted \mathcal{B}_x . x is a trajectory of the controlled system M_C if it is initialised, that is $x(0) \in \mathbb{X}_0$, and if for every time instant $x(t)$ the control action is chosen based on C . At every time instant $x(t)$ is mapped to the specification space as $z(t) = C_z x(t)$, as such the signal $z : \mathbb{T} \rightarrow \mathbb{R}^q$ is generated. We refer to z as a trace of system, and the collection of all possible traces form the behaviour and is denoted as $\mathcal{B}_z(M_C)$. For the deterministic model we want to verify that this set of behaviours satisfies the required formal properties defined next. Similarly, the stochastic model (4.2) generates traces $z(t)$ when controlled. But now the set of traces is associated to a probability measure. For this the goal is to find a controller such that properties defined over $z(t)$ are satisfied with a high likelihood.

4.1.c Formalising properties

We consider system properties expressed in Linear-time Temporal Logic [Baier and Katoen, 2008] over a finite set of atomic propositions $p_i \in AP, i = 1, \dots, |AP|$. Any LTL formula ψ is built recursively via the syntax

$$\psi ::= \text{true} \mid p_i \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc\psi \mid \psi \text{ U } \psi.$$

Examples of LTL formulae are $p_1 \wedge p_2, p_1 \wedge (\bigcirc p_2)$, and $p_1 \text{ U } (\bigcirc p_2)$ for $p_{1,2} \in AP$. This syntax allows to extend the study to more complex propositional formulae such as disjunction (\vee). Of special interest to us is the k -bounded and unbounded invariance (or safety) operator as $\square^k \psi := \bigwedge_{i=0}^k \bigcirc^i \psi$ and $\square \psi := \neg(\text{true U } \neg\psi)$, respectively.

Let $\pi = \pi(0), \pi(1), \pi(2), \dots \in \Sigma^{\mathbb{N}^+}$ be a string composed of letters from the alphabet $\Sigma := 2^{AP}$, and let $\pi_t = \pi(t), \pi(t+1), \pi(t+2), \dots$ be a subsequence (postfix) of π . The satisfaction relation between π and a property ψ , expressed via LTL, is denoted as $\pi \models \psi$ (or equivalently $\pi_0 \models \psi$). The semantics of the satisfaction relation are defined recursively over π_t and the syntax of the LTL formula ψ as described in Chapter 2.

LTL formulae can be used to define properties of dynamical model, this means that any LTL formula relates to the behaviour of a dynamical model over a given time horizon for $t \geq 0$. Let the output $z(t) \in \mathbb{R}^q$ be labeled by a map $\text{Lab} : \mathbb{R}^q \rightarrow \Sigma$. Generally, letters of Σ defined by sets of atomic propositions in AP are associated to polyhedra over \mathbb{R}^p . A given signal z defines a word π as $\pi := \text{Lab}(z)$. If the signal is generated by a model, then also the word π is generated by this model. If all words generated by an initialised model satisfy a LTL property ψ then the property is satisfied by the model. More precisely we say a set of traces \mathcal{B}_z generated by a model satisfies a property, if for all $z \in \mathcal{B}_z$: $\text{Lab}(z) \models \psi$. Consider a distance measure $\mathbf{d} : \mathbb{R}^q \times \mathbb{R}^q \rightarrow \mathbb{R}^+$, $\mathbf{d}(z_1, z_2)$ and define for a given precision ε the point-wise expansion of \mathcal{B}_z as $\mathcal{B}_z^\varepsilon := \{\tilde{z} \mid \exists z \in \mathcal{B}_z, \forall t \geq 0, \mathbf{d}(\tilde{z}(t), z(t)) \leq \varepsilon\}$. Then

we say that a property ψ is satisfied ε -robustly by \mathcal{B}_z if for all $z \in \mathcal{B}_z^\varepsilon : \text{Lab}(z) \models \psi$. A practical solution to this based on the shrinking of polyhedra associated to the atomic propositions is referenced in [Mickelin et al., 2014].

4.1.d Problem statement: Output-based and correct-by-design control

Is it possible to semi-automatically synthesise a correct-by-design controller for high-level temporal logic properties and for systems subject to stochastic disturbances? More precisely, can we find

- a controller C such that once composed with the model (4.1), denoted M_C , it holds that $M_C \models \psi$, or
- a controller C for the stochastic disturbed M (4.2) such that M_C satisfies ψ with quantified accuracy?

4.2 State-based correct-by-design controller synthesis

4.2.a Symbolic models & simulation relations

For a deterministic system decidability of the controller synthesis problem follows from the abstraction of its model to a finite state model. As such the control problem is first solved symbolically over the finite state system before it is refined back to the continuous state model. We introduce the notion of symbolic models or transition systems as follows.

Definition 4.1 (Transition system [Tabuada, 2009]) *The tuple $TS = (\mathcal{X}, \mathcal{X}_0, \mathcal{A}, \rightarrow, \mathcal{Z}, \mathcal{H})$ defines a transition system for which*

- \mathcal{X} is a (possibly infinite) set of states;
- \mathcal{X}_0 is a (possibly infinite) set of initial states;
- \mathcal{A} is a (possibly infinite) set of actions, with elements $u \in \mathcal{A}$;
- $\rightarrow \subseteq \mathcal{X} \times \mathcal{A} \times \mathcal{X}$ is a transition relation;
- \mathcal{Z} is a (possible infinite) set of observations;
- $\mathcal{H} : \mathcal{X} \rightarrow \mathcal{Z}$ is a map assigning to each $x \in \mathcal{X}$ an observation $\mathcal{H}(x) \in \mathcal{Z}$.

A metric transition system is a transition system endowed with a metric over the observation space \mathcal{Z} .

As mentioned before, we will consider an observation space $\mathcal{Z} := \mathbb{R}^q$ mapped to the alphabet Σ with the labelling map Lab .

Not only a finite abstraction of the model can be interpreted at a transition system. Also the LTI model (4.1) can be described as a transition system characterised by a tuple $(\mathbb{R}^n, \mathbb{X}_0, \mathbb{R}^m, \rightarrow, \mathbb{R}^q, C_z)$, with a state space $x \in \mathbb{R}^n$, a set of initial states $x(0) \in \mathbb{X}_0$, and transitions $\rightarrow := \{x, u, x' \mid x' = Ax + Bu\}$. Additionally, C_z assigns observation $z \in \mathbb{R}^q$ to $x \in \mathbb{R}^n$: $z = C_z x$. Note that again for a given labelling map from \mathcal{Z} to a finite alphabet Σ the labelling of the signal $z(t)$ for $t \geq 0$ defines a word π as $\pi := \text{Lab}(z_0), \text{Lab}(z_1), \text{Lab}(z_2), \dots$. As such the system can be verified over the labelled properties. Note that the given transition system $(\mathbb{R}^n, \mathbb{X}_0, \mathbb{R}^m, \rightarrow, \mathbb{R}^q, C_z)$ has uniquely defined transitions, since for every state-action pair there is a unique state transition. Such a transition system is called *deterministic*. In this case, it is also a non-blocking transition systems, since every state $x \in \mathcal{X}$ is associated to a non empty transition relation.

The verification of models can be attained by abstracting the dynamics of a model to a simplified representation, pairs of such concrete and abstract models can be related as follows.

Definition 4.2 (Simulation relation [Tabuada, 2009])

Let $\text{TS}_a = (\mathcal{X}_a, \mathcal{X}_{a0}, \mathcal{A}_a, \rightarrow_a, \mathcal{Z}_a, \mathcal{H}_a)$ and $\text{TS}_b = (\mathcal{X}_b, \mathcal{X}_{b0}, \mathcal{A}_b, \rightarrow_b, \mathcal{Z}_b, \mathcal{H}_b)$ be transition systems with the same output sets $\mathcal{Z}_a = \mathcal{Z}_b$. A binary relation $\mathcal{R} \subset \mathcal{X}_a \times \mathcal{X}_b$ is said to be a simulation relation from TS_a to TS_b if the following conditions are satisfied:

- **Equivalence:** for every $(x_a, x_b) \in \mathcal{R}$ we have $\mathcal{H}_a(x_a) = \mathcal{H}_b(x_b)$;
- **Invariance:** for every $(x_a, x_b) \in \mathcal{R}$ we have that $x_a \xrightarrow{u_a} x_a'$ in TS_a implies the existence of $x_b \xrightarrow{u_b} x_b'$ in TS_b satisfying $(x_a', x_b') \in \mathcal{R}$.

We say that TS_a is simulated by TS_b , or that TS_b simulates TS_a , denoted as $\text{TS}_a \preceq_S \text{TS}_b$, if there exists a simulation relation from TS_a to TS_b and if

- for every $x_{a0} \in \mathcal{X}_a$, there exists $x_{b0} \in \mathcal{X}_b$ with $(x_{a0}, x_{b0}) \in \mathcal{R}$.

The models TS_a and TS_b are bisimilar, i.e., $\text{TS}_a \sim_B \text{TS}_b$, if there exists relation \mathcal{R} that is a simulation relation from TS_a to TS_b (such that $\text{TS}_a \preceq_S \text{TS}_b$) and for which \mathcal{R}^{-1} is also a simulation relation from TS_b to TS_a and $\text{TS}_b \preceq_S \text{TS}_a$.

The simulation relation over the set of transition system implies a relation over the behaviour of the transition systems [Tabuada, 2009], more precisely if $\text{TS}_a \preceq_S \text{TS}_b$ then $\mathcal{B}_z(\text{TS}_a) \subseteq \mathcal{B}_z(\text{TS}_b)$, and if $\text{TS}_a \sim_B \text{TS}_b$ then $\mathcal{B}_z(\text{TS}_a) = \mathcal{B}_z(\text{TS}_b)$.

Approximate versions of simulation relations allow for a more robust interpretation. Defined for metric transitions systems, they hinge on a relaxation of the output equivalence relation. Instead of an exact equivalence a small error ε quantified over the metric output space is allowed [Tabuada, 2009]. Consider two given metric transition systems with a shared output space \mathcal{Z} and a metric \mathbf{d} then an ε -approximate simulation relation is defined as follows.

Definition 4.3 (Approximate simulation relation [Tabuada, 2009])

Let $\text{TS}_a = (\mathcal{X}_a, \mathcal{X}_{a0}, \mathcal{A}_a, \rightarrow_a, \mathcal{Z}_a, \mathcal{H}_a)$ and $\text{TS}_b = (\mathcal{X}_b, \mathcal{X}_{b0}, \mathcal{A}_b, \rightarrow_b, \mathcal{Z}_b, \mathcal{H}_b)$ be transition systems with the same output space $\mathcal{Z}_a = \mathcal{Z}_b$ with metric \mathbf{d} . For $\varepsilon \in \mathbb{R}^+$, a relation $\mathcal{R} \subset \mathcal{X}_a \times \mathcal{X}_b$ is said to be an ε -approximate simulation relation from \mathcal{X}_a to \mathcal{X}_b if the following conditions are satisfied

- **Accuracy:** for every $(x_a, x_b) \in \mathcal{R}$ we have $\mathbf{d}(\mathcal{H}_a(x_a), \mathcal{H}_b(x_b)) \leq \varepsilon$;
- **Invariance:** for every $(x_a, x_b) \in \mathcal{R}$ we have that $x_a \xrightarrow{u_a} x_a'$ in TS_a implies the existence of $x_b \xrightarrow{u_b} x_b'$ in TS_b satisfying $(x_a', x_b') \in \mathcal{R}$.

We say that TS_a is approximately simulated by TS_b , or that TS_b approximately simulates TS_a , denoted by $\text{TS}_a \preceq_{\mathcal{S}}^{\varepsilon} \text{TS}_b$, if there exists an ε -approximate simulation relation from TS_a to TS_b and if

- for every $x_{a0} \in \mathcal{X}_{a0}$, there exists $x_{b0} \in \mathcal{X}_{b0}$ with $(x_{a0}, x_{b0}) \in \mathcal{R}$.

The models TS_a and TS_b are approximately bisimilar, i.e., $\text{TS}_a \sim_{\mathcal{B}}^{\varepsilon} \text{TS}_b$, iff there exists a relation \mathcal{R} that is an ε -approximate simulation relation from TS_a to TS_b such that $\text{TS}_a \preceq_{\mathcal{S}}^{\varepsilon} \text{TS}_b$ and for which \mathcal{R}^{-1} is an ε -approximate simulation relation from TS_b to TS_a such that $\text{TS}_b \preceq_{\mathcal{S}}^{\varepsilon} \text{TS}_a$.

This implies that if $\text{TS}_a \preceq_{\mathcal{S}}^{\varepsilon} \text{TS}_b$ then $\mathcal{B}_z(\text{TS}_a) \subseteq \mathcal{B}_z^{\varepsilon}(\text{TS}_b)$, and if $\text{TS}_a \sim_{\mathcal{B}}^{\varepsilon} \text{TS}_b$ then $\mathcal{B}_z(\text{TS}_a) \subseteq \mathcal{B}_z^{\varepsilon}(\text{TS}_b)$ and $\mathcal{B}_z(\text{TS}_b) \subseteq \mathcal{B}_z^{\varepsilon}(\text{TS}_a)$.

Thus if ψ is ε -robustly satisfied by TS_b and if $\text{TS}_a \preceq_{\mathcal{S}}^{\varepsilon} \text{TS}_b$, then TS_a satisfies ψ . Additionally if $\text{TS}_a \sim_{\mathcal{B}}^{\varepsilon} \text{TS}_b$ and if TS_b does not satisfy the property ψ , then we know that TS_a does not ε -robustly satisfy the property.

This explains the use of simulation relation with respect to the verification of systems. In practice, the simulation relations can also be used for the control synthesis problem. That is, they can be used to first solve the synthesis problem on a simplified, and possibly finite, abstraction (TS_a), before refinement over a concrete, complex model (TS_b). Normally this type of controller refinement is best introduced via the alternating notion of simulation relation. Since we will mainly consider deterministic models it is possible to use the non-alternating notion of simulation relation [Tabuada, 2009], as it coincides under these circumstances.

If \mathcal{R} is a simulation relation (cf. Def. 4.2) from the state space of TS_a to the state space of TS_b and if both are deterministic transition systems, then for every pair of paths starting in \mathcal{R} it holds that for every sequence of actions for TS_a , there exists a corresponding sequence for TS_b such that the observed behaviour is the same [Girard and Pappas, 2009]. More precisely, for the action deterministic models we consider definition 4.2, which suggests the refinement of a controller for TS_a to TS_b via the invariance condition: for ever choice of u_a , picked by the controller for TS_a , there exists a suitable input u_b . To allow for an actual refinement of the controller we cannot assume that the initialisation of the concrete system can be freely chosen. Instead we need to add the requirement that for the initial states the concrete system are in the simulation relation \mathcal{R} . This condition is characterised as the requirement that

- for every $x_{b0} \in \mathcal{X}_{b0}$, there exists $x_{a0} \in \mathcal{X}_{a0}$ with $(x_{a0}, x_{b0}) \in \mathcal{R}$.

The actual control refinement is formalised by introducing a specific type of interconnection relation [Tabuada, 2009] next. Referred to as *interface function* it maps actions of one model to actions for another model based on the current states. Interface functions originate from the work in [Girard and Pappas, 2009] on hierarchical control design based on (approximate) simulation relations: the construction of a controller over a simplified model is *refined* to a concrete model while maintaining the same guarantees over the controlled behaviour. The definition is as follows.

Definition 4.4 (Interface function) *Let $\text{TS}_a = (\mathcal{X}_a, \mathcal{X}_{a0}, \mathcal{A}_a, \rightarrow_a, \mathcal{Z}_a, \mathcal{H}_a)$ and $\text{TS}_b = (\mathcal{X}_b, \mathcal{X}_{b0}, \mathcal{A}_b, \rightarrow_b, \mathcal{Z}_b, \mathcal{H}_b)$ be deterministic transition systems with the same output sets $\mathcal{Z}_a = \mathcal{Z}_b$. A relation $\mathcal{R} \subset \mathcal{X}_a \times \mathcal{X}_b$ is an ε -approximate simulation relation from \mathcal{X}_a to \mathcal{X}_b , and $\mathcal{F} : \mathcal{A}_a \times \mathcal{X}_a \times \mathcal{X}_b \rightarrow \mathcal{A}_b$ is its related interface, if the following conditions are satisfied:*

- **Accuracy:** *for every $(x_a, x_b) \in \mathcal{R}$, $\mathbf{d}(\mathcal{H}_a(x_a), \mathcal{H}_b(x_b)) \leq \varepsilon$;*
- **Invariance:** *for every $(x_a, x_b) \in \mathcal{R}$ we have that $x_a \xrightarrow{u_a} x_a'$ in TS_a implies $x_b \xrightarrow{u_b} x_b'$ in TS_b with $u_b = \mathcal{F}(u_a, x_a, x_b)$, satisfying $(x_a', x_b') \in \mathcal{R}$.*

If in addition for every $x_{b0} \in \mathcal{X}_{b0}$, there exists $x_{a0} \in \mathcal{X}_{a0}$ with $(x_{a0}, x_{b0}) \in \mathcal{R}$, then the feedback composition of TS_a and TS_b exists and is denoted as $\text{TS}_a \times_{\mathcal{F}} \text{TS}_b$.

In practice Definition 4.4 entails that the dynamics corresponding to the feedback-composed models $\text{TS}_a \times_{\mathcal{F}} \text{TS}_b$ do not differ more than ε . Hence, a controller composed on TS_a can be refined to TS_b via the interface \mathcal{F} , without affecting its closed-loop accuracy more than ε .

4.2.b State-based correct-by-design controller synthesis

Suppose that an LTI model $\bar{\mathbf{M}}$: $\bar{x}(t+1) = A\bar{x}(t) + B\bar{u}(t)$ is given, and that it has a finite-valued observation map that induces a partition over the observation space \mathbb{R}^q . Under assumptions on the controllability of the model, on the linear independence of the columns of its input matrix B , and on the observation map [Tabuada and Pappas, 2006, Tabuada, 2009], the LTI model can be bisimulated by a finite transition system. Alternatively, under less stringent conditions it is possible to synthesise a finite approximate bisimulation of the given model [Tabuada, 2009, Mazo Jr et al., 2010]: further, for every controller synthesised on the finite-state abstraction there exists a refined controller for the original model, with the same closed-loop behaviour. For a more extensive treatment of this subject please refer to [Tabuada and Pappas, 2006, Tabuada, 2009].

Let us make this more concrete. Suppose that given a model \mathbf{M} and a specification ψ , we have obtained a controlled model $\bar{\mathbf{M}}_{\mathbf{C}}$ satisfying ψ (ε -robustly). More precisely, $\bar{\mathbf{M}}_{\mathbf{C}} := \mathbf{C} \times \mathbf{M}$ denotes the state-based composition of model \mathbf{M} with the

correct-by-design controller \mathbf{C} , where \mathbf{C} takes as input the state of \mathbf{M} and returns an action to \mathbf{M} . This controlled model has hybrid states (\bar{x}, q) with $\bar{x} \in \mathbb{R}^n$ and $q \in Q$, where Q is a finite set. Its dynamics are defined as

$$\bar{\mathbf{M}}_{\mathbf{C}} : \begin{cases} \bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}_q(\bar{x}(t)) \\ q(t+1) &= \delta(\bar{x}(t), q(t)). \end{cases} \quad (4.3)$$

For the initialisation of this controlled model, let us remark that the discrete states q follow from the states of a finite transition system, approximately bisimilar to the continuous-state model \mathbf{M} , and from the specification ψ . Hence the initialisation of the discrete state q is dependent on the specification ψ and the initial state $\bar{x}(0)$. More precisely, the initialisation is denoted as $(\bar{x}(0), q(0)) \in \bigcup_{q_0 \in Q_0} (\{q_0\} \times \mathbb{X}_0(q))$. Note that $\bar{u}_q(\bar{x}(t))$ is a function that maps the current state to an action.

4.3 The idea: output-based correct-by-design control

We approach the design of correct-by-design controller for partially observable systems by building further on the idea of separation of estimation and control, where a state-based controller is used on a partially observable system by including a state observer in the control design. We leverage available results for full state information by constructing a controller that automatically embeds both the state estimates and the state-based controller into an output-based correct-by-design controller. Thus the controller for the partially observable system \mathbf{M} is designed as a combination of the following elements *state-based controller synthesis*, based on state-of-the-art correct-by-design controller synthesis given in Section 4.2.b; and *control refinement* to the concrete system via state-estimation and feedback correction of the state-deviations. In Figure 4.1 a schematic representation of the approach is given, the controller is build up from a hierarchical refinement, given bottom up, where

3. The observer-based implementation, depicted in the bottom layer, starts with the original physical system \mathbf{M} either deterministic (4.1) or stochastic (4.2). For the given concrete model an observer $\mathbf{O}(\mathbf{M})$ estimates the state of this system, based on which the control action dictated by the layer above is refined to the physical system. We show the existence of a-priori computable bounds on the accuracy loss incurred by additive disturbances on the state transitions and on measurements.
2. In the secondary layer “Control refinement”, a state-based control action $\bar{u}(t)$ is computed for the noiseless state-observable LTI model $\bar{\mathbf{M}}$. For this the model $\bar{\mathbf{M}}$ is abstracted to a symbolic model evolving over a finite state space. This finite state model is shown to be (approximately) (bi-)similar to $\bar{\mathbf{M}}$, such that any control for the finite state model can be refined to the continuous state model.
1. The top-layer consists of the control of this symbolic model with respect to specifications over a finite alphabet.

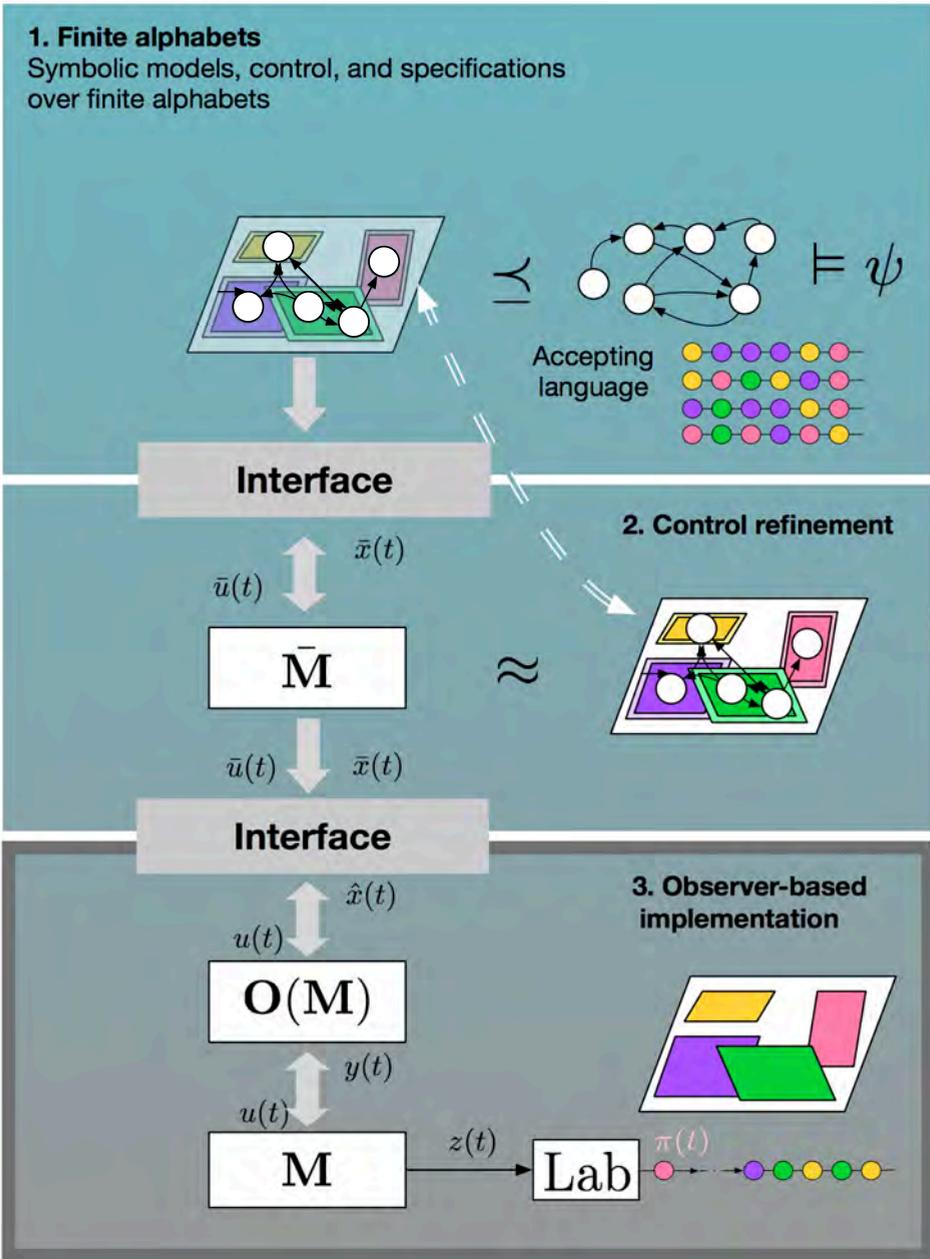


Figure 4.1: The figure gives a schematic representation of the hierarchical control refinement used for correct-by-design control. We contribute to this with a novel third layer.

The first two layers (1. and 2.), given in reverse order, are standard layers within the correct-by-design controller synthesis. A short recap on the ideas behind them has been given in the previous section, i.e., Section 4.2. We newly add the bottom layer in Figure 4.1 in such a way that the design problem really becomes separated into the estimation and control problem. In Section 4.4 we give the methodology for both the noiseless model (4.1) and the noisy model (4.2), and we quantify the accuracy of this implementation layer a-priori. The latter enables us to take into account the expected loss in accuracy when constructing the upper two layers.

4.4 Output-based correct-by-design control

4.4.a Introduction

In this section, we first consider a Linear Time Invariant (LTI) model M not affected by noise on the transitions or measurements as given in (4.1). For this we question whether it is possible to semi-automatically synthesise a correct-by-design controller. Afterwards we consider the case of M given in (4.2) subject to stochastic disturbances on measurements and state transitions.

4.4.b Observer-based control refinement: noiseless case

Consider the model M as given in (4.1) and suppose that there exists a state-based, correct-by-design controller when full state measurements are available, as is the case with closed-loop dynamics denoted by M_C as in (4.3). This control problem is solved in a higher layer of the hierarchical control refinement, now we need to refine the solution of this control problem to the real model. In case we can use full state information the refinement of the controller applied to M_C to the real model M would be very simple, and the error quantification would only be dependent on the difference in initialisation between M_C and M . Since we cannot use full state information, we first define a specific class of allowable interfaces for the control refinement denoted as *sensor-based interfaces*.

Definition 4.5 Let \mathcal{F}_g be an interface function according to Definition 4.4 for TS_a and TS_b , then it is a sensor-based interface if control actions applied to TS_a are refined to TS_b exclusively based on sensor information of TS_b , namely $\mathcal{F}_g : \mathcal{A}_a \times \mathcal{X}_a \times g(\mathcal{X}_b) \rightarrow \mathcal{A}_b$, where g is the sensor function. And if in addition for every $g(x_{b0})$, with $x_{b0} \in \mathcal{X}_{b0}$, there exists $x_{a0} \in \mathcal{X}_{a0}$ with $(x_{a0}, x_{b0}) \in \mathcal{R}$, then the feedback composition of TS_a and TS_b exists and is denoted as $TS_a \times_{\mathcal{F}_g} TS_b$.

In the particular instance of (4.1), a choice of sensor function is $g(x(t)) := Cx(t)$. These structures are of interest to us, as they define the set of interfaces that can be practically implemented for controller refinement on partially observable systems.

Instead of taking the obvious choice of sensor function $Cx(t)$ we design an observer that extends this sensor output with state estimates as in Fig 4.2. Consider

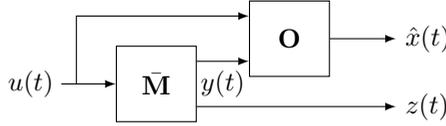


Figure 4.2: Interconnection model/observer, $\bar{\mathbf{M}}\|\mathbf{O}(\bar{\mathbf{M}})$

a Luenberger observer denoted as $\mathbf{O}(\mathbf{M})$:

$$\begin{aligned}\hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + L(\hat{y}(t) - y(t)), \\ \hat{y}(t) &= C\hat{x}(t),\end{aligned}\quad (4.4)$$

with gain matrix L such that $A + LC$ is stable if (A, C) is detectable [Franklin et al., 1990]. The observer is initialised as $\hat{x}(0)$, and uses the outputs from \mathbf{M} to estimate its internal state. The composition of \mathbf{M} with its observer $\mathbf{O}(\mathbf{M})$ is denoted as $\mathbf{M}\|\mathbf{O}(\mathbf{M})$ and portrayed in Figure 4.2. We define a linear, sensor-based interface function between $\bar{\mathbf{M}}_C$ (the state-based, correct-by-design controlled model) and the model/observer interconnection from Figure 4.2 as

$$\mathcal{F}_g(\bar{u}, \bar{x}, \hat{x}) = \bar{u} + F(\hat{x} - \bar{x}), \quad (4.5)$$

where \bar{u} is the action selected by $\bar{\mathbf{M}}_C$ (this role is played by \bar{u}_q in (4.3)). For this linear interface we demand that matrix $A + BF$ is stable. Note that the interface is sensor-based (as defined in Section 4.2), since the state estimate \hat{x} of x can be obtained from the sensor function of $\mathbf{M}\|\mathbf{O}(\mathbf{M})$, thus $g(x, \hat{x}) = \hat{x}$.

The overall controlled model $\bar{\mathbf{M}}_C \times_{\mathcal{F}_g} (\mathbf{M}\|\mathbf{O}(\mathbf{M}))$, denoted as \mathbf{M}_C , is the result of interfacing the two structures discussed above, as depicted in Figure 4.3. This has dynamics evolving over the continuous state space \mathbb{R}^{3n} as:

$$\begin{aligned}\bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}_q(\bar{x}(t)) \\ \hat{x}(t+1) &= (A + LC)\hat{x}(t) + Bu(t) - LCx(t) \\ x(t+1) &= Ax(t) + Bu(t) \\ u(t) &= \mathcal{F}_g(\bar{u}_q(\bar{x}(t)), \bar{x}(t), \hat{x}(t))\end{aligned}\quad (4.6)$$

in combination with the discrete transitions $q(t+1) = \delta(\bar{x}(t), q(t))$ from (4.3).

Remark 4.1 *As depicted in Figure 4.3, we have designed an output-based controller by combining a given state-based controller with an observer. However, unlike classical results where a state-based controller is employed over estimated states from an observer, in this work we have interfaced the state-based controlled model $\bar{\mathbf{M}}_C$ with the model/observer interconnection $\mathbf{M}\|\mathbf{O}(\mathbf{M})$, as in Figure 4.2. This allows one to reason explicitly about the accuracy of the overall output-controlled system, based on the accuracy of the sensor-based interface function. In special cases the proposed architecture is equivalent and can hence be reduced to the classical structure.*

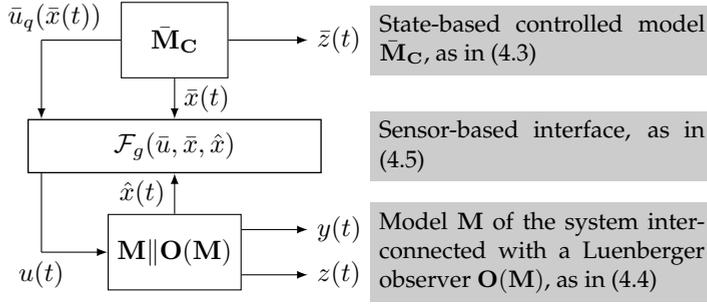


Figure 4.3: Observer-based correct-by-design controller synthesis. The overall interconnection is denoted as \bar{M}_C .

4.4.c Quantification of the accuracy loss: noiseless case

A loss in accuracy is induced by not knowing exactly the state of M ; while the state-based controlled model will exactly satisfy the property it is synthesised for, the observer-based extension introduced in the previous subsection will not have the same output by construction and could hence falsify the property of interest. In this section we will bound the difference between the output of the state-based controlled system and the observer-based controlled system. For this we use the notions of ε -approximate simulation relations and sensor based interface.

The controlled model \bar{M}_C , with trajectories $\bar{x}(t)$ as in (4.3)-(4.6), maps to the specification space as $\bar{z}(t) = C_z \bar{x}(t)$. Let a metric over this space \mathbb{R}^q be defined as $\|\cdot\|_2$. Of interest is the distance between the system output $z(t)$ as in (4.1) and $\bar{z}(t)$, when the system is controlled via the interconnection of Figure 4.3.

Consider the transformed dynamics of (4.6) given as

$$\begin{bmatrix} \bar{x}(t+1) \\ \hat{x}(t+1) - \bar{x}(t+1) \\ x(t+1) - \hat{x}(t+1) \end{bmatrix} = \begin{bmatrix} A & 0 & 0 \\ 0 & (A+BF) & -LC \\ 0 & 0 & (A+LC) \end{bmatrix} \begin{bmatrix} \bar{x}(t) \\ \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix} + \begin{bmatrix} B \\ 0 \\ 0 \end{bmatrix} \bar{u}(t)$$

$$z(t) - \bar{z}(t) = \begin{bmatrix} 0 & C_z & C_z \end{bmatrix} \begin{bmatrix} \bar{x}(t) \\ \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix}.$$

Notice that the error dynamics $z(t) - \bar{z}(t)$ only depend on the difference states $\hat{x}(t) - \bar{x}(t)$ and $x(t) - \hat{x}(t)$ and not on $\bar{x}(t)$ or on the input $\bar{u}(t)$. For ease of notation we introduce the following notation for the evaluation of the error dynamics

$$\begin{bmatrix} \hat{x}(t+1) - \bar{x}(t+1) \\ x(t+1) - \hat{x}(t+1) \end{bmatrix} = \mathcal{A} \begin{bmatrix} \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix}, \quad z(t) - \bar{z}(t) = \mathcal{C} \begin{bmatrix} \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix} (t)$$

with \mathcal{A} and \mathcal{C} defined as

$$\mathcal{A} := \begin{bmatrix} (A+BF) & -LC \\ 0 & (A+LC) \end{bmatrix}, \quad \mathcal{C} := [C_z \quad C_z]. \quad (4.7)$$

Now we can show that the relation

$$\mathcal{R} := \left\{ ((q, \bar{x}), (\hat{x}, x)) \mid \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix} \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix}^T \preceq \mathfrak{X} \right\} \quad (4.8)$$

and (4.5) define, respectively, a simulation relation and (sensor-based) interface function for the models $\bar{\mathbf{M}}_{\mathbf{C}}$ and $\mathbf{M} \parallel \mathbf{O}(\mathbf{M})$. More precisely, we can state the following.

Theorem 4.6 *The function in (4.5) and the relation (4.8) define a sensor-based interface and approximate simulation relation between $\bar{\mathbf{M}}_{\mathbf{C}}$ and $\mathbf{M} \parallel \mathbf{O}(\mathbf{M})$, as given in (4.1),(4.3), (4.4), with precision ε , if \mathfrak{X} is a symmetric positive matrix such that*

$$\varepsilon \geq \sqrt{\text{tr}(\mathcal{C}\mathfrak{X}\mathcal{C}^T)} \quad (4.9)$$

$$\text{s.t.} \quad \begin{bmatrix} \hat{x}(0) - \bar{x}(0) \\ x(0) - \hat{x}(0) \end{bmatrix} \begin{bmatrix} \hat{x}(0) - \bar{x}(0) \\ x(0) - \hat{x}(0) \end{bmatrix}^T - \mathfrak{X} \prec 0 \quad \forall x(0) \in \mathbb{X}_0, \quad (4.10)$$

$$\text{and} \quad \mathcal{A}\mathfrak{X}\mathcal{A}^T - \mathfrak{X} \prec 0. \quad (4.11)$$

Thus the distance between $\bar{z}(t)$ and $z(t)$ is bounded by ε if there exists a \mathfrak{X} for which (4.10) and (4.11) are satisfied. Note that the inequalities (4.9)-(4.11) could also be given in their non-strict form as we did in [Haesaert et al., 2015a], the used strict formulation has the additional benefit that it enforces the refinement to be stabilising. A stability assumption on matrices $A + BF$ and $A + LC$ guarantees the existence of such \mathfrak{X} [Franklin et al., 1990]. Further note that since both $\bar{x}(0)$ and $\hat{x}(0)$ are included in the design space, it would not make much sense to select $\bar{x}(0) \neq \hat{x}(0)$ for the initialisation. Hence, the accuracy depends on the initial states of the models only via $x(0) - \hat{x}(0)$. The second inequality (4.10) in Theorem 4.6 needs to hold for all $x(0) \in \mathbb{X}_0$ in case the initial state $x(0)$ is only known up to a set \mathbb{X}_0 . The proof of the theorem is given as follows.

Proof: [Theorem 4.6] To show that the relation (4.8) is an ε -approximate simulation relation and that the function (4.5) is sensor-based interface function between the state space of $\bar{\mathbf{M}}_{\mathbf{C}}$ and $\mathbf{M} \parallel \mathbf{O}(\mathbf{M})$ we need to show that the accuracy condition, the invariance condition and the initialisation condition are satisfied. For the accuracy consider the output deviation for a given combination of states

$$z - \bar{z} = C_z(\hat{x} - \bar{x}) + C_z(x - \hat{x})$$

$$\|z - \bar{z}\|_2^2 = \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix}^T \mathcal{C}^T \mathcal{C} \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix} = \text{tr} \left(\mathcal{C} \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix} \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix}^T \mathcal{C}^T \right).$$

Therefore if $(\bar{x}, \hat{x}, x) \in \mathcal{R}$ then $\|z - \bar{z}\|_2^2 \leq \text{tr}(\mathcal{C}\mathfrak{X}\mathcal{C}^T) \leq \varepsilon^2$ and $\|z - \bar{z}\|_2 \leq \varepsilon$. Further, to show invariance, suppose that $(\bar{x}(t), \hat{x}(t), x(t)) \in \mathcal{R}$ then

$$\begin{bmatrix} \hat{x}(t+1) - \bar{x}(t+1) \\ x(t+1) - \hat{x}(t+1) \end{bmatrix} = \begin{bmatrix} A + BF & -LC \\ 0 & A + LC \end{bmatrix} \begin{bmatrix} \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix} = \mathcal{A} \begin{bmatrix} \hat{x}(t) - \bar{x}(t) \\ x(t) - \hat{x}(t) \end{bmatrix}$$

thus $(\bar{x}(t+1), \hat{x}(t+1), x(t+1)) \in \mathcal{R}$ if $\mathcal{A}\bar{x}\mathcal{A}^T \preceq \bar{x}$, which holds due to (4.11). Since \mathcal{A} implicitly includes the control refinements given by (4.5), we now have shown that \mathcal{R} is an approximate simulation relation for which the inputs can be refined with the function (4.5). Furthermore, given that both systems evolve deterministically, we can show that (4.5) is an interface if for all initialisations of $\mathbf{M} \parallel \mathbf{O}(\mathbf{M})$ the initialisation of $\bar{\mathbf{M}}_{\mathbf{C}}$ satisfies the relation. This condition is represented by inequality (4.10). \square

As a result of Theorem 4.6, if $\bar{\mathbf{M}}_{\mathbf{C}}$ ε -robustly satisfies ψ , then the sensor-based interface (4.5) is such that composed model $\bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_g} (\mathbf{M} \parallel \mathbf{O}(\mathbf{M}))$ satisfies ψ .

4.4.d Observer-based control refinement: stochastic disturbances

We now consider the case with stochastic disturbances, that is, the systems are given as (4.2). With reference to the previous section, the control design strategy is as follows:

1. Let $\bar{\mathbf{M}}$ be a noiseless version of \mathbf{M} in (4.2), and $\bar{\mathbf{M}}_{\mathbf{C}}$ be the composition of $\bar{\mathbf{M}}$ with its correct-by-design controller as in (4.3);
2. Design a state observer $\mathbf{O}(\mathbf{M})$ for \mathbf{M} , cf. (4.4);
3. Design a linear interface function \mathcal{F}_g stabilising $A + BF$;
4. Implement the control structure in Figure 4.3, and denote the resulting controlled stochastic model as $\mathbf{M}_{\mathbf{C}} := \bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_g} (\mathbf{M} \parallel \mathbf{O}(\mathbf{M}))$.

The initial conditions for $\mathbf{M}_{\mathbf{C}}$, namely $\bar{x}(0), \hat{x}(0)$, are selected as part of the control design problem: as discussed earlier, we pick $\bar{x}(0) = \hat{x}(0)$. Further, let $q(0)$ be any discrete state such that $(\bar{x}(0), q(0)) \in \bigcup_{q_0 \in Q_0} (\{q_0\} \times \mathbb{X}_0(q))$.

In order to analyse the behaviour of the controlled stochastic model $\mathbf{M}_{\mathbf{C}}$ with respect to a metric of interest, let us embed $\mathbf{M}_{\mathbf{C}}$ into the formalism of deterministic transition systems (cf. Definition 4.1) as in [Zamani et al., 2014]. The model can be represented as a symbolic transition system $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$, with states encompassing random variables $\mathbf{x}_{\mathbf{C}}(t)$ representing the distribution of $x_{\mathbf{C}}(t) \sim \mathbf{x}_{\mathbf{C}}(t)$, with $x_{\mathbf{C}}(t) \in \mathbb{R}^{3n}$ as in (4.6). Consider the metric output space \mathcal{Z} , to which the states are mapped as $\mathbf{z}_{\mathbf{C}}(t) = C_{\mathbf{z}}\mathbf{x}_{\mathbf{C}}(t)$. Further consider the metric $\mathbf{d}^*(\mathbf{z}_1, \mathbf{z}_2) = \mathbb{E}(\|\mathbf{z}_1 - \mathbf{z}_2\|_2)$, with $\|\cdot\|_2$ the Euclidean norm. Denote the set of all transition systems with the metric output space \mathcal{Z} as \mathcal{T}^* . The correct-by-design controlled model $\bar{\mathbf{M}}_{\mathbf{C}}$ can be trivially embedded in \mathcal{T}^* via singleton distributions: we denote the corresponding symbolic transition system $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$. We can now show that $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ is approximately bisimilar to $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$.

Theorem 4.7 *Transition system $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ is approximately bisimulated by $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ with precision ε , subject to*

$$\varepsilon \geq \sqrt{\text{tr}(\mathcal{C}\bar{\mathcal{X}}\mathcal{C}^T)} \quad (4.12)$$

$$\text{s.t.} \quad \begin{bmatrix} 0 & 0 \\ 0 & (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & P_0 \end{bmatrix} - \mathfrak{X} \prec 0 \quad (4.13)$$

$$A\mathfrak{X}A^T + B\mathfrak{B}^T - \mathfrak{X} \prec 0 \quad (4.14)$$

with $\mathcal{B} := \begin{bmatrix} -LD_w \\ B_w + LD_w \end{bmatrix}$ and $\mathfrak{X} \succ 0$.

Note that (4.14) is known to admit positive matrices \mathfrak{X} for which ε is finite if $A + BF$ and $A + LC$ are both stable matrices [Franklin et al., 1990].

Proof: [Theorem 4.7] The composition of $\bar{\mathbf{M}}_{\mathbf{C}}$ with $\mathbf{M} \parallel \mathbf{O}(\mathbf{M})$ over the interface (4.5) gives the continuous dynamics of as $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$

$$\begin{aligned} \bar{\mathbf{x}}(t+1) &= A\bar{\mathbf{x}}(t) + B\bar{\mathbf{u}}_q(\bar{\mathbf{x}}(t)) \\ \hat{\mathbf{x}}(t+1) &= (A + LC)\hat{\mathbf{x}}(t) + B\mathbf{u}(t) - LC\mathbf{x}(t) - LD_w\mathbf{w}(t) \\ \mathbf{x}(t+1) &= A\mathbf{x}(t) + B\mathbf{u}(t) + B_w\mathbf{w}(t) \\ \mathbf{u}(t) &= \bar{\mathbf{u}}_q(\bar{\mathbf{x}}(t)) + F(\hat{\mathbf{x}}(t) - \bar{\mathbf{x}}(t)) \end{aligned}$$

with output $\mathbf{z}(t) = C_z\mathbf{x}(t)$. Consider the relation \mathcal{R} , between the state spaces of $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ and $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$, defined as

$$\mathcal{R} := \left\{ ((\mathbf{q}', \bar{\mathbf{x}}'), (\mathbf{q}, \bar{\mathbf{x}}, \hat{\mathbf{x}}, \mathbf{x})) \mid \begin{bmatrix} \bar{\mathbf{x}}' \\ \mathbf{q}' \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{x}} \\ \mathbf{q} \end{bmatrix} \wedge \mathbb{E} \left[\begin{bmatrix} \hat{\mathbf{x}} - \bar{\mathbf{x}} \\ \mathbf{x} - \hat{\mathbf{x}} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}} - \bar{\mathbf{x}} \\ \mathbf{x} - \hat{\mathbf{x}} \end{bmatrix}^T \right] \preceq \mathfrak{X} \right\}$$

where $\bar{\mathbf{x}}'$ is the continuous state of $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ and $\mathbf{x}_{\mathbf{C}} := [\bar{\mathbf{x}}^T \hat{\mathbf{x}}^T \mathbf{x}^T]^T$ the continuous state of $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$. The outputs \mathbf{z}' and \mathbf{z} are similarly defined. Now we want to show that $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ is approximately bisimulated by $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ with precision ϵ with respect to \mathcal{R} . Applying the congruence transform with $\Delta_{\mathbf{x}}(t) = \hat{\mathbf{x}}(t) - \bar{\mathbf{x}}(t)$ and $\mathbf{e}(t) = \mathbf{x}(t) - \hat{\mathbf{x}}(t)$ on the continuous dynamics of $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ gives :

$$\begin{aligned} \bar{\mathbf{x}}(t+1) &= A\bar{\mathbf{x}}(t) + B\bar{\mathbf{u}}_q(\bar{\mathbf{x}}(t)) \\ \Delta_{\mathbf{x}}(t+1) &= (A + BF)\Delta_{\mathbf{x}}(t) - LC\mathbf{e}(t) - LD_w\mathbf{w}(t) \\ \mathbf{e}(t+1) &= (A + LC)\mathbf{e}(t) + (B_w + LD_w)\mathbf{w}(t) \\ \mathbf{z}(t) &= C_z\bar{\mathbf{x}}(t) + C_z\Delta_{\mathbf{x}}(t) + C_z\mathbf{e}(t). \end{aligned}$$

This is needed to show next that \mathcal{R} and \mathcal{R}^{-1} are approximate simulation relations for, respectively, the state spaces of $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ and $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ and for $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ and $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$. For all $((\mathbf{q}', \bar{\mathbf{x}}'), (\mathbf{q}, \bar{\mathbf{x}}, \hat{\mathbf{x}}, \mathbf{x})) \in \mathcal{R}$: the metric $\mathbb{E} [\|\mathbf{z}' - \mathbf{z}\|_2]$ can be written as

$$\begin{aligned} \mathbb{E} [\|C_z\bar{\mathbf{x}}' - C_z\mathbf{x}\|_2] &\leq \sqrt{\mathbb{E} [(C_z\bar{\mathbf{x}}' - C_z\mathbf{x})^T (C_z\bar{\mathbf{x}}' - C_z\mathbf{x})]} \\ &= \sqrt{\text{tr} \mathbb{E} [C_z(\bar{\mathbf{x}}' - \mathbf{x})(\bar{\mathbf{x}}' - \mathbf{x})^T C_z^T]} \end{aligned}$$

The inequality follows from the fact that the square root is a concave function. Note that $(\bar{\mathbf{x}}' - \mathbf{x}) = ((\bar{\mathbf{x}}' - \bar{\mathbf{x}}) - (\hat{\mathbf{x}} - \bar{\mathbf{x}}) - (\mathbf{x} - \hat{\mathbf{x}}))$, and since $\bar{\mathbf{x}}' - \bar{\mathbf{x}} = 0$ (due to

the relation \mathcal{R} , the metric is bounded from above by

$$\text{tr} \left(C \mathbb{E} \begin{bmatrix} \hat{\mathbf{x}} - \bar{\mathbf{x}} \\ \mathbf{x} - \hat{\mathbf{x}} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}} - \bar{\mathbf{x}} \\ \mathbf{x} - \hat{\mathbf{x}} \end{bmatrix}^T C^T \right)^{\frac{1}{2}} \leq \text{tr} (C \mathfrak{X} C^T)^{\frac{1}{2}}.$$

The above holds for both \mathcal{R} and \mathcal{R}^{-1} .

For the next condition we have to prove invariance of \mathcal{R} (and \mathcal{R}^{-1}). More specifically if $((\mathbf{q}'(t), \bar{\mathbf{x}}'(t)), (\mathbf{q}(t), \bar{\mathbf{x}}(t), \hat{\mathbf{x}}(t), \mathbf{x}(t))) \in \mathcal{R}$ then for every transition of $\bar{\mathbf{M}}_{\mathbf{C}}$: $(\mathbf{q}'(t), \bar{\mathbf{x}}'(t)) \rightarrow (\mathbf{q}'(t+1), \bar{\mathbf{x}}'(t+1))$ there exists a transient in $\mathbf{M}_{\mathbf{C}}$ for which $((\mathbf{q}'(t+1), \bar{\mathbf{x}}'(t+1)), (\mathbf{q}(t+1), \bar{\mathbf{x}}(t+1), \hat{\mathbf{x}}(t+1), \mathbf{x}(t+1))) \in \mathcal{R}$. Note that, since both models are controlled, all non-determinism in the control actions have been resolved. Therefore showing that this holds for \mathcal{R} also leads to the conclusion that this holds for \mathcal{R}^{-1} .

Since $\mathbf{M}_{\mathbf{C}}$ is composed of $\bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_g} (\mathbf{M} \parallel \mathbf{O}(\mathbf{M}))$ we know that for every transition $(\mathbf{q}'(t), \bar{\mathbf{x}}'(t)) \rightarrow (\mathbf{q}'(t+1), \bar{\mathbf{x}}'(t+1))$ in $\bar{\mathbf{M}}_{\mathbf{C}}$ there is an equivalent transition $(\mathbf{q}(t), \bar{\mathbf{x}}(t)) \rightarrow (\mathbf{q}(t+1), \bar{\mathbf{x}}(t+1))$ in $\mathbf{M}_{\mathbf{C}}$. Based on the continuous dynamics we also know that the other states evolve as follows

$$\begin{aligned} & \mathbb{E} \left[\begin{bmatrix} \hat{\mathbf{x}}(t+1) - \bar{\mathbf{x}}(t+1) \\ \mathbf{x}(t+1) - \hat{\mathbf{x}}(t+1) \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}(t+1) - \bar{\mathbf{x}}(t+1) \\ \mathbf{x}(t+1) - \hat{\mathbf{x}}(t+1) \end{bmatrix}^T \right] \\ &= \mathcal{A} \mathbb{E} \left[\begin{bmatrix} \hat{\mathbf{x}}(t) - \bar{\mathbf{x}}(t) \\ \mathbf{x}(t) - \hat{\mathbf{x}}(t) \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}(t) - \bar{\mathbf{x}}(t) \\ \mathbf{x}(t) - \hat{\mathbf{x}}(t) \end{bmatrix}^T \right] \mathcal{A}^T + \mathcal{B} \mathcal{B}^T \prec \mathfrak{X}, \end{aligned}$$

where the latter inequality follows from (4.14). Therefore every transition of $\bar{\mathbf{M}}_{\mathbf{C}}$ can be mimicked by $\mathbf{M}_{\mathbf{C}}$. The proof that every transition in $\mathbf{M}_{\mathbf{C}}$ has an equivalent transition in $\bar{\mathbf{M}}_{\mathbf{C}}$ goes along the same lines.

We have shown that \mathcal{R} and \mathcal{R}^{-1} are simulation relations. By considering the initialisation, we can now additionally show that both \mathcal{R} and \mathcal{R}^{-1} define an approximate simulation between the systems. For this observe that the initialisation is deterministic within the frame work of \mathcal{T}^* and : $\bar{\mathbf{x}}'(0) = \bar{\mathbf{x}}(0)$, $\hat{\mathbf{x}}(0) - \bar{\mathbf{x}}(0) = 0$ and $\mathbb{E} \left[(\mathbf{x}(0) - \hat{\mathbf{x}}(0)) (\mathbf{x}(0) - \hat{\mathbf{x}}(0))^T \right] = (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T + P_0$. Therefore based on (4.13) it follows that $\text{TS}^*(\bar{\mathbf{M}}_{\mathbf{C}})$ and $\text{TS}^*(\mathbf{M}_{\mathbf{C}})$ are in an approximate bisimulation relation. \square

4.4.e A first note on the choice of design variables

Thus far we have assumed that L and F are chosen so that they stabilise $A - LC$ and $A + BF$. As long as the model is detectable and stabilisable these gains exist [Franklin et al., 1990]. In general we would like to choose F and L such that ϵ is minimal. A direct implementation of an optimisation of ϵ subject to (4.9)-(4.11) is difficult since the matrix inequalities are nonlinear in F and L . Omitting the initialisation, the computation of the precision level defined in (4.12) together with (4.14) for given L and F is equivalent to $\epsilon = \lim_{t \rightarrow \infty} \sqrt{\mathbb{E} \|\Delta z(t)\|_2^2}$ for a given white

noise sequences $w(t)$. As such the optimisation problem leading to L and F can be recast in the familiar LQG stochastic control problem [Witsenhausen, 1971] for which it is known that the optimal observer gain L and the optimal state-feedback gain F can be computed separately.

Note that since there is no trade-off between the state error and the magnitude of the control gain, the state-feedback gain will push the control to deadbeat control [Franklin et al., 1990]: this behaviour can be easily remedied by extending the observation space $z = C_z x$ with $D_z u$, such that the extended performance signal becomes $z_e(t) = [z^T(t) \quad z_u^T(t)]^T$, with $z_u(t) = D_z u(t)$, or equivalently with $z_u(t) = D_z (u(t) - \bar{u}(t))$. We remark that by defining the refinement with strict inequalities (marginal) stability of both $(A + LC)$ and $(A + BF)$ is guaranteed. The use of non-strict inequalities both on (4.14) and on the \mathfrak{X} in (4.13) would allow for non-stabilising choices of F .

In the next chapter, we will show that there exists a set of separated matrix inequalities replacing (4.14) that enable the computationally constructive design of the correct-by-design control via linear matrix inequalities.

4.5 Case study in smart buildings

We are interested in the advanced energy management of an office building. As a motivation for output-based controllers, consider a building that is divided in two connected zones, each with a radiator regulating the heat in each zone via the controlled boiler water temperature [Holub and Macek, 2013]. Due to a sensor fault in the second zone, only the temperature in the first zone and the ambient (outside) temperatures are measured. The temperature fluctuations in the two zones and the ambient temperature are modelled via M as [Holub and Macek, 2013]

$$x(t+1) = \Xi x(t) + \Gamma u(t) + B_w w(t) \quad (4.15)$$

$$y(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} x(t) + D_w w(t), \quad (4.16)$$

$$z(t) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} x(t),$$

with stable dynamics

$$\Xi = \begin{bmatrix} 0.8725 & 0.0625 & 0.0375 \\ 0.0625 & 0.8775 & 0.0250 \\ 0 & 0 & 0.9900 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 0.0650 & 0 \\ 0 & 0.0600 \\ 0 & 0 \end{bmatrix},$$

where $x_{1,2}(t)$ are the temperatures in zone 1 and 2, respectively; $x_3(t)$ is the deviation of the ambient temperature from its mean; and $u(t) \in \mathbb{R}^2$ is the control input.

Note that since Ξ is stable, it follows that (Ξ, Γ) is stabilisable and $(\Xi, \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix})$ is detectable. The constants in matrix Ξ are selected to represent the heat exchange

rate between the individual zones and the heat loss rate of each zone to the ambient; those in Γ represent the rate of heat supplied by the radiators to the two zones, respectively. The disturbance is modelled as independent and identically distributed standard normal distributions $w(t)$, rescaled by

$$B_w = \begin{bmatrix} .05 & -.02 & 0 & 0 & 0 \\ -.02 & .05 & 0 & 0 & 0 \\ 0 & 0 & 0.1 & 0 & 0 \end{bmatrix} \text{ and } D_w = \begin{bmatrix} 0 & 0 & 0 & .05 & 0 \\ 0 & 0 & 0 & 0 & .05 \end{bmatrix}.$$

The upper block in B_w represents random heat transfers, caused for example by people moving within and between zones, whereas the lower element in the third column represents the stochastic nature of the fluctuation in the outside temperature. The values in D_w define the standard deviation of the additive disturbance on the temperature sensors in the first zone and in the ambient. $y(t)$ is the stochastic signal that can be measured, whereas the specification is defined over $z(t)$ (zone temperatures). Note that the noise acting on the measurements is independent of the noise on the state transitions, because $D_w B_w^T = 0$.

We consider the realistic situation where a sensor failure in zone 2 is discovered. At time $t = 0$ the initial state variables is not known to the control system but can be modelled as $x(0) \sim \mathcal{N}(x_0, P_0)$ with $x_0 = [16 \ 16 \ 0]^T$ and

$$P_0 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & .3 \\ 0 & 0.3 & 2 \end{bmatrix}.$$

The objective is to design an output-based, correct-by-design controller, such that the temperature trajectories $z(t) = (x_1(t), x_2(t))$ eventually both take values in the interval $[20.5, 21]^2$, and remain within this interval thereafter.¹ The controller is initialised with $\bar{x}(0) = \hat{x}(0) = [16 \ 16 \ 0]^T$.

The dynamics of the noiseless model $\bar{\mathbf{M}}$ are solely governed over the first two states, where the correct-by-design controller for the given specification is designed. We synthesise $\bar{\mathbf{M}}_{\mathbf{C}}$ by PESSOA [Mazo Jr et al., 2010], where the discrete-time dynamics are further discretised over state and action spaces: we have selected a state quantisation of .05 over the range $[15, 25]^2$, and an input quantisation of .05 over $[10, 30]^2$. Figure 4.4 displays (continuous blue line) the state trajectory of the obtained correct-by-design system $\bar{\mathbf{M}}_{\mathbf{C}}$: it can be observed that the controller regulates the model to eventually remain within the target region.

Next, we are interested in refining the designed controller to the concrete (noisy) model of the system based on noisy output measurements of the first zone and of the ambient. As a first attempt we implement the controller based on a feedforward architecture, where $\mathcal{F}_{ff} := \bar{u}(t)$. This is what we would obtain applying the results in [Zamani et al., 2014]. It can be observed in Figure 4.4 (circled red realisation) that a trajectory $(x_1(t), x_2(t))$ in $\bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_{ff}} \mathbf{M}$ deviates substantially from the desired temperature range. In Table 4.1 the accuracy of this feedforward interface is given. As a second design, we implement the structure in Figure 4.3, where the

¹This property can be formally expressed as an “eventually always” specification in LTL.

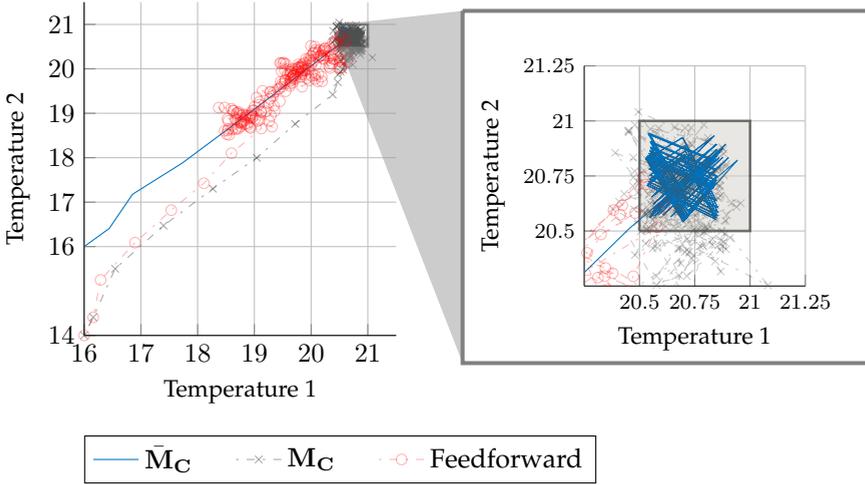


Figure 4.4: Simulation outcomes for controlled models: \bar{M}_C denotes state-based control of the noiseless model realisation [Mazo Jr et al., 2010]; M_C is the output-based control of the Gaussian process model (4.15); Feedforward denotes feedforward design using \bar{M}_C .

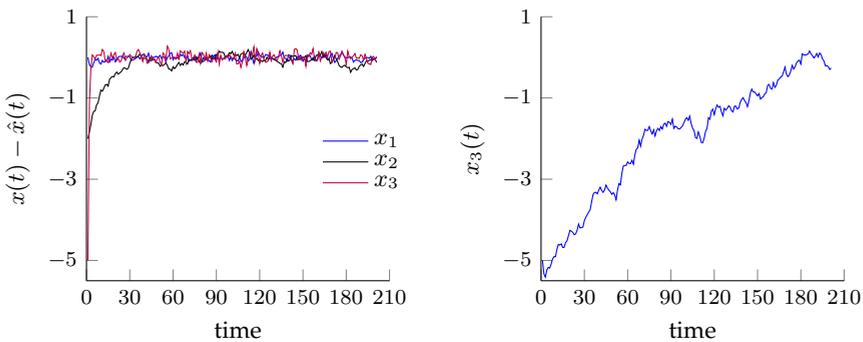


Figure 4.5: (Left plot) Error in state estimation for M_C ; (Right plot) Deviation from mean ambient temperature.

Table 4.1: Error Bounds – Accuracy of the controlled systems based on the interface. With respect to the given initialisation the accuracy is defined by ε_{x_0} and for $t \rightarrow \infty$ the accuracy is given as ε_∞ . The estimates $\hat{\varepsilon}_{x_0,100}$ and ε_∞ are computed as $\sqrt{\hat{\mathbb{E}}_{1:100} \|z(t) - \bar{z}(t)\|_2^2}$ and $\sqrt{\hat{\mathbb{E}}_{10^2:4 \times 10^3} \|z(t) - \bar{z}(t)\|_2^2}$ respectively, with the empirical mean computed as $\mathbb{E}_{i:j} x = \frac{1}{j-i} \sum_{k=i}^j x(k)$.

	ε_{x_0}	ε_∞	$\hat{\varepsilon}_{x_0,100}$	$\hat{\varepsilon}_\infty$
$\bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_{ff}} \mathbf{M}$	3.9618	0.4890	1.9961	0.4845
$\mathbf{M}_{\mathbf{C}}$	2.1194	0.1284	0.5184	0.1240

gains F, L , are selected as detailed Section 4.4.e to be the Kalman gain and the optimal gain F , respectively. The resulting F is equal to the optimal LQ gain. Since we presume in the design that the initialisation is optimal, the observer gain is the Kalman gain. The resulting design values are

$$L = \begin{bmatrix} -0.5201 & -0.0333 \\ 0.2239 & -0.0262 \\ -0.0022 & -0.8196 \end{bmatrix} \text{ and } F = \begin{bmatrix} -13.42 & -0.964 & -0.5759 \\ -1.044 & -14.63 & -0.414 \end{bmatrix}.$$

A trajectory (crossed grey line in Figure 4.4) realised from

$$\mathbf{M}_{\mathbf{C}} := \bar{\mathbf{M}}_{\mathbf{C}} \times_{\mathcal{F}_g} (\mathbf{M} \parallel \mathbf{O}(\mathbf{M}))$$

and based on the previous noise realisation ends up close to the desired temperature range. This substantial improvement with respect to the feedforward interface is also quantified in Table 4.1. Figure 4.5 displays the error of the state estimation $x(t) - \hat{x}(t)$ of $\mathbf{M}_{\mathbf{C}}$ (upper plot): it can be observed that the estimated state converges to a region close to the exact state. The lower plot in Figure 4.5 provides a simulation of the deviation of the ambient temperature from its mean.

4.6 Conclusions

In this work, we have shown that correct-by-design controllers can be extended to work on stochastic partially-observable LTI systems, as long as the LTI system is detectable and stabilisable. We have developed a hierarchical design method that hinges on the design of a state-observer and interface, verified with matrix inequalities. Bounds on the accuracy loss within the hierarchical control refinement have been expressed via the expected 2-norm. The computation of the accuracy loss has been defined based on matrix inequalities. These inequalities have been shown to be similar to those for LQG control problems. The next chapter focusses on this type of matrix inequalities, and develops new theoretical results that enable the constructive synthesis of correct-by-design control.

List of Symbols

General

d Distance measure or metric, $d : \mathbb{Z} \times \mathbb{Z} \rightarrow [0, \infty)$

LTI models

M LTI model as in (4.1) or (4.2).
 \mathbb{Z} Observation space where specifications are defined
 \mathcal{B}_z Set of traces in \mathbb{Z}
 $\mathcal{B}_z^\varepsilon$ Point-wise expansion of \mathcal{B}_z
 Σ Finite alphabet with $\Sigma := 2^{AP}$

Lab Labelling map $\text{Lab} : \mathbb{Z} \rightarrow \Sigma$; induces partitioning of \mathbb{Z}
 g The sensor function, generally $y(t) = Cx(t) + [D_w w(t)]$
 C Controller
 M_C Controlled model
 \bar{M} State-observable and deterministic abstraction of M
 \bar{M}_C Controlled \bar{M}
 $\text{TS}_a \times_{\mathcal{F}} \text{TS}_b$ Feedback composition over interface \mathcal{F}

$\mathbf{O}(M)$ Luenberger observer for M
 L Observer gain

\mathcal{F} Interface function, cf. Definiton 4.4
 \mathcal{F}_g Sensor-based interface function
 F State-feedback gain in interface

Transition systems

TS Transition system
 \mathcal{Z} Observation space for transition system
 \mathbb{R} Relation over Cartesian product of the state spaces of two models
 \preceq_S Preorder over the class of transition systems sharing an observation space. $\text{TS}_a \preceq_S \text{TS}_b$: TS_a is simulated by TS_b
 \sim_B Equivalence relation over the class of transition systems sharing an observation space. $\text{TS}_a \sim_B \text{TS}_b$: TS_a and TS_b are bisimilar
 \preceq_S^ε Preorder over the class of transition systems sharing an observation space. $\text{TS}_a \preceq_S^\varepsilon \text{TS}_b$: TS_a is approximately simulated by TS_b
 \sim_B^ε Equivalence relation over the class of transition systems sharing an observation space. $\text{TS}_a \sim_B^\varepsilon \text{TS}_b$: TS_a and TS_b are approximately bisimilar

Temporal logic

ψ	LTl specifications
π	word (or string) composed of letters of the alphabet Σ
true	True
AP	Set of atomic propositions
p_i	Atomic proposition in AP
\neg	Negation
\wedge	And operator
\vee	Or operator
\bigcirc	Temporal modality for next
U	Temporal modality for until
\square	Temporal modality always, referring to the unbounded invariance (or safety)
\square^k	Temporal modality referring to the k bounded invariance

When asked what it was like to set about proving something, the mathematician likened proving a theorem to seeing the peak of a mountain and trying to climb to the top. One establishes a base camp and begins scaling the mountain's sheer face, encountering obstacles at every turn, often retracing one's steps and struggling every foot of the journey. Finally when the top is reached, one stands examining the peak, taking in the view of the surrounding countryside and then noting the automobile road up the other side!

Robert J. Kleinhenz

5

A separation theorem for guaranteed performance through matrix inequalities

The usage of convex optimisation programs that leverage linear matrix inequalities allows for numerical solutions to the design of an output-feedback controller with respect to H_2 performance. As decreed by the classical separation theorem for the related LQG control problem, the H_2 control problem admits an optimal solution in terms of optimal solutions to the separate optimal state-estimation and state-feedback design problems. This chapter details a new and alternative proof of this separation theorem. Developed for the H_2 output-feedback control problem, the proof is built up using only techniques for (linear) matrix inequalities. Additionally, sub-optimal solutions to the output-feedback control problem are analysed. We show that feasible but sub-optimal solutions of the state-feedback and the state-estimation problem yield a complete solution with guaranteed H_2 performance. Further, we briefly touch upon the implications of these results on the potential reduction in computational complexity and their usage for multi-objective control design. We extend the applicability of these results to output-based hierarchical control refinement. As such we present a constructive approach to the output-based controller refinement.

5.1 Introduction

In control theory, the separation principle usually refers to a controller synthesis methodology in which a state estimator is designed independently from a state-feedback regulator so as to result in a controller that processes measurements to control inputs. A classical example is the linear quadratic Gaussian control problem for a linear time invariant plant in which an optimal controller is obtained as the interconnection of a Kalman filter that minimises the asymptotic covariance of a state-estimation error and the optimal solution of the linear quadratic regulator problem. Or alternatively, consider the design of an output-feedback controller minimising the H_2 -norm of the closed-loop system. In these cases, the separation principle leads to an optimal controlled system and one refers to the *separation theorem* instead of a principle.

There has been a consistent trend to perform computations for optimal control design problems in the realm of convex optimisation programs defined through the feasibility conditions of linear matrix inequalities. Such a phrasing in the context of convex optimisation is well known for the H_2 output-feedback control design problem, but until today it has been restricted to yield a full, that is un-separated, parameterisation of the controller design. In this chapter, we show that the feasibility conditions of the fully parameterised H_2 output-feedback problem can be replaced by separated feasibility conditions related to the estimator and state-feedback designs for a control structure parameterised with the estimator and feedback gains. In particular, we establish separable feasibility conditions related to estimator and state-feedback designs expressible as matrix inequalities. As a result of this, we obtain a new and alternative proof of the classical separation theorem newly phrased and proven within the realm of matrix inequalities and convex optimisation. As a part of this proof, we also show that any pair of suboptimal solutions to the state estimator and feedback problems gives a suboptimal solution to the output-feedback problem.

These results have particular relevance for multi-objective design problems where upper-bounds on achievable or realised performance are explored to meet additional design specifications. Separation of feasibility tests of underlying convex optimisation problems may provide substantial insight and simplify the design process for structured and multi-objective controllers. This is also shown in the secondary part of this chapter. Where we show leverage the separation results for the automatised or algorithmically implementable hierarchical control refinement in the output-based setting introduced in the previous chapter. We tackle the design as a lexicographic optimisation problem. This is a type of multi-objective optimisation that relies on an ordering of objectives. More precisely, we prioritise the accuracy loss, but we allow for a small deviation from its optimal to yield a better performance.

The separation theorem was first proven for continuous-time models in the late 1960 by Wonham [1968]. For discrete-time models the separation theorem can be proved by using dynamic programming; early work on this topic includes the work of Joseph and Tou [1961] and Striebel [1965]. An excellent survey paper that followed the different proofs of the separation theorem was written by Witsen-

hausen [1971]. A more recent proof, given by Davis and Zervos [1995], exploits Lagrange multipliers to prove the separation theorem for quadratic performance objectives in the discrete-time case.

The solution of the H_2 optimal control problem is related to the non-negative solutions of discrete-time algebraic Riccati equations. Analysis of these associated discrete-time algebraic Riccati equations can be found in [Caines and Mayne, 1970b, Ionescu and Weiss, 1993] (with a correction of [Caines and Mayne, 1970b] in [Caines and Mayne, 1970a]), which provides sufficient conditions for minimality, convergence, uniqueness and stability of solutions to the discrete-time matrix Riccati equations.

Structure of the chapter In the next section, we elucidate the problem statement with respect to the H_2 -performance bounds. Based on the background review of optimal observer and state-feedback design in Section 5.3, Section 5.4 contains the main result and proves the separation theorem via matrix inequalities. Section 5.5 discusses a number of corollaries of the main results. Then, in Section 5.6 we give the control refinement problem present in case of hierarchical controller design. As specific cases of this problem the computational problems present in the previous chapter are recovered.

Subsequently, in Section 5.7 results for the control refinement problem are given. Conclusions are summarised in Section 5.8.

5.2 Problem statement for guaranteed H_2 performance

5.2.a The H_2 norm

The H_2 norm of a discrete-time stable system is given by

$$\|T\|_{H_2} = \sqrt{\frac{1}{2\pi} \int_{-\pi}^{\pi} \text{tr} [T(e^{j\omega})T(e^{j\omega})^*] d\omega}$$

where $T(z) = \mathcal{C}(zI - \mathcal{A})^{-1}\mathcal{B}$ is the transfer function and $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are real matrices of appropriate dimension. If this realisation is minimal, then it is well known that

$$\|T\|_{H_2}^2 = \text{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T)$$

where $\mathcal{Y} \succ 0$ is the positive-definite solution of the Lyapunov equation $\mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T = 0$. Upper-bounds on the H_2 norm are characterised as follows.

Proposition 5.1 *The following statements are equivalent:*

1. $\|T\|_{H_2} < \lambda$ and all the eigenvalues of \mathcal{A} are in the open unit disk;
2. there exists a $\mathcal{Y} \succ 0$ such that

$$\mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0, \quad \text{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T) < \lambda^2; \quad (5.1a)$$

3. there exists an $\mathcal{X} \succ 0$ such that

$$\mathcal{X} - \mathcal{A}^T \mathcal{X} \mathcal{A} - \mathcal{C}^T \mathcal{C} \succ 0, \quad \text{tr}(\mathcal{B}^T \mathcal{X} \mathcal{B}) < \lambda^2; \quad (5.1b)$$

4. there exists a \mathcal{Y} and a Z such that

$$\begin{aligned} \begin{bmatrix} \mathcal{Y} & \mathcal{Y} \mathcal{C}^T \\ \mathcal{C} \mathcal{Y} & Z \end{bmatrix} \succ 0, \quad \text{tr} Z < \lambda^2, \\ \mathcal{Y} - \mathcal{A} \mathcal{Y} \mathcal{A}^T - \mathcal{B} \mathcal{B}^T \succ 0; \end{aligned} \quad (5.1c)$$

5. there exists an \mathcal{X} and a Z such that

$$\begin{aligned} \begin{bmatrix} \mathcal{X} & \mathcal{X} \mathcal{B} \\ \mathcal{B}^T \mathcal{X} & Z \end{bmatrix} \succ 0, \quad \text{tr} Z < \lambda^2, \\ \mathcal{X} - \mathcal{A}^T \mathcal{X} \mathcal{A} - \mathcal{C}^T \mathcal{C} \succ 0. \end{aligned} \quad (5.1d)$$

5.2.b System description

Let there be a system \mathbf{M} whose dynamics are given by the mathematical model

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + B_w w(t) \\ y(t) &= Cx(t) + D_w w(t) \\ z(t) &= C_z x(t) + D_z u(t) \end{aligned} \quad (5.2)$$

where $x(t) \in \mathbb{R}^n$ is the state, $u(t) \in \mathbb{R}^m$ is the control input, $y(t) \in \mathbb{R}^p$ is the measured output available for control, and $z(t) \in \mathbb{R}^q$ is the (unmeasured) performance output. A, B, C are real matrices of appropriate dimensions. The system is subject to stochastic disturbances $w(t) \in \mathbb{R}^{w_n}$ on its state transitions and on its measurements, modelled via a white noise sequence with a standard Gaussian distribution, i.e., $\mathcal{N}(0, I_{w_n})$. The signals $z(t) \in \mathbb{R}^q$ are used to define performance. In contrast to the measured output $y(t)$, the structure of which is physically specified by the sensors attached to the system, the choice of C_z and D_z can be adapted to the design requirements.

Throughout this chapter, we will make the following assumptions.

Assumption 5.1

1. (A, B) is stabilisable and (C, A) is detectable;
2. $D_w B_w^T = 0$ and $D_z^T C_z = 0$.

The last condition, targeting both the independence of the noise terms and of the performance indices, is used throughout the chapter for notational brevity. Additionally, the results on optimality are proven via Riccati equations under the following additional hypotheses.

Assumption 5.2

1. $D_z^T D_z \succ 0$ and $D_w D_w^T \succ 0$;
2. (C_z, A) is observable and (A, B_w) is controllable.

5.2.c Controller and controlled system

Consider the set of possible controllers \mathcal{K} for which a given element $\mathbf{K} \in \mathcal{K}$ is defined as

$$\begin{aligned} x_{\mathbf{K}}(t+1) &= A_{\mathbf{K}} x_{\mathbf{K}}(t) + B_{\mathbf{K}} y(t) \\ u_{\mathbf{K}}(t) &= C_{\mathbf{K}} x_{\mathbf{K}}(t). \end{aligned} \quad (5.3)$$

The interconnection of system (5.2) with controller (5.3) is defined by setting $u_{\mathbf{K}} = u$ and yields the controlled system

$$\begin{aligned} \xi(t+1) &= \mathcal{A}\xi(t) + \mathcal{B}w(t) \\ z(t) &= \mathcal{C}\xi(t) \end{aligned} \quad (5.4)$$

with

$$\left[\begin{array}{c|c} \mathcal{A} & \mathcal{B} \\ \hline \mathcal{C} & 0 \end{array} \right] = \left[\begin{array}{cc|c} A & BC_{\mathbf{K}} & B_w \\ B_{\mathbf{K}}C & A_{\mathbf{K}} & B_{\mathbf{K}}D_w \\ \hline C_z & D_z C_{\mathbf{K}} & 0 \end{array} \right]. \quad (5.5)$$

5.2.d Controller design

Consider the objective to choose the control parameters $(A_{\mathbf{K}}, B_{\mathbf{K}}, C_{\mathbf{K}})$ such that the controlled system (5.4) is stable and has minimal H_2 norm:

$$\begin{aligned} \lambda_{\inf} &:= \inf_{\lambda, \mathcal{Y}, \mathbf{K} \in \mathcal{K}} \lambda \\ \text{subject to:} & \quad \sqrt{\text{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T)} < \lambda \\ & \quad \mathcal{Y} \succ 0, \quad \mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0. \end{aligned} \quad (5.6)$$

By Proposition 5.1, the equivalent dual problem is formulated as

$$\begin{aligned} \lambda_{\inf} &= \inf_{\lambda, \mathcal{X}, \mathbf{K} \in \mathcal{K}} \lambda \\ \text{subject to:} & \quad \sqrt{\text{tr}(\mathcal{B}^T \mathcal{X} \mathcal{B})} < \lambda \\ & \quad \mathcal{X} \succ 0, \quad \mathcal{X} - \mathcal{A}^T \mathcal{X} \mathcal{A} - \mathcal{C}^T \mathcal{C} \succ 0. \end{aligned} \quad (5.7)$$

Generally the objective is twofold: one aims to compute the optimal value λ_{\inf} and to find, if it exists, an optimal controller \mathbf{K}^* such that, for each $\lambda > \lambda_{\inf}$, there exists a \mathcal{Y} (respectively \mathcal{X}) for which the inequalities in (5.6) (respectively (5.7)) are satisfied. Viewed as a synthesis problem in the controller parameters $(A_{\mathbf{K}}, B_{\mathbf{K}}, C_{\mathbf{K}})$ for some fixed $\lambda > 0$, the H_2 optimal control problem admits a numerically efficient solution through polynomial time methods [Boyd et al., 1994, Scherer et al., 1997]

by conversion to a convex optimisation problem with constraints given as linear matrix inequalities.

Considering the optimal control problem, we first notice that the formulation of the linear quadratic Gaussian (LQG) problem for linear time-invariant systems reduces to this H_2 control problem, and is solvable via linear matrix inequalities (LMI's). Surprisingly, the separation theorem for the stochastic interpretation of the LQG problem, c.f. [Witsenhausen, 1971], has not yet been deduced using a formulation with matrix inequalities.

Following the separation principle, the optimal solution of the H_2 control problem is the composition of the *Kalman filter* (which represents the mean of the sufficient statistic and the optimal state estimator) and an *optimal static state-feedback controller*. In this work we want to prove that this principle also holds for the separation of matrix inequalities. More precisely, we want to develop a separation theorem on the basis of matrix inequalities in which feasibility of the separated matrix inequalities defines a controller, which yields guaranteed H_2 performance as it is also feasible with respect to the associated matrix inequalities (5.6), or equivalently (5.7). We note that feasible solutions of the separated matrix inequalities define sub-optimal solutions to the state-estimation and state-feedback problems. As such, we relate both feasible (i.e. sub-optimal) solutions and optimal solutions of the separate design problems to the original output-feedback control problem.

Thus, we analyse under which conditions (sub-)optimal solutions of the *state-estimation problem*

$$\inf_{Q \succ 0, L} \operatorname{tr} C_z Q C_z^T \quad (5.8)$$

$$\text{s.t. } Q - (A + LC)Q(A + LC)^T - (B_w + LD_w)(B_w + LD_w)^T \succ 0 \quad (5.9)$$

and of the *state-feedback problem*

$$\inf_{F, P \succ 0} \operatorname{tr} B_w^T P B_w \quad (5.10)$$

$$\text{s.t. } P - (A + BF)^T P (A + BF) - (C_z + D_z F)^T (C_z + D_z F) \succ 0 \quad (5.11)$$

can be used to find a (sub-)optimal solution to the output-feedback problems of (5.6) and (5.7). These two problems will be detailed in the next section.

5.3 Background: optimal control and linear matrix inequalities

5.3.a State-estimation problem

Extracted from system (5.2), we consider the following equations

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + B_w w(t) \\ y(t) &= Cx(t) + D_w w(t). \end{aligned} \quad (5.12)$$

For this system we design a state observer as

$$\begin{aligned}\hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + L(\hat{y}(t) - y(t)), \\ \hat{y}(t) &= C\hat{x}(t).\end{aligned}\tag{5.13}$$

The objective is to find the observer gain L that minimises the asymptotic variance of the estimation error

$$\lim_{t \rightarrow \infty} \mathbb{E}[C_z(x(t) - \hat{x}(t))(x(t) - \hat{x}(t))^T C_z^T]$$

for some given choice of C_z . Recall that the state-error covariance matrix $Q_t := \mathbb{E}[(x(t) - \hat{x}(t))(x(t) - \hat{x}(t))^T]$ then evolves according to

$$\begin{aligned}Q_{t+1} &= (A + LC)Q_t(A + LC)^T \\ &\quad + (B_w + LD_w)(B_w + LD_w)^T.\end{aligned}$$

At each time instance choosing the estimator gain as $L_t = -AQ_t C^T (CQ_t C^T + D_w D_w^T)^{-1}$ yields the smallest variance. Of interest to us is the choice of a constant L leading to the smallest asymptotic variance. The latter is computed as the positive semi-definite solution of the fixed point equation

$$Q = (A + LC)Q(A + LC)^T + (B_w + LD_w)(B_w + LD_w)^T,$$

and it exists for any L for which $(A + LC)$ is stable (Schur). We refer to matrices as Schur when all eigenvalues are in the open unit disk. To find the optimal gain, we consider the positive semi-definite solutions of the corresponding discrete-time algebraic Riccati equation (DTARE)

$$AQA^T - Q + B_w B_w^T - AQC^T(CQC^T + D_w D_w^T)^{-1}CQA^T = 0.\tag{5.14}$$

We say that Q is a solution to (5.14) if $Q = Q^T$, $(CQC^T + D_w D_w^T)$ is invertible and if Q satisfies the equation (5.14). Of interest are the solutions that give (semi-)stabilising dynamics, referred to as strong solutions [Bitmead et al., 1985] and defined next.

Definition 5.2 (Strong solution) *A real symmetric positive semi-definite solution Q^* of the DTARE (5.14) is called a strong solution if the corresponding closed loop matrix*

$$A - L^*C \text{ with } L^* := -AQ^*C^T(CQ^*C^T + D_w D_w^T)^{-1}$$

*has all its eigenvalues in the closed unit disk. A strong solution is called a stabilising if $A - L^*C$ has all its eigenvalues in the open unit disk (i.e. Schur).*

Let Q^* be the strong solution of the DTARE. Then the optimal value of the estimation problem (5.8)-(5.9) is given by $\text{tr}(C_z Q^* C_z^T)$. If Q^* is the stabilising solution of the DTARE, the optimum is actually attained and the optimal observer gain yielding the minimal asymptotic variance is given as $L^* := -AQ^*C^T(CQ^*C^T + D_w D_w^T)^{-1}$, the well known *Kalman gain*.

Recall that the strong solution Q^* exists if (C, A) is detectable and if $D_w D_w^T \succ 0$ and $D_w B_w^T = 0$ hold (cf. Assumptions 5.1- 5.2.1); we emphasise that the last property, i.e. $D_w B_w^T = 0$, just serves to simplify notations all throughout the paper. If, in addition, (A, B_w) is stabilisable, then the strong solution is actually stabilising and can be characterised as the unique positive semi-definite solution of (5.14). In case that (A, B_w) is controllable then Q^* is positive definite. We refer to [Bitmead et al., 1985, Caines and Mayne, 1970b] for details on this Riccati equation. Results on existence of (stabilising) strong solutions, not dependent on $D_w D_w^T \succ 0$ can be found inter alia in [Ionescu and Weiss, 1993, Stoorvogel and Saberi, 1998].

In summary, the relation of Q^* with problem (5.8)-(5.9) can be expressed as follows.

Proposition 5.3 *Under Assumption 5.1- 5.2.1 the DTARE (5.14) admits a strong solution Q^* and*

$$\inf_{Q \succ 0, L \text{ s.t. (5.9)}} \text{tr}(C_z Q C_z^T) = \text{tr}(C_z Q^* C_z^T).$$

Thus, instead of employing the Riccati equation (5.14) the optimal estimator gain can also be found through the feasibility of inequalities in (5.9). For this we can leverage efficient numerical solution techniques based on interior point methods. These polynomial time algorithms [Boyd et al., 1994] allow us to obtain stabilising estimator gains that achieve performance levels arbitrary close to the optimal one.

Further, we introduce the following properties of interest for the derivation in the next section under the standing Assumption 5.1 and for $D_z^T D_z \succ 0$ and $D_w D_w^T \succ 0$.

Proposition 5.4 *Let Assumptions 5.1- 5.2.1 hold. Let Q^* be the strong solution of (5.14). Then the following properties hold*

1. *if Q and L satisfy (5.9), then $Q \succ Q^*$;*
2. *for any $\varepsilon > 0$ there exists Q and L satisfying (5.9) such that $Q^* + \varepsilon I \succ Q$;*
3. *if Q and L satisfy (5.9), then also Q and the updated gain*

$$L^+ := -AQC^T(CQC^T + D_w D_w^T)^{-1}$$

satisfy (5.9).

5.3.b Optimal state-feedback controller

The case that the full state is available for control can be distilled from system (5.2) by taking $y = x$ and results in the equations

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + B_w w(t) \\ z(t) &= C_z x(t) + D_z u(t). \end{aligned}$$

For a linear static state-feedback $u(t) = Fx(t)$ the closed-loop dynamics are

$$\begin{aligned}\xi_s(t+1) &= (A + BF)\xi_s(t) + B_w w(t) \\ z(t) &= (C_z + D_z F)\xi_s(t).\end{aligned}$$

In this setting the objective is to minimise $\lim_{t \rightarrow \infty} \mathbb{E}[z(t)^T z(t)]$. Note that $\lim_{t \rightarrow \infty} \mathbb{E}[z(t)^T z(t)]$ can be written as

$$\text{tr} \left((C_z + D_z F) \lim_{t \rightarrow \infty} \mathbb{E}[\xi_s(t)\xi_s(t)^T] (C_z + D_z F)^T \right).$$

If $(A + BF)$ is stable (Schur, i.e., has all eigenvalues in the open unit disk) then

$$\lim_{t \rightarrow \infty} \mathbb{E}[\xi_s(t)\xi_s(t)^T] = W$$

where $W \succeq 0$ satisfies

$$W = (A + BF)W(A + BF)^T + B_w B_w^T.$$

The state-feedback problem amounts to minimising the output variance $\lim_{t \rightarrow \infty} \mathbb{E}[z(t)^T z(t)]$ and is equal to

$$\begin{aligned}\inf_{W \succ 0, F} \text{tr} \left((C_z + D_z F)W(C_z + D_z F)^T \right) \\ \text{s.t. } W - (A + BF)W(A + BF)^T - B_w B_w^T \succ 0.\end{aligned}\quad (5.15)$$

By Proposition 5.1 this is reformulated to (5.10)-(5.11) as

$$\begin{aligned}\inf_{P \succ 0, F} \text{tr } B_w^T P B_w \\ \text{s.t. } P - (A + BF)^T P (A + BF) - (C_z + D_z F)^T (C_z + D_z F) \succ 0.\end{aligned}$$

From the similarity with the previous subsection, under Assumptions 5.1 and 5.2, this problem admits an optimal gain $F^* = -(B^T P^* B + D_z^T D_z)^{-1} B^T P^* A$ where P^* denotes the stabilising solution of the DTARE

$$\begin{aligned}A^T P A - P + C_z^T C_z \\ - A^T P B (B^T P B + D_z^T D_z)^{-1} B^T P A = 0.\end{aligned}\quad (5.16)$$

In parallel with the state-estimation problem, if (A, B) is stabilisable and $D_z^T D_z \succ 0$ as well as $D_z^T C_z = 0$ (cf. Assumption 5.1-5.2.1) then the strong solution of (5.16) exists. If, in addition, (C_z, A) is detectable then the strong solution is stabilising (i.e. $A + BF^*$ has its eigenvalues in the open unit disk) and can be characterised as the unique positive semi-definite solution of (5.16). Under the conditions of both Assumption 5.1 and 5.2, the stabilising solution P^* is positive definite.

Similar to Proposition 5.3, under the standing Assumption 5.1 and $D_z^T D_z \succ 0$,

(5.16) admits a strong solution P^* , and

$$\inf_{F, P \succ 0 \text{ s.t. (5.11)}} \text{tr}(B_w^T P B_w) = \text{tr}(B_w^T P^* B_w). \quad (5.17)$$

5.3.c Solving the separate problems with linear matrix inequalities

As mentioned before the individual state-estimation and state-feedback problems can be rewritten as convex optimisations subject to linear matrix inequalities. Consider first the *state-feedback problem*. The state-feedback problem, cf. (5.11)-(5.10), can be rewritten as a convex optimisation problem with linear matrix inequalities, that is

$$\begin{aligned} & \inf_{W, R, Z, \gamma} \gamma \\ & \text{subject to} \quad \begin{bmatrix} Z & B_w^T \\ B_w & W \end{bmatrix} \succ 0, \quad \text{tr}(Z) \leq \gamma, \\ & \quad \begin{bmatrix} W & (AW + BR)^T & (C_z W + D_z R)^T \\ (AW + BR) & W & 0 \\ (C_z W + D_z R) & 0 & I \end{bmatrix} \succ 0. \end{aligned} \quad (5.18)$$

The last inequality is derived from (5.11) as follows

$$\begin{aligned} P - \begin{bmatrix} (A + BF)^T & (C_z + D_z F)^T \end{bmatrix} \begin{bmatrix} P & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} (A + BF) \\ (C_z + D_z F) \end{bmatrix} & \succ 0 \\ \begin{bmatrix} P & (A + BF)^T & (C_z + D_z F)^T \\ (A + BF) & P^{-1} & 0 \\ (C_z + D_z F) & 0 & I \end{bmatrix} & \succ 0. \end{aligned}$$

Define $W := P^{-1}$ and $R := FW$ and do a congruence transform with $\text{diag}(\begin{bmatrix} W & I & I \end{bmatrix})$. This gives

$$\begin{bmatrix} W & (AW + BR)^T & (C_z W + D_z R)^T \\ (AW + BR) & W & 0 \\ (C_z W + D_z R) & 0 & I \end{bmatrix} \succ 0,$$

that is, a linear matrix inequality in W and R . We introduce $Z = B_w^T P B_w \succ 0$ and rewrite it as

$$\begin{bmatrix} Z & B_w^T \\ B_w & W \end{bmatrix} \succ 0,$$

to obtain the formulation in (5.18).

In the same way, the *state-estimation problem* can now be rewritten with $M := Q^{-1}$

and $S = M\tilde{L}$ as

$$\inf_{M, S, Z, \gamma} \gamma \quad (5.19)$$

$$\text{subject to } \begin{bmatrix} Z & C_z \\ C_z^T & M \end{bmatrix}, \quad \text{tr}(Z) \leq \gamma \quad (5.20)$$

$$\begin{bmatrix} M & (MA + SC) & (MB_w + SD_w) \\ (MA + SC)^T & M & 0 \\ (MB_w + SD_w) & 0 & I \end{bmatrix} \succ 0. \quad (5.21)$$

The aforementioned problems are feasible under the assumption of detectability of the system matrices (C, A) for (5.18) and the stabilisability of (A, B) for (5.19). Optimal solutions can be found under Assumptions 5.1-5.2.

5.4 Separation theorem via matrix inequalities

5.4.a Introduction

In this section, we derive the separation theorem via matrix inequalities. More precisely, we prove that optimal solutions of the separate problems (5.8)-(5.9) and (5.10)-(5.11) yield, once combined, an optimal solution to the H_2 problem. As an important side result, we relate feasible solutions of the separate problems (5.8) and (5.10) to output-feedback controllers with bounded H_2 performance by showing how solutions of (5.8) and (5.10) define feasible solutions of the H_2 problem (5.6). We start from the standard H_2 -controller design problem in (5.6), that is, a minimisation of λ (quantifying the performance) over the set of allowed controllers \mathbf{K} subject to a set of strict matrix inequalities.

Firstly, we determine a lower bound on the infimum of the controller design problem by loosening the matrix inequalities to *necessary* conditions (cf. Section 5.4.b). The obtained lower bound is a function of the separate control problems in estimation and feedback. The results are built upon a transformation of the matrix inequalities in (5.6) that yields a structurally more beneficial set of constraints.

Secondly in Section 5.4.c, starting again from these new matrix inequalities, we quantify performance λ of controllers composed from feasible solutions of the separated state-estimation and feedback problems.

The combination of the above results allows us to recover optimality (cf. Section 5.4.d), and yields a new separation theorem. Under assumptions on the solution of the corresponding Riccati equations we show that the combination of the optimal solutions of the state-estimation problem (5.8) and the state-feedback problem (5.10), defines the optimal controller \mathbf{K}^* .

5.4.b Lower bound

As a first step to the separation principle via LMIs, we determine a lower bound for the infimum in the H_2 -controller design problem in (5.6).

Theorem 5.5 *Let Q^* and P^* be strong solutions of the DTARE for the filtering problem (5.14) and for the feedback problem (5.16), respectively. Define*

$$\lambda_0^2 := \text{tr}(C_z Q^* C_z) + \text{tr}(A Q^* C^T (C Q^* C^T + D_w D_w^T)^{-1} C Q^* A^T P^*). \quad (5.22)$$

Then for any $\mathbf{K} \in \mathcal{K}$, \mathcal{Y} , $\lambda > 0$ subject to (5.6), i.e.

$$\sqrt{\text{tr}(C \mathcal{Y} C^T)} < \lambda, \quad \mathcal{Y} - A \mathcal{Y} A^T - B B^T \succ 0 \text{ and } \mathcal{Y} \succ 0$$

λ is lower bounded by λ_0 , that is, $\lambda > \lambda_0$.

The above theorem is newly proven via matrix inequalities and uses a transform that will also be crucial in the subsequent results in Section 5.4.d and Section 5.4.c.

Proof: Suppose $\lambda > \lambda_{\text{inf}}$ and let $\mathbf{K} \in \mathcal{K}$ be a controller that achieves (5.6). By expanding the controller realisation $(A_{\mathbf{K}}, B_{\mathbf{K}}, C_{\mathbf{K}})$ with stable uncontrollable or stable unobservable modes we can assume that $\dim(A_{\mathbf{K}}) \geq \dim(A)$. Let $\mathcal{Y} \succ 0$ satisfy (5.6) and partition \mathcal{Y} accordingly with A as

$$\mathcal{Y} = \begin{bmatrix} Y & U \\ U^T & \bar{Y} \end{bmatrix}.$$

Then U is wide and we can assume that U has full row rank (possibly after a small perturbation of \mathcal{Y} , since the inequalities in (5.6) are strict). Define $V := -\bar{Y}^{-1} U^T$ and $X := Y + UV$. Then

$$\mathcal{S} := \begin{bmatrix} I & V^T \\ 0 & -V^T \end{bmatrix} \text{ and } \mathcal{R} := \begin{bmatrix} X & 0 \\ Y - X & U \end{bmatrix}$$

satisfy $\mathcal{S} \mathcal{Y} = \mathcal{R}$ and

$$\mathcal{R} \mathcal{S}^T = \mathcal{S} \mathcal{R}^T = \begin{bmatrix} X & 0 \\ 0 & Y - X \end{bmatrix}.$$

Since V has full column rank, \mathcal{S} will have full row rank. By congruence, this means that $\mathcal{Y} \succ 0$ implies that $\mathcal{S} \mathcal{Y} \mathcal{S}^T = \mathcal{R} \mathcal{S}^T \succ 0$, and thus also that $Y - X \succ 0$.

Part I: Keeping equivalence.

Based on Proposition 5.1 we can expand the matrix inequalities in (5.6), from the form of (5.1a) to (5.1c) with auxiliary Z , subject to $\text{tr} Z < \lambda^2$. Note that $\mathcal{Y} - A \mathcal{Y} A^T - B B^T \succ 0$ can be written as

$$\mathcal{Y} - B B^T - A \mathcal{Y} [\mathcal{Y}^{-1}] (A \mathcal{Y})^T \succ 0$$

and expanded further with a Schur complement. Pre- and post-multiplying the resulting inequalities with the appropriate full rank matrices of \mathcal{S} , specifically

$\text{diag}(\mathcal{S}, I_q)$ and $\text{diag}(\mathcal{S}, \mathcal{S})$ and their respective transposes, gives

$$\begin{bmatrix} \mathcal{R}\mathcal{S}^T & \mathcal{R}\mathcal{C}^T \\ \mathcal{C}\mathcal{R}^T & Z \end{bmatrix} \succ 0, \quad \begin{bmatrix} \mathcal{R}\mathcal{S}^T - \mathcal{S}\mathcal{B}\mathcal{B}^T\mathcal{S}^T & \mathcal{S}\mathcal{A}\mathcal{R}^T \\ \mathcal{R}\mathcal{A}^T\mathcal{S}^T & \mathcal{R}\mathcal{S}^T \end{bmatrix} \succ 0.$$

Written out we obtain, respectively,

$$\begin{bmatrix} X & 0 & XC_z^T \\ 0 & Y-X & (Y-X)C_z^T + UC_{\mathbf{K}}^T D_z^T \\ C_z X & C_z(Y-X) + D_z C_{\mathbf{K}} U^T & Z \end{bmatrix} \succ 0, \quad (5.23)$$

and

$$\begin{bmatrix} X - B_w B_w^T - V^T B_{\mathbf{K}} D_w D_w^T B_{\mathbf{K}}^T V & * & * & * \\ V^T B_{\mathbf{K}} D_w D_w^T B_{\mathbf{K}}^T V & Y - X - V^T B_{\mathbf{K}} D_w D_w^T B_{\mathbf{K}}^T V & * & * \\ X A^T + X C^T B_{\mathbf{K}}^T V & -X C^T B_{\mathbf{K}}^T V & X & * \\ (Y-X)A^T + U A_{\mathbf{K}}^T V + U C_{\mathbf{K}}^T B^T + (Y-X)C^T B_{\mathbf{K}}^T V & -U A_{\mathbf{K}}^T V - (Y-X)C^T B_{\mathbf{K}}^T V & 0 & Y-X \end{bmatrix} \succ 0. \quad (5.24)$$

We now introduce the following abbreviations

$$\begin{aligned} \Delta &:= Y - X, & F &:= C_{\mathbf{K}} U^T \Delta^{-1}, \\ L &:= V^T B_{\mathbf{K}}, & \text{and,} & A_{\delta}^T &:= \Delta^{-1} U A_{\mathbf{K}}^T V - (A + BF + LC)^T. \end{aligned}$$

As such (5.23) and (5.24) are simplified and given, in order, as follows

$$\begin{bmatrix} X & 0 & XC_z^T \\ 0 & \Delta & \Delta C_z^T + \Delta F^T D_z^T \\ C_z X & C_z \Delta + D_z F \Delta & Z \end{bmatrix} \succ 0. \quad (5.25)$$

and

$$\begin{bmatrix} X - B_w B_w^T - L D_w D_w^T L^T & * & * & * \\ L D_w D_w^T L^T & \Delta - L D_w D_w^T L^T & * & * \\ X(A+LC)^T & -X C^T L^T & X & * \\ \Delta A_{\delta}^T & \Delta(A - A_{\delta} + BF)^T & 0 & \Delta \end{bmatrix} \succ 0$$

By taking a Schur complement over the latter inequality, we obtain an $(n+n) \times (n+n)$ matrix inequality,

$$\begin{bmatrix} \text{(block 1)} & \text{(off diag.)} \\ \text{(off diag.)}^T & \text{(block 2)} \end{bmatrix} \succ 0 \quad (5.26)$$

with

$$\begin{aligned} \text{(block 1)} &:= X - (A + LC)X(A + LC)^T - B_w B_w^T - L D_w D_w^T L^T - A_{\delta} \Delta A_{\delta}^T, \\ \text{(off diag.)} &:= -A_{\delta} \Delta (A - A_{\delta} + BF)^T + L D_w D_w^T L^T + (A + LC)X C^T L^T, \\ \text{(block 2)} &:= \Delta - (A + BF - A_{\delta}) \Delta (A - A_{\delta} + BF)^T - L(D_w D_w^T + C X C^T) L^T. \end{aligned}$$

Feasibility of (5.6) for a given $(Z, \mathcal{Y}, A_{\mathbf{K}}, B_{\mathbf{K}}, C_{\mathbf{K}})$ implies that the obtained inequalities (5.25) and (5.26) are feasible. Hence the inequalities derived in this part represent necessary conditions for (5.6). For the solutions with a square invertible U , (5.25) and (5.26) represent both necessary and sufficient conditions for (5.6).

Part II: Extracting necessary conditions from (5.26)

We consider (5.26) and obtain two necessary conditions. The first allows us to argue that $X \succ Q^*$ is guaranteed. This follows by looking at the upper diagonal block element of (5.26), which implies

$$X - (A + LC)X(A + LC)^T - B_w B_w^T - LD_w D_w^T L^T \succ 0. \quad (5.27)$$

Both (5.27) and $X \succ 0$ are necessary conditions of X . Therefore based on Proposition 5.4 it follows that $X \succ Q^*$. Remark that, as mentioned before, $B_w D_w^T = 0$ is assumed for notational brevity.

A second necessary condition with respect to state-feedback is recovered next. We distill from (5.26) a necessary condition as

$$\begin{bmatrix} I \\ I \end{bmatrix}^T \begin{bmatrix} \text{(block 1)} & \text{(off diag.)} \\ \text{(off diag.)}^T & \text{(block 2)} \end{bmatrix} \begin{bmatrix} I \\ I \end{bmatrix} \succ 0, \quad (5.28)$$

and obtain the following inequality

$$\begin{aligned} \Delta + X - A(\Delta + X)A^T - B_w B_w^T \\ - (A\Delta F^T B^T) - (A\Delta F^T B^T)^T - (BF)\Delta(BF)^T \succ 0. \end{aligned}$$

With as goal the replacement of X by Q^* in this inequality, we now re-introduce Y , $C_{\mathbf{K}}$ and U via $F = C_{\mathbf{K}} U^T \Delta^{-1}$ and $\Delta = Y - X$, this gives

$$\begin{aligned} Y - AY A^T - BC_{\mathbf{K}} U^T (Y - X)^{-1} U C_{\mathbf{K}}^T B^T \\ - (AUC_{\mathbf{K}}^T B^T) - (AUC_{\mathbf{K}}^T B^T)^T - B_w B_w^T \succ 0. \end{aligned}$$

Since $Y - X \succ 0$, this is equivalent to

$$\begin{bmatrix} Y - X & \\ UC_{\mathbf{K}}^T B^T & Y - AY A^T - (AUC_{\mathbf{K}}^T B^T)^* - (AUC_{\mathbf{K}}^T B^T)^T - B_w B_w^T \end{bmatrix} \succ 0.$$

We add the following positive semi-definite matrix

$$\begin{bmatrix} X - Q^* & & 0 \\ 0 & AQ^* A^T - Q^* + B_w B_w^T - AQ^* C^T (CQ^* C^T + D_w D_w^T)^{-1} CQ^* A^T \end{bmatrix}$$

and get

$$\begin{bmatrix} Y - Q^* & & * \\ UC_{\mathbf{K}}^T B^T & (Y - Q^*) - A(Y - Q^*)A^T - (AUC_{\mathbf{K}}^T B^T)^* - (AUC_{\mathbf{K}}^T B^T)^T \\ & & -AQ^* C^T (CQ^* C^T + D_w D_w^T)^{-1} CQ^* A^T \end{bmatrix} \succ 0.$$

With the substitutions $\Delta_* := Y - Q^*$ and $F_* := C_{\mathbf{K}}^T U^T \Delta_*^{-1}$ and taking the Schur complement, we obtain the second necessary inequality

$$\Delta_* - (A + BF_*)\Delta_*(A + BF_*)^T - AQ^* C^T (CQ^* C^T + D_w D_w^T)^{-1} CQ^* A^T \succ 0. \quad (5.29)$$

Since $AQ^* C^T (CQ^* C^T + D_w D_w^T)^{-1} CQ^* A^T$ is positive semi-definite, we can ex-

press this matrix as $\tilde{B}_w \tilde{B}_w^T$ and obtain

$$\Delta_* - (A + BF_*)\Delta_*(A + BF_*)^T - \tilde{B}_w \tilde{B}_w^T \succ 0. \quad (5.30)$$

We emphasise the relevant fact that (5.29) assumes the form of the inequality (5.15).

Part III: Extracting necessary conditions from (5.25)

To compute the lower bound λ_0 on the H_2 controller design problem we now also consider the inequality (5.25) related to performance. Again we want to get inequalities for F_* and Δ_* . Therefore we re-introduce Y , C_K and U via $F = C_K U^T \Delta^{-1}$ and $\Delta = Y - X$, which gives

$$\begin{bmatrix} X & & * \\ 0 & Y - X & * \\ C_z X & -C_z(Y - X) - D_z C_K U^T & Z \end{bmatrix} \succ 0.$$

Via subsequent Schur complements, this is equivalent to

$$Z - C_z Y C_z^T - C_z U C_K^T D_z^T - (C_z U C_K^T D_z^T)^T - D_z C_K U^T (Y - X)^{-1} U C_K^T D_z^T \succ 0,$$

and

$$\begin{bmatrix} (Y - X) & U C_K^T D_z^T \\ D_z C_K U^T & Z - C_z Y C_z^T - C_z U C_K^T D_z^T - (C_z U C_K^T D_z^T)^T \end{bmatrix} \succ 0.$$

Since we assured that $X \succ Q^*$ is a necessary condition, we infer

$$\begin{bmatrix} (Y - Q^*) & * \\ D_z C_K U^T & Z - C_z Y C_z^T - C_z U C_K^T D_z^T - (C_z U C_K^T D_z^T)^T \end{bmatrix} \succ 0.$$

Taking the Schur complement and substituting Δ_* and F_* gives

$$Z - (C_z + D_z F_*)\Delta_*(C_z + D_z F_*)^T - C_z Q^* C_z^T \succ 0.$$

Thus $Z \succ (C_z + D_z F_*)\Delta_*(C_z + D_z F_*)^T + C_z Q^* C_z^T$ and since the trace operation (tr) is a monotonic operation with respect to (\succeq) it follows that any feasible λ is bounded from below as

$$\lambda^2 > \text{tr}((C_z + D_z F_*)\Delta_*(C_z + D_z F_*)^T) + \text{tr}(C_z Q^* C_z^T)$$

for the given choices of F_* and Δ_* . Given the necessary constraint (5.29) a lower bound on λ^2 is found by solving the following optimisation problem

$$\begin{aligned} \inf_{\Delta_* \succ 0, F_*} \quad & \text{tr}(C_z + D_z F_*)\Delta(C_z + D_z F_*)^T \\ \text{s.t. (5.29):} \quad & \Delta - (A + BF_*)\Delta_*(A + BF_*)^T \\ & - A Q^* C^T (C Q^* C^T + D_w D_w^T)^{-1} C Q^* A^T \succ 0. \end{aligned}$$

The above infimum is equal to the optimal value of a state feedback controller

design problem with gain F and a disturbance input matrix \tilde{B}_w . As iterated in Section 5.3.b and (5.15), this state-feedback problem can be rewritten into its dual form based on Proposition 5.1

$$\begin{aligned} & \inf_{P \succ 0, F_*} \operatorname{tr} (AQ^*C^T(CQ^*C^T + D_wD_w^T)^{-1}CQ^*A^TP) \\ \text{s.t. } & P - (A + BF_*)^T P (A + BF_*) \\ & \quad - (C_z + D_zF_*)^T (C_z + D_zF_*) \succ 0. \end{aligned}$$

Since P^* is the strong solution of the related DTARE, (5.17) shows that the optimal value is $\operatorname{tr}(\tilde{B}_w^T P^* \tilde{B}_w) = \operatorname{tr}(\tilde{B}_w \tilde{B}_w^T P^*)$. By recalling the definition of $\tilde{B}_w \tilde{B}_w^T$, we end up with

$$\operatorname{tr}(AQ^*C^T(CQ^*C^T + D_wD_w^T)^{-1}CQ^*A^TP^*) = \lambda_0$$

being a lower bound on λ^2 , as was to be shown. \square

5.4.c Sufficiency and controller construction

Theorem 5.6 *Let us be given feasible solutions P, F and Q, \tilde{L} of the inequalities*

$$P \succ 0, \quad P - (A + BF)^T P (A + BF) - (C_z + D_zF)^T (C_z + D_zF) \succ 0, \quad (5.31)$$

$$Q \succ 0, \quad Q - (A + \tilde{L}C)Q(A + \tilde{L}C)^T - (B_w + \tilde{L}D_w)(B_w + \tilde{L}D_w)^T \succ 0. \quad (5.32)$$

Then for any λ subject to

$$\operatorname{tr}(AQC^T(CQC^T + D_wD_w^T)^{-1}CQA^TP) + \operatorname{tr}(C_zQC_z^T) < \lambda^2 \quad (5.33)$$

there exists an \mathcal{Y} such that the matrix inequalities

$$\begin{aligned} & \sqrt{\operatorname{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T)} < \lambda \\ & \mathcal{Y} \succ 0, \quad \mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0 \end{aligned} \quad (5.34)$$

are satisfied with the controller $\mathbf{K} \in \mathcal{K}$ composed as

$$\begin{bmatrix} A_{\mathbf{K}} & B_{\mathbf{K}} \\ C_{\mathbf{K}} & 0 \end{bmatrix} = \begin{bmatrix} A + BF + LC & -L \\ F & 0 \end{bmatrix} \quad (5.35)$$

for the updated gain $L := -AQC^T(CQC^T + D_wD_w^T)^{-1}$.

This theorem is one of the two key novel contributions of this section. It introduces a way to use the sub-optimal solutions to the separate state-estimation and state-feedback problems. More precisely, it shows that feasible but sub-optimal solutions to the separate problems can be combined and that the combined solution will give guarantees on the obtained H_2 -performance. This setting avoids conditions on uniqueness of solutions of Riccati equations. Specifically, only require-

ments on the stabilisability and detectability of, respectively, (A, B) and (C, A) are necessary for Theorem 5.6. Thus for these limited conditions the above theorem provides a complete and explicit synthesis procedure with guaranteed H_2 performance based on the separate designs of the state-feedback and the state-estimator gains. We note that the inequalities (5.31) and (5.32) are not linear in the design variables (P, F) and (Q, \tilde{L}) . In combination with the nonlinear objective (5.33), these inequalities cannot immediately be employed to simultaneously compute F and \tilde{L} via the convenient use of linear matrix inequalities and convex optimisation. The actual computation of the designs of F and \tilde{L} subject to (5.31) and (5.32) with respect to (5.33) is tackled in detail in the second part of this chapter.

Let us note that the update of the state-estimation gain L is chosen such that it decouples the effect of the state-estimation and the state-feedback. To observe this consider first that the proof is built on the same transform as in Theorem 5.5. This yields a block-diagonal matrix which is For the chosen controller structure, there exists a \mathcal{Y} such that the inequality $\mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0$ once transformed into (5.26) the off diagonal block equates to zero (cf. off.diag = 0) due to the update of L . Thus (5.26) can be given as a block diagonal matrix. The remaining blocks on the diagonal represent the filtering and feedback inequalities.

Proof: Similar to before, we choose $\mathcal{Y} \in \mathbb{S}^{2n \times 2n}$ to be partitioned accordingly with \mathcal{A} as

$$\mathcal{Y} = \begin{bmatrix} Y & U \\ U^T & \bar{Y} \end{bmatrix}.$$

We will first apply a transform to \mathcal{Y} , then we define the appropriate values in the transformed problem. We again obtain (5.25) and (5.26) from (5.34) with the transforms of the proof of Theorem 5.5. For this we select L and F as introduced in Theorem 5.6. We note that this indeed reflects the chosen controller \mathbf{K} (5.35) for $U = (Y - X)$, and $V = -I$, that is $B_{\mathbf{K}} = -L$, $C_{\mathbf{K}} = F$ and $A_{\mathbf{K}} = A + BF + LC$. We note that for this choice of U, V , the applied transform is a congruence transform and we can construct \mathcal{Y} satisfying (5.34), if we can find X, Δ , and Z satisfying the inequalities (5.25) and (5.26). Note in (5.26) that the auxiliary variable A_{δ} equates to zero by the choice of $A_{\mathbf{K}}$, that is $A_{\delta} = 0$. Thus (block 1), (block 2) and (off diag.) become,

$$\begin{aligned} \text{(block 1): } & X - (A + LC)X(A + LC)^T - B_w B_w^T - L D_w D_w^T L^T, \\ \text{(block 2): } & \Delta - (A + BF)\Delta(A + BF)^T - L(D_w D_w^T + CXC^T)L^T, \\ \text{(off diag.): } & L(CXA^T + (CXC^T + D_w D_w^T)L^T). \end{aligned}$$

By choosing $X = Q$, the off diagonal term disappears, due to the specific structure of L . The matrix inequality is now decoupled

$$\begin{bmatrix} \text{(block 1b)} & 0 \\ 0 & \text{(block 2b)} \end{bmatrix} \succ 0. \quad (5.36)$$

with (block 1b) $\succ 0$ and (block 2b) $\succ 0$ given as

$$\begin{aligned} \text{(block 1b)} &:= Q - (A + LC)Q(A + LC)^T - B_w B_w^T - LD_w D_w^T L^T, \\ \text{(block 2b)} &:= \Delta - (A + BF)\Delta(A + BF)^T - L(D_w D_w^T + CQC^T)L^T. \end{aligned}$$

The choice of Q and L implies that (block 1b) $\succ 0$. This can be shown as follows. Given (5.32), i.e.,

$$Q - (A - \tilde{L}C)Q(A - \tilde{L}C)^T - B_w B_w^T - \tilde{L}D_w D_w^T \tilde{L}^T \succ 0$$

and given that the matrix $(A + LC)Q(A + LC)^T + LD_w D_w^T L^T$ as a function of L is smallest for $L = -AQC^T(CQC^T + D_w D_w^T)^{-1}$ it also follows that

$$Q - (A + LC)Q(A + LC)^T - B_w B_w^T - LD_w D_w^T L^T \succ 0$$

holds (c.f. Proposition 5.4) and that (block 1b) $\succ 0$ is satisfied.

For the given controller structure in (5.35), we only need to prove the existence of $\Delta, Z \succ 0$ such that (5.25), (block 2b) $\succ 0$ and $\text{tr}(Z) < \lambda^2$ are satisfied. Via a Schur complement, (5.25) can be written equivalently as

$$Z \succ (C_z + D_z F)\Delta(C_z + D_z F)^T + C_z Q C_z^T. \quad (5.37)$$

Introducing Z_1 , the inequalities can be trivially rewritten as

$$\begin{aligned} \text{tr}(Z_1) + \text{tr}(C_z Q C_z^T) &< \lambda^2 \\ Z_1 &\succ (C_z + D_z F)\Delta(C_z + D_z F)^T \\ \Delta &\succ 0, \quad \Delta - (A + BF)\Delta(A + BF)^T - L(CQC^T + D_w D_w^T)L^T \succ 0. \end{aligned}$$

Based on Proposition 5.1 the conditions on Z_1 , Δ and B_K can be expressed equivalently as

$$\begin{aligned} \lambda^2 &> \text{tr}(L(CQC^T + D_w D_w^T)L^T R) + \text{tr}(C_z Q C_z^T) \\ R &\succ 0, \quad R - (A + BF)^T R(A + BF) - (C_z + D_z F)^T (C_z + D_z F) \succ 0. \end{aligned}$$

This holds for $R = P$. We have shown the existence of X, U, Δ and Z subject to the inequalities (5.25) and (5.26) with $\text{tr}(Z) < \lambda^2$. Based on the definition of \mathcal{S} with $V = -I$ we can now recover \mathcal{Y} as $\mathcal{Y} = \mathcal{S}^{-1}\mathcal{R}$.

The author would like to acknowledge features of the non commutative algebra package NCAAlgebra [Helton et al., 2015] developed for mathematica for the discovery of this separation of matrix inequalities. \square

Remark 5.1 For notational brevity results in this paper assume that $D_w B_w^T = 0$ and $D_z C_z = 0$. All theorems and proofs can be adapted to account for $D_w B_w^T \neq 0$ and $D_z C_z \neq 0$. Specifically, in Theorem 5.5 the condition $D_z^T C_z = 0$ is not necessary. Additionally, the theorem can be extended to the case that $D_w B_w^T \neq 0$ by changing to the

appropriate gain update

$$L := -(AQC^T + B_w D_w^T) (CQC^T + D_w D_w^T)^{-1}.$$

In (5.26) of the proof, the off-diagonal block matrices will have an additional term. Again the off-diagonal matrices will equate to zero with this adapted gain update. Note that this again decouples the inequality in (5.26) into its blocks on the diagonal, that is, the filtering block and the state-feedback block on the diagonal.

5.4.d Optimality: the separation theorem

As a consequence of Theorem 5.5 and Theorem 5.6, we now prove the separation theorem. More precisely, we show that an optimal controller \mathbf{K}^* can be constructed based on the optimal observer gain of the estimation problem (5.8) and the optimal feedback gain for the feedback problem (5.10).

Theorem 5.7 (Separation theorem) *Let Assumption 5.1 and Assumption 5.2.1 hold. Consider P^* and Q^* as strong solutions of the DTARE (5.14) and (5.16). Then infimum λ_{inf} of the H_2 control design problem (5.6) is given as $\lambda_{\text{inf}} = \lambda_0$ with λ_0 defined in (5.22). If, in addition, Assumption 5.2.2 holds, then the controller (5.35) constructed with $F = F^* := -(B^T P^* B + D_z^T D_z)^{-1} B^T P^* A$ and $L = L^* := -A Q^* C^T (C Q^* C^T + D_w D_w^T)^{-1}$ is the optimal controller \mathbf{K}^* .*

Note that the optimal controller \mathbf{K}^* , is required to be stabilising. Assumptions 5.1-5.2 define a sufficient condition for this. The proof follows a standard Lyapunov-based perturbation argument.

Proof: Note that $\lambda_{\text{inf}} \geq \lambda_0$ follows trivially from Theorem 5.5. To prove the reversed inequality $\lambda_0 \geq \lambda_{\text{inf}}$ we choose any $\lambda > \lambda_0$ and show that $\lambda > \lambda_{\text{inf}}$ is true. For this purpose we exploit item (2) in Proposition 5.4 and its dual version. Due to the definition of λ_0 in (5.22) and since $\lambda > \lambda_0$, there exist P, F with (5.31) and Q, \tilde{L} with (5.32) such that (5.33) is valid. By Theorem 5.6 we hence conclude $\lambda > \lambda_{\text{inf}}$ as was to be shown.

To show that \mathbf{K}^* is the optimal controller, we need to show that given \mathbf{K}^* for each $\lambda > \lambda_{\text{inf}}$ there exists a \mathcal{Y} for which the inequalities in (5.6) are satisfied.

Conditional on assumptions 5.1 and 5.2 we have that $Q^* \succ 0$ and \mathcal{A} is stable. This follows from the stabilising estimator and feedback gain obtained based which, once combined into \mathcal{A} as given in (5.35), yield stabilising dynamics. For $\mathcal{Y} = \mathcal{Y}^* + \mathcal{Y}_\varepsilon$ and $\lambda_\varepsilon := \sqrt{\lambda^2 - \lambda_{\text{inf}}^2}$ the inequalities in (5.6) hold if

$$\begin{aligned} \sqrt{\text{tr}(\mathcal{C}\mathcal{Y}_\varepsilon\mathcal{C}^T)} &< \lambda_\varepsilon \\ \mathcal{Y}_\varepsilon &\succ 0, \quad \mathcal{Y}_\varepsilon - \mathcal{A}\mathcal{Y}_\varepsilon\mathcal{A}^T \succ 0 \end{aligned} \tag{5.38}$$

and

$$\begin{aligned} \sqrt{\text{tr}(\mathcal{C}\mathcal{Y}^*\mathcal{C}^T)} &\leq \lambda_{\text{inf}} \\ \mathcal{Y}^* &\succ 0, \quad \mathcal{Y}^* - \mathcal{A}\mathcal{Y}^*\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succeq 0. \end{aligned} \tag{5.39}$$

Since \mathcal{A} is stable, for any λ_ε there exists a \mathcal{Y}_ε such that (5.38) holds. Thus \mathbf{K}^* is optimal if it is stabilising and if (5.39) holds for some \mathcal{Y}^* . Similarly to the proofs of Theorem 5.5 and Theorem 5.6 we pick a congruence transform with S and we pick \mathcal{Y}^* as follows

$$\mathcal{Y}^* := \begin{bmatrix} \Delta + Q^* & \Delta \\ \Delta & \Delta \end{bmatrix} \text{ such that } \mathcal{R}\mathcal{S}^T = \begin{bmatrix} Q^* & 0 \\ 0 & \Delta \end{bmatrix}.$$

We assume that $\Delta \succ 0$, this allows us to repeat the first steps of the proof of Theorem 5.6. Since the subsequent Schur-decompositions therein and simplifications can be trivially extended to the current non-strict inequality as long as $Q^* \succ 0$ and $\Delta \succ 0$. The resulting inequalities, equivalent to (5.39) are

$$\text{tr}((C_z + D_z F)\Delta(C_z + D_z F)^T) + \text{tr}(C_z Q^* C_z^T) \leq \lambda_{\text{inf}}^2 \quad (5.40)$$

$$\begin{bmatrix} (\text{block } 1b^*) & 0 \\ 0 & (\text{block } 2b^*) \end{bmatrix} \succeq 0. \quad (5.41)$$

with $(\text{block } 1b^*) \succeq 0$ and $(\text{block } 2b^*) \succeq 0$ given as

$$\begin{aligned} (\text{block } 1b^*) &:= Q^* - (A + LC)Q^*(A + LC)^T - B_w B_w^T - L D_w D_w^T L^T, \\ (\text{block } 2b^*) &:= \Delta - (A + BF)\Delta(A + BF)^T - L(D_w D_w^T + C Q^* C^T)L^T. \end{aligned}$$

The former, $(\text{block } 1b^*) \succeq 0$, holds by construction since $(\text{block } 1b^*) = 0$. For the latter we choose $\Delta \succ 0$ such that $(\text{block } 2b^*) = 0$; this choice also satisfies (5.40). If the Lyapunov equation $(\text{block } 2b^*) = 0$ does not give a positive definite Δ , a simple Lyapunov argument can be used to get $\Delta \succ 0$ together with an additional slack variable ε_δ on (5.39). \square

5.5 Corollaries for H_2 performance

This section provides the dual versions of the results for the H_2 controller problem defined via (5.7). As a corollary of Theorem 5.6 and 5.7 we have the following result.

Corollary 5.8 *Let us be given feasible solutions P, F and Q, \tilde{L} such that (5.31) and (5.32), then there exists a matrix $\Delta \succ 0$ such that (5.6) is satisfied with the controller $\mathbf{K} \in \mathcal{K}$ composed as (5.35) and with*

$$\mathcal{Y} := \begin{bmatrix} Q + \Delta & \Delta \\ \Delta & \Delta \end{bmatrix}.$$

Corollary 5.9 *Let Q^* and P^* be positive semi-definite, strong solutions of the DTARE for the filtering problem (5.14) and for the feedback problem (5.16), respectively. Consider*

$$\lambda_0^2 := \text{tr}(B_w^T P^* B_w) + \text{tr}(A^T P^* B (B^T P^* B + D_z^T D_z)^{-1} B^T P^* A Q^*). \quad (5.42)$$

Then for any $\mathbf{K} \in \mathcal{K}$, $\mathcal{X} \succ 0$, $\lambda > 0$ subject to (5.7):

$$\sqrt{\text{tr}(B^T \mathcal{X} B)} < \lambda, \mathcal{X} \succ 0 \text{ and } \mathcal{X} - \mathcal{A}^T \mathcal{X} \mathcal{A} - \mathcal{C}^T \mathcal{C} \succ 0$$

λ is lower bounded by λ_0 , that is, $\lambda > \lambda_0$.

Proof: The result follows by repeating the proof of Lemma 5.9 for the dual optimisation problem. \square

Corollary 5.10 Given feasible solutions P, \tilde{F} and Q, L such that

$$\begin{aligned} P \succ 0, \quad P - (A + B\tilde{F})^T P (A + B\tilde{F}) - (C_z + D_z \tilde{F})^T (C_z + D_z \tilde{F}) \succ 0, \\ Q \succ 0, \quad Q - (A + LC)Q(A + LC)^T - (B_w + LD_w)(B_w + LD_w)^T \succ 0. \end{aligned}$$

Then for any λ subject to

$$\text{tr}(A^T P B (B^T P B + D_z^T D_z)^{-1} B^T P A Q) + \text{tr}(B_w^T P B_w) < \lambda^2$$

there exists an \mathcal{X} such that (5.7) is satisfied with the controller \mathbf{K} composed as

$$\begin{bmatrix} A_{\mathbf{K}} & B_{\mathbf{K}} \\ C_{\mathbf{K}} & 0 \end{bmatrix} = \begin{bmatrix} A + BF + LC & -L \\ F & 0 \end{bmatrix} \quad (5.43)$$

with the updated gain $F := -(B^T P B + D_z^T D_z)^{-1} B^T P A$.

Additionally, we can again make a statement on the structure of \mathcal{X} satisfying (5.7) in the above corollary.

Corollary 5.11 Under the same conditions as in Corollary 5.10 there exists a matrix $\Delta \succ 0$ such that (5.7) is satisfied with the controller $\mathbf{K} \in \mathcal{K}$ composed as (5.43) and with

$$\mathcal{X} := \begin{bmatrix} P + \Delta & \Delta \\ \Delta & \Delta \end{bmatrix}.$$

Note that Δ in the above statement, is obtained via solving the dual of the state-estimation problem.

Theorem 5.12 (Separation theorem) Let Assumption 5.1 and Assumption 5.2.1 hold. Consider P^* and Q^* as strong solutions of the DTARE (5.14) and (5.16). Then infimum λ_{inf} of the H_2 control design problem (5.6) is given as $\lambda_{\text{inf}} = \lambda_0$ with λ_0 defined in (5.22).

If, in addition, Assumption 5.2 holds then the controller (5.43) constructed with $F := F^* = -(B^T P^* B + D_z^T D_z)^{-1} B^T P^* A$ and $L := L^* = -A Q^* C^T (C Q^* C^T + D_w D_w^T)^{-1}$ is the optimal controller \mathbf{K}^* .

5.6 Problem statement for hierarchical control

5.6.a System description & hierarchical controller design

System description. Consider again the system \mathbf{M} whose dynamics are given by the mathematical model (5.2), that is,

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t) + B_w w(t) \\ y(t) &= Cx(t) + D_w w(t) \\ z(t) &= C_z x(t) + D_z u(t). \end{aligned} \quad (5.44)$$

The system is subject to stochastic disturbances $w(t) \in \mathbb{R}^{w_n}$ on its state transitions and on its measurements, modelled as a white noise sequence with unit covariance, i.e., $\mathbb{E}[w(t)w(t)^T] = I_{w_n}$. We still assume that the noise on state and measurement is decoupled, that the criteria on u and x are decoupled, and that we have dynamics that are stabilisable and detectable (cf. Assumption 5.1). Unlike the previous part of this chapter, we generalise the results beyond Assumption 5.2. When necessary a perturbation term can be used to achieve Assumption 5.2. The choice of C_z and D_z is related to the design requirements. Of interest to us is the case of correct-by-design control (see Chapter 4) for which these matrices are chosen based on the temporal logic properties of interest. Alternatively, as standard, they can also define a performance metric that weights the control versus state deviations.

Hierarchical controller design. Consider a deterministic version of \mathbf{M} , denoted $\bar{\mathbf{M}}$, given as

$$\begin{aligned} \bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}(t) \\ y(t) &= C\bar{x}(t) \\ z(t) &= C_z \bar{x}(t) + D_z \bar{u}(t). \end{aligned} \quad (5.45)$$

Let $\{(\bar{x}(t), \bar{u}(t))\}_{t \geq 0}$ be a trajectory of the behaviour of $\bar{\mathbf{M}}$ that is $\{(\bar{x}(t), \bar{u}(t))\}_{t \geq 0} \in \mathcal{B}(\bar{\mathbf{M}})$. For this deterministic model the control task may be solved in a first step. This yields a planned trajectory within $\mathcal{B}(\bar{\mathbf{M}})$. Figure 5.1, depicts how such a trajectory can then be refined to yield a full hierarchical control implementation for the original system \mathbf{M} .

The goal of the control refinement is to follow the deterministic (pre-planned) trajectory with the trajectory of the noisy system \mathbf{M} , that is, $\{(x(t), u(t))\}_{t \geq 0}$. We are interested in a-priori bounds on the accuracy loss, expressed via C_z and D_z , that hold for any trajectory in $\mathcal{B}(\bar{\mathbf{M}})$. More precisely, we question how to design a control refinement, such that for any trajectory in $\mathcal{B}(\bar{\mathbf{M}})$, the controlled trajectory of \mathbf{M} follows with quantified precision.

As in Figure 5.1, we define a controller \mathbf{K} parameterised in L and F , which are,

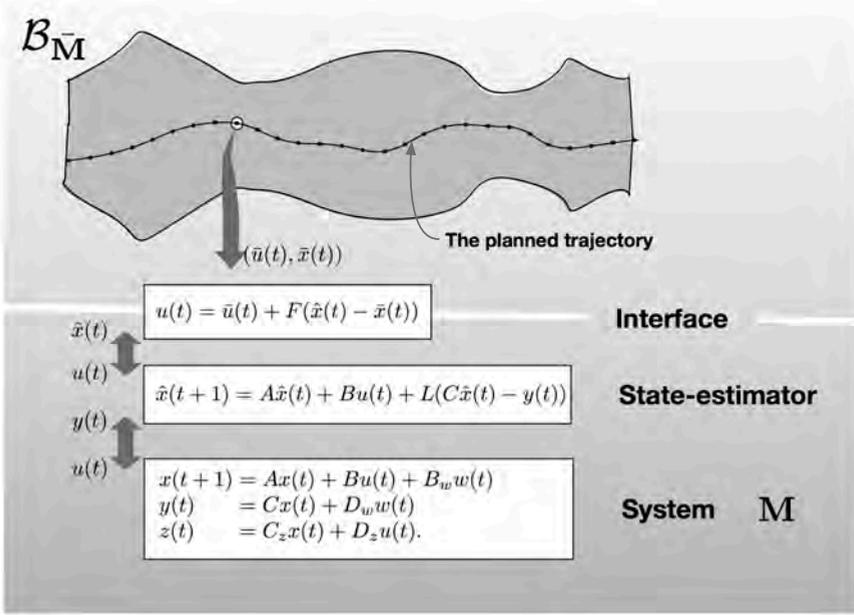


Figure 5.1: The figure portrays the control refinement problem within a hierarchical control setup. The top part of the figure depicts the behaviour of the deterministic abstraction of M . The control task at hand is first solved on \bar{M} and yields a planned trajectory within $\mathcal{B}_{\bar{M}}$. Then, as a next step in the hierarchy, this is refined to the real system based on a state-estimator and an interface function.

respectively, the state-estimator gain and the state-feedback gain as

$$\begin{aligned}\hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + L(C\hat{x}(t) - y(t)), \\ u(t) &= \bar{u}(t) + F(\hat{x}(t) - \bar{x}(t)).\end{aligned}$$

Note that we can write the controller explicitly as a function of its inputs (\bar{x}, \bar{u}, y) and output u , that is

$$\begin{aligned}\hat{x}(t+1) &= (A + BF + LC)\hat{x}(t) - Ly(t) + B(\bar{u}(t) - F\bar{x}(t)), \\ u(t) &= \bar{u}(t) + F(\hat{x}(t) - \bar{x}(t)).\end{aligned}\tag{5.46}$$

The performance of the controller is now quantified with respect to $z(t) - \bar{z}(t)$ where $\bar{z}(t) = C_z \bar{x}(t) + D_z \bar{u}(t)$. Denote the parameterised set of possible controllers \mathcal{K}_p indexed by L and F with elements $\mathbf{K} \in \mathcal{K}_p$ defined as (5.46). We assume, without loss of generality, that we are free to choose the initialisation of $\bar{x}(0)$ equal to the initial state of the observer $\hat{x}(0)$.

To analyse the performance, we analyse the interconnection of the controlled system (5.44)-(5.46) together with the transitions of \bar{M} over the trajectory of interest.

This yields

$$\begin{aligned}
 \bar{x}(t+1) &= A\bar{x}(t) && + B\bar{u}(t) \\
 \hat{x}(t+1) &= -BF\bar{x}(t) + (A+LC+BF)\hat{x}(t) - LCx(t) + B\bar{u}(t) - LD_w w(t) \\
 x(t+1) &= -BF\bar{x}(t) + BF\hat{x}(t) + Ax(t) && + B\bar{u}(t) + B_w w(t)
 \end{aligned}$$

$$z(t) - \bar{z}(t) = -(C_z + D_z F)\bar{x}(t) + D_z F\hat{x}(t) + C_z x(t).$$

We use $z(t) - \bar{z}(t)$ to quantify the deviation of the trajectory tracking based on the implemented control refinement.

Similarly to the state transform of the specific case presented in Chapter 4, we can again take a state transform and separate the state (x, \hat{x}, \bar{x}) into the transitions of \bar{x} based on $\bar{\mathbf{M}}$, and error dynamics $\xi_e = \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix}$. More precisely

$$\begin{aligned}
 \bar{x}(t+1) &= A\bar{x}(t) + B\bar{u}(t) \\
 \xi_e(t+1) &= \mathcal{A}_e \xi_e(t) + \mathcal{B}_e w(t) \\
 z(t) - \bar{z}(t) &= \mathcal{C}_e \xi_e(t)
 \end{aligned} \tag{5.47}$$

$$\text{with } \left[\begin{array}{c|c} \mathcal{A}_e & \mathcal{B}_e \\ \hline \mathcal{C}_e & 0 \end{array} \right] = \left[\begin{array}{cc|c} A+BF & -LC & -LD_w \\ 0 & A+LC & B_w + LD_w \\ \hline C_z + D_z F & C_z & 0 \end{array} \right]. \tag{5.48}$$

We note that (5.47) consists of the dynamics of the deterministic system driven by $\bar{u}(t)$, composed with the error dynamics ξ_e . The latter dynamics are solely governed by the noise $w(t)$. Thus we have separated the deterministic control problem of \bar{x} , which could be for example a navigation problem, from the noise rejection problem of ξ_e .

Next, we quantify the accuracy and performance of a controller $\mathbf{K} \in \mathcal{K}_p$ as a function of the expected difference in $z(t) - \bar{z}(t)$. Consecutively, we seek a constructive formulation of the criteria, such that we can automatically construct a controller $\mathbf{K} \in \mathcal{K}_p$.

Accuracy. We want to quantify how close the controlled system will stay to the trace over the whole time horizon. Similarly to Chapter 4 and for a given choice of F and L , we quantify the accuracy loss with respect to the expected deviation

$$\sqrt{\mathbb{E}[\|z(t) - \bar{z}(t)\|_2^2]}$$

for all $\{(\bar{x}(t), \bar{u}(t))\}_{t \geq 0}$ and uniformly over time¹. For the quantification of the accuracy loss, we need to take into account the potential scenarios for the initialisation of \mathbf{M} . More precisely, we could know the initial state of \mathbf{M} exactly; or if we do not know it exactly we can either assign a Gaussian distribution to it or we can define a bounded set \mathbb{X}_0 to which it belongs.

¹Remark that in the previous chapter the accuracy expressed as $\mathbb{E}[\|z(t) - \bar{z}(t)\|_2]$ was bounded based on $\mathbb{E}[\|z(t) - \bar{z}(t)\|_2^2]$.

Firstly, when the initial state of \mathbf{M} , that is $x(0)$, is known exactly, then the accuracy loss is upper bounded with ε , if there exists an $\mathfrak{X} \in \mathbb{S}^{2n}$ subject to

$$\sqrt{\text{tr}(\mathcal{C}_e \mathfrak{X} \mathcal{C}_e^T)} \leq \varepsilon, \quad (5.49a)$$

$$\mathfrak{X} \succ 0, \quad (5.49b)$$

$$\mathfrak{X} - \mathcal{A}_e \mathfrak{X} \mathcal{A}_e^T - \mathcal{B}_e \mathcal{B}_e^T \succ 0. \quad (5.49c)$$

This follows by initialising the observer as $\hat{x}(0) = x(0)$, which thereby also yields an initialisation of the error dynamics $\xi_e(0) = 0$. The derivation of the matrix inequalities trivially follows the line reasoning given in Chapter 4.

When the initial state of \mathbf{M} is not known exactly, then we need to also take into account the transient behaviour. Firstly, suppose the initial state of \mathbf{M} can be defined as a random variable $x(0) \sim \mathcal{N}(x_0, P_0)$. The accuracy loss is upper bounded uniformly over time with ε , if there exists a $\mathfrak{X} \in \mathbb{S}^{2n}$ subject to (5.49a), (5.49c) and

$$\mathfrak{X} \succ \begin{bmatrix} 0 & 0 \\ 0 & (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & P_0 \end{bmatrix}. \quad (5.50)$$

Inequality (5.50) replaces the weaker positive definiteness condition (5.49b). Suppose that the initialisation of \mathbf{M} is known up to a set \mathbb{X}_0 , that is $x(0) \in \mathbb{X}_0$. Then the accuracy is at least ε , if there exists a $\mathfrak{X} \in \mathbb{S}^{2n}$ subject to (5.49a), (5.49c) and

$$\mathfrak{X} \succ \begin{bmatrix} 0 & 0 \\ 0 & (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T \end{bmatrix}, \forall x_0 \in \mathbb{X}_0. \quad (5.51)$$

Remark 5.2 Remark that for $D_z = 0$, we recover the set of conditions represented by matrix inequalities in Theorem 4.7 of Chapter 4. As such these conditions equivalently represent the accuracy of an approximate bisimulation. Similarly, for $D_z = 0$, $B_w = 0$, and $D_w = 0$, we recover the Lyapunov-like matrix inequalities in Theorem 4.6 of Chapter 4.

Even though accuracy with respect to the state deviation, or the loss thereof, is key in the control refinements, a design that solely optimises with respect to this metric is impractical. In fact, considering an accuracy metric with respect to the weighted deviation of $x - \bar{x}$ only ($D_z = 0$), as already mentioned in Chapter 4, would mean that the accuracy of the control actions will be not penalised. Even though the control actions will still be bounded, unlike in the continuous time case, this induces large control actions and deadbeat like behaviour in the controlled system. To avoid this corner cases, we often need to add a secondary criteria on the performance of the controller next to the accuracy.

Performance. We define the performance of the control refinement based on the asymptotic variance of $x(t) - \bar{x}(t)$ and $u(t) - \bar{u}(t)$ for all $\bar{u}(t)$. For this, we introduce a secondary set of matrices $C_{z,p}$ and $D_{z,p}$, which quantify the desired trade-off the state-deviation with the deviation of the control action $u(t) - \bar{u}(t)$. More precisely,

on the controlled system we define the performance (deviation) signal as

$$z_p(t) - \bar{z}_p(t) := C_{z,p}(x(t) - \bar{x}(t)) + D_{z,p}(u(t) - \bar{u}(t))$$

and we introduce

$$\mathcal{C}_p := \begin{bmatrix} C_{z,p} + D_{z,p}F & C_{z,p} \end{bmatrix}.$$

Given \mathcal{C}_p , we can now quantify the H_2 performance of the output based controller based on $(\mathcal{A}_e, \mathcal{B}_e, \mathcal{C}_p)$ as λ subject to the existence of $\mathfrak{X} \in \mathbb{S}^{2n}$ such that

$$\begin{aligned} \sqrt{\text{tr}(\mathcal{C}_p \mathfrak{X} \mathcal{C}_p^T)} &< \lambda, & \mathfrak{X} &\succ 0 \\ \mathfrak{X} - \mathcal{A}_e \mathfrak{X} \mathcal{A}_e^T - \mathcal{B}_e \mathcal{B}_e^T &\succ 0. \end{aligned} \quad (5.52)$$

5.6.b Problem statement: construction of the control refinement

We now investigate how we to construct a controller $\mathbf{K} \in \mathcal{K}_p$ such that we obtain

- minimal accuracy loss, and
- optimal performance.

Furthermore, we want to be able to take into account the different types of initial scenarios, that is,

- a known initial state,
- an initial state known up-to Gaussian distribution, or
- an initial state known up to a bounded set \mathbb{X}_0 .

A design of \mathbf{K} as a direct optimisation of the given accuracy bounds 5.49 or optimality bounds 5.52 is difficult. The matrix inequalities are parameterised in L and F , as such they are non-linearly parameterised. Additionally, because they are not fully parameterised as was the case for the general H_2 -norm with control structure (5.3), it follows that the transformation to a linear parameterisation [Scherer and Weiland, 2000], standardly used for these control designs, cannot be applied. Hence instead, we will exploit the separated feasibility matrices obtained in the preceding part of this chapter to develop a constructive approach to the design of the controller with guaranteed accuracy and, if applicable, performance.

5.7 Computational approach to output-based control refinement

5.7.a Introduction

We consider the design of $\mathbf{K} \in \mathcal{K}_p$ with quantified bounds on the accuracy loss and the performance. Of interest are computational constructive approaches towards the design of \mathbf{K} . We seek a way of defining the construction of \mathbf{K} as a

(sequence) of convex optimisation problems that can be efficiently solved. Firstly in Section 5.7.b, we evaluate the case where the initialisation of the system \mathbf{M} is known and we give the controller synthesis with respect to the accuracy loss (5.49) as a basic result of the H_2 -control problem. In the same subsection, we also investigate controller synthesis with respect to both the accuracy loss and the performance (5.52), as a lexicographic optimisation problem. Then in Section 5.7.c, we show how the controller synthesis can be extended to include the initialisation for the case that the initial state of \mathbf{M} is not fully specified, as given by the condition in (5.50) or in (5.51).

5.7.b Controller synthesis via LMI's

Consider the greatest lower bound on the attainable accuracy subject to (5.49), this quantity ε_{acc} defined formally as the infimum of the accuracy problem as

$$\varepsilon_{acc} := \inf_{\varepsilon, \mathbf{x}, \mathbf{K} \in \mathcal{K}_p} \varepsilon \text{ s.t. (5.49).}$$

We can show that this infimum is equal to the infimum of the H_2 control design problem in (5.6) for equal system matrices.

More precisely, let there be strong solutions of the DTARE for the state-estimation and feedback problem, respectively (5.14) and (5.16), then the infimum of (5.6) is equal to the infimum of the accuracy problem.

Corollary 5.13 *The infimum of the accuracy problem ε_{acc} is equal to the infimum of (5.6).*

Proof: The proof follows directly from the state transform introduced in the proof of Corollary 5.14, which transforms the problem to the original H_2 -problem. Then the infimum of the accuracy problem ε_{acc} is equal to the infimum of (5.6) defined over \mathcal{K}_p . Since ε_{acc} is still a lower bound and since for every $\varepsilon > \varepsilon_{acc}$ there exist a controller $\mathbf{K} \in \mathcal{K}_p$ such that (5.6) is satisfied it follows that ε_{acc} is the infimum. \square

As before, if the strong solutions of the DTARE of the state-feedback and the state-estimators are both stabilising and positive definite, then they define an optimal controller \mathbf{K} with respect to the accuracy loss.

Lexicographic optimisation. We now consider the design L and F with respect to both accuracy and performance. The goal is to find a controller in the bounded region of the minimal accuracy loss that is optimal with respect to the H_2 -norm λ as in (5.52). More precisely, we relax the optimality of the accuracy with a small positive value $\gamma > 0$ and we define this set of controllers close to the optimal accuracy as

$$\mathcal{K}_p^\gamma := \{\mathbf{K} \in \mathcal{K}_p \mid \varepsilon < \varepsilon_{acc} + \gamma \text{ with } \varepsilon \text{ s.t. (5.49)}\}.$$

Now the optimal controller, within this lexicographic optimisation problem, is defined as

$$\inf_{\mathfrak{X}, \varepsilon, \mathbf{K} \in \mathcal{K}_p^\gamma} \lambda \quad \text{subject to (5.52).}$$

We propose Algorithm 4, where the design of L and F is done with respect to the performance λ and subject to the bound on the accuracy (5.52), hence delivering guarantees on both the performance and the accuracy. We compute the minimised the accuracy loss and compute the Kalman gain L^* and Q^* . Then we relax the achieved accuracy with γ and select an updated state-feedback gain F that optimises the H_2 -performance criteria.

Algorithm 4 Refinement construction

- 1: Compute ε_{acc} , for strong solution Q^* and L^* by solving the accuracy problem (5.49) as a standard H_2 problem via the separation theorem
- 2: Fix the choice of estimator gain $L := L^*$, $Q := Q^*$, and find F, Δ and λ minimising $\lambda > 0$ subject to

$$\begin{aligned} \text{s.t. } & \text{tr}((C_{z,p} + D_{z,p}F)\Delta(C_{z,p} + D_{z,p}F)^T) + \text{tr}(C_{z,p}QC_{z,p}) < \lambda, \\ & \text{tr}((C_z + D_zF)\Delta(C_z + D_zF)^T) + \text{tr}(C_zQC_z) < \varepsilon_{acc} + \gamma, \\ & \Delta - (A + BF)\Delta(A + BF)^T - AQC^T(D_wD_w^T + CQC^T)^{-1}CQA^T \succ 0, \\ & \Delta \succ 0. \end{aligned}$$

- 3: Construct refinement with $L := L^*$ and F
-

We note that line 1 can be solved based on the DTARE as long as it considers only a solution based on strong solutions P^* and Q^* of the corresponding DTAREs. Otherwise, even under the detectability and stability requirements, the lack of stabilisability requirement could make \mathcal{K}_p^γ empty as it is defined based on strict inequalities. Of course this can also be computed based on an LMI optimisation.

5.7.c Controller synthesis via LMI's with initialisation

Accuracy & Initialisation. Let us extend the control refinement or control synthesis to the case where the initial state of the system is either defined by a probability distribution or a bounded set. That is, we want to find $\varepsilon, \mathfrak{X}, F, L, \Delta_x$ such that

$$\text{s.t. } \text{tr}(C_e \mathfrak{X} C_e^T) \leq \varepsilon^2 \tag{5.53a}$$

$$\mathfrak{X} - A_e \mathfrak{X} A_e^T - B_e B_e^T \succ 0, \tag{5.53b}$$

$$\mathfrak{X} \succ \begin{bmatrix} 0 & 0 \\ 0 & \Delta_x \end{bmatrix}, \tag{5.53c}$$

$$\text{and } \Delta_x \text{ s.t. } \begin{cases} \Delta_x \succ (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T, \forall x_0 \in \mathbb{X}_0, \text{ or} \\ \Delta_x \succ (x_0 - \hat{x}(0))(x_0 - \hat{x}(0))^T + P_0. \end{cases} \tag{5.53d}$$

Again ε defines the accuracy of this controller refinement. We apply the result of Theorem 5.6 to this case as follows.

Corollary 5.14 *Given feasible solutions of P, F and Q, \tilde{L} subject to*

$$P \succ 0, \quad P - (A + BF)^T P (A + BF) - (C_z + D_z F)^T (C_z + D_z F) \succ 0 \quad (5.54)$$

$$Q \succ \Delta_x, \quad Q - (A + \tilde{L}C)Q(A + \tilde{L}C)^T - (B_w + \tilde{L}D_w)(B_w + \tilde{L}D_w)^T \succ 0. \quad (5.55)$$

Then for any $\varepsilon > 0$ such that

$$\text{tr}(AQC^T(CQC^T + D_w D_w^T)^{-1}CQA^T P) + \text{tr}(C_z Q C_z^T) \leq \varepsilon^2 \quad (5.56)$$

there exists an \mathfrak{X} such that (5.53a),(5.53c) and (5.53b) are satisfied for controller $\mathbf{K} \in \mathcal{K}_p$ with gains F and the updated estimator gain $L := -AQC^T(CQC^T + D_w D_w^T)^{-1}$.

Note that the constraint on the initialisation has been changed into a lower bound on Q , see $Q \succ \Delta_x$ in (5.55). To prove this we apply a state transform to the system and obtain the same system dynamics as in Theorem 5.6.

Proof: [of Corollary 5.14] We transform (5.53a),(5.53c) and (5.53b) to, respectively,

$$\sqrt{\text{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T)} < \varepsilon, \quad \mathcal{Y} - \begin{bmatrix} \Delta_x & 0 \\ 0 & 0 \end{bmatrix} \succ 0 \quad (5.57)$$

$$\mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0 \quad (5.58)$$

based on the state transform

$$\mathcal{A} = \begin{bmatrix} I & I \\ I & 0 \end{bmatrix} \mathcal{A}_e \begin{bmatrix} I & I \\ I & 0 \end{bmatrix}^{-1} = \begin{bmatrix} A & BF \\ -LC & A + LC + BF \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} I & I \\ I & 0 \end{bmatrix} \mathcal{B}_e = \begin{bmatrix} B_w \\ -LD_w \end{bmatrix},$$

$$\mathcal{C} = \mathcal{C}_e \begin{bmatrix} I & I \\ I & 0 \end{bmatrix}^{-1} = [C_z \quad D_z F], \quad \text{with transformed state } \begin{bmatrix} I & I \\ I & 0 \end{bmatrix} \begin{bmatrix} \hat{x} - \bar{x} \\ x - \hat{x} \end{bmatrix} = \begin{bmatrix} x - \bar{x} \\ \hat{x} - \bar{x} \end{bmatrix}.$$

\mathcal{Y} is related to \mathfrak{X} as

$$\mathcal{Y} = \begin{bmatrix} I & I \\ I & 0 \end{bmatrix} \mathfrak{X} \begin{bmatrix} I & I \\ I & 0 \end{bmatrix}^T.$$

We can now leverage Theorem 5.6, hence based on (5.54) and (5.55) we have that there exists a \mathcal{Y} such that $\sqrt{\text{tr}(\mathcal{C}\mathcal{Y}\mathcal{C}^T)} < \varepsilon$ and $\mathcal{Y} - \mathcal{A}\mathcal{Y}\mathcal{A}^T - \mathcal{B}\mathcal{B}^T \succ 0$. Furthermore, based on Corollary 5.8, we can construct \mathcal{Y} as

$$\begin{bmatrix} Q + \Delta & \Delta \\ \Delta & \Delta \end{bmatrix},$$

which based on $Q \succ \Delta_x$ implies that

$$\mathcal{Y} - \begin{bmatrix} \Delta_x & 0 \\ 0 & 0 \end{bmatrix} \succ 0.$$

Based on the inverse state transform we retain the guarantee that we can also find an \mathfrak{X} such that (5.53a),(5.53c) and (5.53b) are satisfied, which is what we wanted to prove. \square

Note that we are unable to construct both gains (P, F) and (Q, L) based on the above corollary as a single convex optimisation. The reason is that the objective ϵ depends in (5.56) nonlinearly on the variables Q and P . Given Q and \tilde{L} such that the inequalities in (5.55) are satisfied, then one can optimise with respect to (P, F) .

Thus, as in Algorithm 4, we can design the controller in two steps, first design both L and F with respect to accuracy and then update F via the lexicographic optimisation to take performance into account. For the first step we need to take the initialisation into account. In other words the selection of L and F , depends on the design of Δ_x and we will give an approach, where we first

- compute Δ_x of (5.53d), then
- optimise F and L with respect to accuracy, and if necessary
- include the performance criteria.

Compute Δ_x of (5.53d). Note that both the choice of Δ_x and that of $\hat{x}(0)$ (and $\bar{x}(0)$, i.e., $\bar{x}(0) = \hat{x}(0)$) based on the possible set of $x(0) \in \mathbb{X}_0$ is a design choice, which is subject to a semi-definite constraint. As such denote \hat{x}_0 as the to-be-chosen initialisation $\hat{x}(0) := \hat{x}_0$ and $\bar{x}(0) := \hat{x}_0$. Now on can give several examples where the semi-definite constraint over \mathbb{X}_0 ,

$$\forall x_0 \in \mathbb{X}_0 : \quad \Delta_x \succ (x_0 - \hat{x}_0)(x_0 - \hat{x}_0)^T, \quad (5.59)$$

can be rewritten into an LMI. As an example, consider the case that \mathbb{X}_0 is a polytope, then there exists a finite set of vertices, denoted \mathcal{V} , such that $\mathbb{X}_0 = \text{conv}(\mathcal{V})$. The condition in (5.59) is equal to the requirement that $\mathbb{X}_0 - \hat{x}_0$ is a subset of the ellipsoid defined by Δ_x . The constraint in (5.59) can then be replaced by a finite number of matrix inequalities constraints

$$\forall x_{i0} \in \mathcal{V} : \quad \Delta_x \succ (x_{i0} - \hat{x}_0)(x_{i0} - \hat{x}_0)^T, \quad (5.60)$$

with design variables \hat{x}_0 and Δ_x . Each of these inequalities can be rewritten as

$$\forall x_{i0} \in \mathcal{V} : \quad \begin{bmatrix} \Delta_x & (x_{i0} - \hat{x}_0) \\ (x_{i0} - \hat{x}_0) & I \end{bmatrix} \succ 0. \quad (5.61)$$

Given Δ_x compute L and F . To find a control refinement $\mathbf{K} \in \mathcal{K}_p$ based on accuracy alone, we want to select L and F minimising the accuracy loss

$$\text{tr}(AQC^T(CQC^T + D_w D_w^T)^{-1}CQA^T P) + \text{tr}(C_z Q C_z^T)$$

and subject to (5.54) and (5.55). For any pair L, \mathfrak{X} , the objective is a weighted trace of P . Hence P and F can be obtained as a strong solution of the (generalised)

DTARE. Alternatively, it can also be implemented as a LMI-based convex optimisation. For the optimal² P, F the inequality in (5.54) becomes an equality and the objective can be rewritten as

$$\begin{aligned} & \text{tr} \left(QP - QA^T PA + QF^T (D_z^T D_z + B^T PB) F \right) \\ & \quad + \text{tr} \left(AQC^T (D_w D_w^T + CQC^T)^{-1} CQA^T P \right). \end{aligned}$$

For the operations remember that $\text{tr}(AB) = \text{tr}(BA)$ and $\text{tr}(A+B) = \text{tr}(A) + \text{tr}(B)$. The objective is still non-linear in Q . Replacing the inequality in (5.55) with an equality in the above objective yields

$$\text{tr} (B_w B_w^T P) + \text{tr} (QF^T (D_z^T D_z + B^T PB) F).$$

Where the above objective is linear in Q . In all solutions of (5.55) will give strict inequalities due to the presence of $Q \succ \Delta_x$. We can still use the above formula, which is linear in Q , as an upper bound. More precisely,

$$\begin{aligned} & \text{tr} (QP - QA^T PA + QF^T (D_z^T D_z + B^T PB) F) \\ & \quad + \text{tr} \left(AQC^T (D_w D_w^T + CQC^T)^{-1} CQA^T P \right) \\ & \leq \text{tr} (B_w B_w^T P) + \text{tr} (QF^T (D_z^T D_z + B^T PB) F). \end{aligned}$$

Based on this we give following practical design of F and L with respect to accuracy alone. We follow a procedure also implemented in Matisse [Girard and Pappas, 2007] that uses the DTARE to compute a stabilising solution which solves the (5.55). For a given Δ_x , we introduce $Q_\delta := Q - \Delta_x$. Then (5.55) becomes

$$\begin{aligned} Q_\delta \succ 0, \quad & (Q_\delta + \Delta_x) - (A + \tilde{L}C)(Q_\delta + \Delta_x)(A + \tilde{L}C)^T \\ & - (B_w + \tilde{L}D_w)(B_w + \tilde{L}D_w)^T \succ 0. \end{aligned}$$

The second constraint can be written as³

$$\begin{aligned} & Q_\delta - (A + \tilde{L}C)Q_\delta(A + \tilde{L}C)^T \\ & \quad \succ (A + \tilde{L}C)\Delta_x(A + \tilde{L}C)^T - \Delta_x + (B_w + \tilde{L}D_w)(B_w + \tilde{L}D_w)^T \\ & Q_\delta - (A + \tilde{L}C)Q_\delta(A + \tilde{L}C)^T \\ & \quad \succ [I \quad \tilde{L}] \left(\begin{bmatrix} A\Delta_x A^T - \Delta_x & A\Delta_x C^T \\ C\Delta_x A^T & C\Delta_x C^T \end{bmatrix} + \begin{bmatrix} B_w B_w^T & B_w D_w^T \\ D_w B_w^T & D_w D_w^T \end{bmatrix} \right) \begin{bmatrix} I \\ \tilde{L}^T \end{bmatrix}. \end{aligned}$$

Of interest is a positive definite upper bound for the right side of the above in-

²For approximate solutions subject to the LMI that are close enough to the optimal the reasoning still holds.

³Note that we have assumed for notational brevity that $B_w D_w^T = 0$.

equality, that is

$$\begin{bmatrix} A\Delta_x A^T - \Delta_x & A\Delta_x C^T \\ C\Delta_x A^T & C\Delta_x C^T \end{bmatrix} + \begin{bmatrix} B_w B_w^T & B_w D_w^T \\ D_w B_w^T & D_w D_w^T \end{bmatrix} \preceq \begin{bmatrix} \tilde{B}_w \\ \tilde{D}_w \end{bmatrix} \begin{bmatrix} \tilde{B}_w \\ \tilde{D}_w \end{bmatrix}^T. \quad (5.62)$$

The smallest \tilde{B}_w, \tilde{D}_w can be computed by writing the left side of the inequality into its eigenvalue decomposition and equating any negative eigenvalues to zero. A similar procedure has been implemented in Matisse [Girard and Pappas, 2007] for Lyapunov inequalities with a lower bound on the Lyapunov matrix. Remark that $\tilde{B}_w \tilde{D}_w^T \neq 0$. A sufficient condition for (5.55) is

$$Q_\delta - (A + \tilde{L}C)Q_\delta(A + \tilde{L}C)^T - (\tilde{B}_w + \tilde{L}\tilde{D}_w)(\tilde{B}_w + \tilde{L}\tilde{D}_w)^T \succ 0 \text{ with } Q_\delta \succ 0.$$

Therefore we approximate the greatest lower bound on the accuracy with $\hat{\varepsilon}_{acc}$ by computing the smallest strong solution⁴ of the corresponding generalised DTARE the resulting Q_δ^* is then used to compute $Q := Q_\delta^* + \Delta_x$ and \tilde{L} as

$$\tilde{L} := -(AQ_\delta^* C^T + \tilde{B}_w \tilde{D}_w^T)(CQ_\delta^* C^T + D_w D_w^T)^{-1}.$$

If $A - \tilde{L}C$ is not stabilising then disturbing (5.62) such that it becomes strictly positive definite will force the existence of a stabilising strong solution. The obtained pair Q and \tilde{L} gives a non-strict satisfaction of the strict inequalities, still, remark that all obtained results on guarantees still hold. More precisely the stabilising solution allows us to add a (trivially small) Lyapunov matrix with respect to $(A - \tilde{L}C)$ such that the strict inequality is satisfied. By updating the estimator gain L to

$$L := -AQC^T(CQC^T + D_w D_w^T)^{-1}$$

we obtain a controller refinement design L and F based on accuracy alone, where the attained accuracy computed with (5.56) is denoted as $\hat{\varepsilon}_{acc}$.

Inclusion of performance. As before we include the performance metric by redesigning F with respect to the performance and the accuracy. Again, we relax the attained accuracy in the previous part with γ and we define this set of controllers close to the optimal accuracy as

$$\mathcal{K}_p^\gamma := \{\mathbf{K} \in \mathcal{K}_p \mid \varepsilon < \hat{\varepsilon}_{acc} + \gamma \text{ with } \varepsilon \text{ s.t. (5.53)}\}.$$

Thus for the choice of estimator gain L, Q , find F, Δ and λ minimising $\lambda > 0$

⁴subject to its existence

subject to

$$\begin{aligned} \text{s.t. } & \text{tr}((C_{z,p} + D_{z,p}F)\Delta(C_{z,p} + D_{z,p}F)^T) + \text{tr}(C_{z,p}QC_{z,p}) < \lambda, \\ & \text{tr}((C_z + D_zF)\Delta(C_z + D_zF)^T) + \text{tr}(C_zQC_z) < \hat{\varepsilon}_{acc} + \gamma, \\ & \Delta - (A + BF)\Delta(A + BF)^T - AQC^T(D_wD_w^T + CQC^T)^{-1}CQA^T \succ 0, \\ & \Delta \succ 0. \end{aligned}$$

5.8 Future work and conclusions

In this chapter, we have proved a novel result on the separation problem for the synthesis of controllers resulting in closed-loop systems with a guaranteed bound on their H_2 performance. We have taken an approach using linear matrix inequalities to show that such a controller can, in fact, be designed by separately designing a state-feedback controller and an estimator. As such, the main result of this chapter provides a separation theorem in which an estimator with guaranteed H_2 performance is designed independently from a state-feedback controller with guaranteed H_2 performance so as to result in an output feedback controller that achieves guaranteed H_2 performance for the controlled system. Hence, this result generalises the separation theorem from optimal to bounded H_2 performance. The setting of feasibility tests in terms of matrix inequalities is the novel feature of this contribution. This avoids conditions on uniqueness of solutions of Riccati equations.

The results of this chapter provide a first instance in which feasibility tests (by semi-definite programs) of a state-feedback design and estimator design define a controller with guaranteed H_2 performance. These results allow for relaxations on the optimal H_2 performance to meet additional design specifications. The design problem introduced in Chapter 4 presents a synthesis problem with an inherent observer-feedback structure. For this case we have used the developed theory to reason about the constructive design of control refinements (operating mainly in polynomial time) for which we have guaranteed bounds on the accuracy and when required also on the performance.

5.A Operations on linear matrix inequalities

Definition 5.15 ([Scherer and Weiland, 2000]) *A linear matrix inequality (LMI) is an inequality $F(x) \prec 0$ where F is an affine function mapping a finite dimensional vector space \mathbb{R}^n with elements $x = [x_1 \ x_2 \ \dots \ x_n]^T$ to either the set \mathbb{H} or the set \mathbb{S} of symmetric matrices, i.e.,*

$$F(x) := F_0 + F_1x_1 + \dots + F_nx_n.$$

Note that for general control applications, as in this thesis, the linear matrix inequalities are described by functions of matrix variables $X \in \mathbb{R}^{n_1 \times n_2}$ rather than scalar design variables x . Of course these matrix functions $F(X)$ can again be rewritten as functions of scalar design variables and the number then depends on the dimensionality of the matrix X and the structure imposed on it.

We consider some properties of these matrix inequalities first. Let $M \in \mathbb{H}^n$ and T be nonsingular matrix then

$$M \prec 0 \Leftrightarrow T^\dagger M T \prec 0$$

Proposition 5.16 (Schur complement [Scherer and Weiland, 2000]) Let $F : \mathbb{X}^m \rightarrow \mathbb{H}^n$ be an affine function which is partitioned according to

$$F(x) = \begin{bmatrix} F_{11}(x) & F_{12}(x) \\ F_{21}(x) & F_{22}(x) \end{bmatrix} \quad (5.63)$$

where $F_{11}(x)$ is square. Then following statements are equivalent

- $F(x) \succ 0$;
- $F_{11}(x) \succ 0$, and $F_{22}(x) - F_{21}(x)(F_{11}(x))^{-1}F_{12}(x) \succ 0$;
- $F_{22}(x) \succ 0$, and $F_{11}(x) - F_{12}(x)(F_{22}(x))^{-1}F_{21}(x) \succ 0$.

Part II

Verification

Errors using inadequate data are much less than those using no data at all.

Charles Babbage

6

Data-driven and model-based verification

via Bayesian inference and reachability analysis

In this chapter a measurement-driven and model-based formal verification approach, applicable to dynamical systems with partly unknown dynamics, is developed. We provide a new principled method, grounded on Bayesian inference and on reachability analysis respectively, to compute the confidence that a physical system driven by external inputs and accessed under noisy measurements verifies a given property expressed as a temporal logic formula. A case study discusses the bounded- and unbounded-time safety verification of a partly unknown system, encompassed within a class of linear time-invariant dynamical models with inputs and output measurements.

6.1 Introduction

The strength of formal techniques, such as model checking, is bound to the fundamental requirement of having access to a given model, obtained from the knowledge of the behaviour of the underlying system of interest. In practice, for most physical systems the dynamical behaviour is known only in part: this holds in particular for biological systems [Abate et al., 2012] or for classes of engineered

systems where, as a consequence, the use of uncertain control models built from data is a common practice [Hjalmarsson, 2005].

Only limited work within the formal methods community deals with the verification of models with partly unknown dynamics. Classical results [Batt et al., 2007, Henzinger and Wong-Toi, 1996] consider verification problems for non-stochastic models described by differential equations with bounded parametric uncertainty. Similarly, but for continuous-time *probabilistic* models, [Bortolussi and Sanguinetti, 2014, Brim et al., 2013] explore the parameter space with the objective of model verification (respectively statistical or probabilistic). Whenever full state measurements of the system are available, Statistical Model Checking (SMC) [Sen et al., 2004b, Legay et al., 2010] replaces numerical model-based procedures with empirical testing of formalised properties. SMC is limited to fully observable stochastic systems with little or no non-determinism, and may require the gathering a large set of measurements. Extensions towards the inclusion of non-determinism have been studied in [Henriques et al., 2012, Legay and Sedwards, 2013], with preliminary steps towards Markov decision processes. Related to SMC techniques, but bound to finite state models, [Chen and Nielsen, 2012, Mao and Jaeger, 2012, Sen et al., 2004a] assume that the system is encompassed by a finite-state Markov chain and efficiently use data to learn the corresponding model and to verify it. Similarly, [Bartocci et al., 2013, Bortolussi and Sanguinetti, 2013] employ machine learning techniques to infer finite-state Markov models from data over given logical formulae.

An alternative approach, allowing both partly unknown dynamics over uncountable (continuous) variables and noisy output measurements, is the usage of a Bayesian framework relating the confidence in a formal property to the uncertainty of a model built from data. When applied on nonlinearly parameterised, linear time invariant (LTI) models this approach introduces heavy computational issues, which can only be mitigated via statistical methods [Gyori et al., 2014]. Instead, in order to obtain reliable and numerical solutions, we propose the use of linearly parameterised model sets defined through orthonormal basis functions to represent these partially unknown systems. This is a broadly used framework in system identification [Heuberger et al., 2005, Hjalmarsson, 2005]: while maintaining the beneficial computational aspects of linear parameterisations, the choice of orthonormal basis functions allows for the incorporation of prior knowledge on the system behaviour. Practically, this has been widely used for the modelling of physical systems, such as the thermal dynamics of buildings [Virk and Loveday, 1994].

This work investigates the verification of temporal logic properties over partially unknown systems, using both prior modelling knowledge and data drawn from the system in a Bayesian setting. Building on [Haesaert et al., 2015c,b], we provide a complete framework and newly extend the modelling class in [Haesaert et al., 2015c] to multi-input multi-output models. The core of this work is further set apart from [Haesaert et al., 2015b], which focused on the use of data as experiments to ameliorate the verification procedure – the optimal design of experiments to efficiently decide on quantitative system properties will be explored in the next chapter.

The chapter is structured as follows. Section 6.2 grounds the problem on models and notations. Section 6.3 focuses on the model-based property verification step. Extensions are discussed in Section 6.4.

6.2 General framework and problem statement

In this section we overview a new methodology to assess the confidence in whether a system \mathbf{S} satisfies a given specification ψ , formulated in a suitable temporal logic, by integrating the partial knowledge of the system dynamics with data obtained from a measurement setup around the system.

Let us further clarify this framework. Let us denote with \mathbf{S} a physical system, or equivalently its associated dynamical behaviour. A signal input $u(t) \in \mathbb{U}, t \in \mathbb{N}$, captures how the environment acts on the system. Similarly, an output signal $y_0(t) \in \mathbb{Y}$ indicates how the system interacts with the environment, or alternatively how the system can be measured. Note that the input and output signals are assumed to take values over continuous domains. The system dynamics can be described via mathematical models, which quantify the behavioural relation between its inputs and outputs. The knowledge of the behaviour of the system is often limited or uncertain, making it impossible to analyse its dynamics by means of a “true” model. In this case, a-priori available knowledge allows to construct a model set \mathcal{G} with elements $\mathbf{M} \in \mathcal{G}$: this model class encompasses the uncertainty on the underlying system by means of a parameterisation $\theta \in \Theta$, $\mathcal{G} = \{\mathbf{M}(\theta) | \theta \in \Theta\}$. The unknown “true” model $\mathbf{M}(\theta^0)$ representing \mathbf{S} , is assumed to be an element of \mathcal{G} , namely $\theta^0 \in \Theta$. Model sets \mathcal{G} obtained through first principles and with unknown parameters adhere to this standard setup.

Samples can be drawn from the underlying physical system via a measurement setup, as depicted in Figure 6.1. An experiment consists of a finite number (N_s) of input-output samples drawn from the system, and is denoted by

$$Z^{N_s} = \{u(t)_{ex}, \tilde{y}(t)_{ex}\}_{t=1}^{N_s},$$

where $u(t)_{ex} \in \mathbb{U}$ (in general a continuous domain) is the input for the experiment and $\tilde{y}(t)_{ex}$ is a (possibly noisy) measurement of $y_0(t)_{ex}$. In general, the measurement noise can enter non-additively and be a realisation of a stationary stochastic process.¹ We assume that at the beginning of the measurement procedure (say at $t = 0$), the initial condition of the system, encompassed by the initial state $x(0)_{ex}$ of models in \mathcal{G} , is either known, or, when not known, has a structured uncertainty distribution that is based on the knowledge of past inputs and/or outputs. As reasonable, we implicitly consider only well-defined problems, such that for any model $\mathbf{M}(\theta)$ representing the system, given an input signal $u(t)_{ex}$ and an (uncertainty distribution for) $x(0)_{ex}$, the probability density distribution of the measured signal can be fully characterised.

¹Notice that the operating conditions of the experiment, that is the input signal $u(t)_{ex}$, the initial state $x(0)_{ex}$, and the measurements $\tilde{y}(t)_{ex}$, have been indexed with “*ex*” to distinguish them from the conditions of interest for verification (“*ver*”), to be discussed shortly.

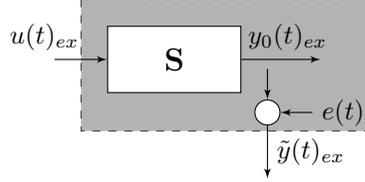


Figure 6.1: System (smaller white box) and measurement setup (grey box). In the measurement setup the output $\tilde{y}(t)_{ex}$ includes the system output $y_0(t)_{ex}$ and the measurement noise $\epsilon(t)$. Data collected from experiments comprises the input $u(t)_{ex}$ and the measured output $\tilde{y}(t)_{ex}$ signals.

The end objective is to analyse the behaviour of system S . We consider properties encoded as specifications ψ and expressed in a temporal logic of choice (to be detailed shortly). Let us remark that the behaviour of S to be analysed is bound to a set of operating conditions that are pertinent to the verification problem and that will be indexed by “*ver*”: this comprises the set of possible input signals $u(t)_{ver}$ (e.g., a white or coloured noise signal, or a non-deterministic signal $u(t)_{ver} \in \mathbb{U}_{ver} \subseteq \mathbb{U}$), and of the set of initial states $x(0)_{ver} \in \mathbb{X}_{ver}$ for the mathematical models $M(\theta)$ reflecting past inputs and/or outputs of the system. The system satisfies a property if the “true” model representing it satisfies it, namely $S \models \psi$ if and only if $M(\theta^0) \models \psi$.

In this chapter we consider the satisfaction of a property $M(\theta) \models \psi$ as a *binary-valued mapping* from the parameter space Θ . More generally, when in addition to the measurements of the system also its internal transitions are disturbed by stochastic noise (known as process noise), then property satisfaction is a mapping from the parameter space Θ to the interval $[0, 1]$, and quantifies the probability that the model $M(\theta)$ satisfies the property. This mapping generalises the definition of the satisfaction function discussed in [Bortolussi and Sanguinetti, 2014], and is now stated as follows.

Definition 6.1 (Satisfaction Function) *Let \mathcal{G} be a set of models M that is indexed by a parameter $\theta \in \Theta$, and let ψ be a formula in a suitable temporal logic. The satisfaction function $f_\psi : \Theta \rightarrow [0, 1]$ associated with ψ is*

$$f_\psi(\theta) = \mathbf{P}(M(\theta) \models \psi). \quad (6.1)$$

Let us assume that the satisfaction function f_ψ is measurable and entails a decidable verification problem (e.g., a model checking procedure) for all $\theta \in \Theta$ and properties ψ of interest. In this chapter we consider the verification of partly unknown physical systems with respect to a subset of linear time temporal logic properties. We are in a position to state the following.

Problem Description 6.1 *For a partly unknown physical system S , under prior knowledge on the system given as a parameterised model class \mathcal{G} supporting an uncertainty*

distribution over the parameterisation, gather possibly noisy data drawn from the measurement setup and verify properties on \mathbf{S} expressed in a temporal logic of choice, with a formal quantification of the confidence of the assertion.

6.2.a A Bayesian framework for data-driven modelling and verification

Consider Problem 6.1. Denote with $\mathbb{P}(\cdot)$ and $p(\cdot)$ respectively a probability measure and a probability density function, both defined over a continuous domain. We employ Bayesian probability calculus [Lindley, 2000] to express the confidence in a property as a measure of the uncertainty distribution defined over the set \mathcal{G} . By adopting the Bayesian framework, uncertainty distributions are handled as probability distributions of random variables. Therefore the confidence in a property is computed as a probability measure $\mathbb{P}(\cdot)$ via the densities $p(\cdot)$ over the uncertain variables.

Proposition 6.2 (Bayesian Confidence) *Given a specification ψ and a data set Z^{N_s} , the confidence that $\mathbf{S} \models \psi$ can be quantified via inference as*

$$\mathbb{P}(\mathbf{S} \models \psi \mid Z^{N_s}) = \int_{\Theta} f_{\psi}(\theta) p(\theta \mid Z^{N_s}) d\theta, \quad (6.2)$$

where f_{ψ} is the satisfaction function given in (6.1). The a-posteriori uncertainty distribution $p(\theta \mid Z^{N_s})$, given the data set Z^{N_s} , is based on parametric inference over θ as

$$p(\theta \mid Z^{N_s}) = \frac{p(Z^{N_s} \mid \theta) p(\theta)}{\int_{\Theta} p(Z^{N_s} \mid \theta) p(\theta) d\theta}, \quad (6.3)$$

which assumes the knowledge of an uncertainty distribution $p(\theta)$ over the parameter set Θ , representing prior knowledge.

The statement can be formally derived based on standard Bayesian calculus, as in [Lindley, 2000]. We have chosen to employ a Bayesian framework, as per (6.3), since it allows to reason explicitly over the uncertain knowledge on the system and to work with the data acquired from the measurement setup. This leads to the efficient incorporation of the available knowledge and to its combination with the data acquisition procedure, in order to compute the confidence on the validity of a given specification over the underlying system. As a special instance, this result can be employed for Bayesian hypothesis testing [Zuliani et al., 2013]. As long as the mapping f_{ψ} is measurable, the models in the model set (and hence the system represented by it) can be characterised by either probabilistic or non-probabilistic dynamics.

Remark 6.1 *In statistical model checking [Legay et al., 2010, Sen et al., 2004b], the objective is to replace the computationally tolling verification of a system over bounded-time properties by the empirical (statistical) testing of the relevant specifications over finite executions drawn from the system. In contrast, our setup tackles the problem of efficiently*

incorporating data with prior knowledge, for the formal (deductive) verification of the behaviour of a system with partly unknown dynamics. As such our overall verification approach is, as claimed, both data-driven and model-based. Moreover, by separating the operational conditions of an experiment from those of importance for the verification procedure, the system can be verified over non-deterministic quantities, encompassing both controller and disturbance inputs, as well as modelling errors.

6.2.b Existing computational approaches

In the literature the satisfaction function is related to the exploration of a parameter set over the validity of a formal property $f_\psi(\theta)$, and has been studied for autonomous models in continuous time in [Batt et al., 2007, Frehse et al., 2008, Henzinger and Wong-Toi, 1996].

Bayesian inference is widely applicable to different types of properties and models, however its computational complexity might in practice limit its implementation. Analytical solutions to the inference equation (6.3) can be found if the prior is a conjugate distribution. For linear dynamical systems, closed-form solutions are given *inter alia* in [Peterka, 1981].

In general (6.2)-(6.3) in Proposition 6.2 lack analytical solutions, and the assessment of the satisfaction function (6.1) may be computationally intensive. Statistical methods such as the one proposed in [Gyori et al., 2014] on a similar Bayesian approach lead to involved computations and introduce additional uncertainty from Monte Carlo techniques.

In contrast with the reviewed literature, in the next section we propose a novel computational approach over discrete-time linear time-invariant systems. By exploiting linear parameterisations, analytical solutions of both the parameter inference and the satisfaction function are characterised, over properties expressed within a fragment of a temporal logic.

6.3 LTL verification of LTI systems

Consider a system \mathbf{S} that can be represented by a class of finite-dimensional dynamical models that evolve in discrete-time, and are linear time-invariant (LTI). We focus the study to non probabilistic dynamics. These models depend on input and output signals ranging over \mathbb{R}^m and \mathbb{R}^p , respectively, and on variables $x_{\mathbf{S}}(t)$ taking values in an Euclidean space, $x_{\mathbf{S}}(t) \in \mathbb{X} \subseteq \mathbb{R}^n$, where n , the state dimension, is the model order. The behaviour of such a system is encompassed by state-space models $(A_{\mathbf{S}}, B_{\mathbf{S}}, C_{\mathbf{S}}, D_{\mathbf{S}})$ as

$$\mathbf{S} : \begin{cases} x_{\mathbf{S}}(t+1) &= A_{\mathbf{S}}x_{\mathbf{S}}(t) + B_{\mathbf{S}}u(t), \\ y_{\mathbf{S}}(t) &= C_{\mathbf{S}}x_{\mathbf{S}}(t) + D_{\mathbf{S}}u(t), \end{cases} \quad (6.4)$$

where matrices $A_{\mathbf{S}}, B_{\mathbf{S}}, C_{\mathbf{S}}, D_{\mathbf{S}}$ are of appropriate dimensions. The experimental measurement setup, as depicted in Figure 6.1, consists of the signals $u(t)_{ex}$ and

$\tilde{y}(t)_{ex} = y_0(t)_{ex} + e(t)$, representing the inputs and the measured outputs, respectively, and where $e(t)$ is an additive zero-mean, white, Gaussian-distributed measurement noise with covariance Σ_e that is uncorrelated from the inputs. N_s samples are collected within a data set $Z^{N_s} = \{u(t)_{ex}, \tilde{y}(t)_{ex}\}_{t=1}^{N_s}$.

6.3.a Formalisation of properties

System properties are expressed, over a finite set of atomic propositions $p_i \in AP$, $i = 1, \dots, |AP|$, in Linear Temporal Logic [Baier and Katoen, 2008]. Any LTL formula ψ is built up recursively via the syntax

$$\psi ::= \text{true} \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc\psi \mid \psi \text{ U } \psi.$$

Let $\pi = \pi(0), \pi(1), \pi(2), \dots \in \Sigma^{\mathbb{N}^+}$ be a string composed of letters from the alphabet $\Sigma = 2^{AP}$, and let $\pi_t = \pi(t), \pi(t+1), \pi(t+2), \dots$ be a subsequence (postfix) of π . The satisfaction relation between π and ψ is denoted as $\pi \models \psi$ (or equivalently $\pi_0 \models \psi$). The semantics of the satisfaction relation are defined recursively over π_t and the syntax of the LTL formula ψ as follows:

$$\begin{aligned} (\text{true}) \pi_t \models \text{true} &\Leftrightarrow \text{true} \\ (\text{atomic prop.}) \pi_t \models p &\Leftrightarrow p \in \pi(t) \\ (\text{negation}) \pi_t \models \neg\psi &\Leftrightarrow \pi_t \not\models \psi \\ (\text{conjunction}) \pi_t \models \psi_1 \wedge \psi_2 &\Leftrightarrow \pi_t \models \psi_1 \text{ and } \pi_t \models \psi_2 \\ (\text{next}) \pi_t \models \bigcirc\psi &\Leftrightarrow \pi_{t+1} \models \psi \\ (\text{until}) \pi_t \models \psi_1 \text{ U } \psi_2 &\Leftrightarrow \exists i \in \mathbb{N} : \pi_{t+i} \models \psi_2, \\ &\text{and } \forall j \in \mathbb{N} : \\ &0 \leq j < i, \pi_{t+j} \models \psi_1 \end{aligned}$$

This syntax allows to extend the study to more complex propositional formulae (such as disjunction or implication). Denote the k -bounded and unbounded invariance (or safety) operator as $\square^k\psi = \bigwedge_{i=0}^k \bigcirc^i\psi$ and $\square\psi = \neg(\text{true U } \neg\psi)$, respectively.

It is of interest to refer formal properties expressed as LTL formulae to the input-output behaviour of a dynamical model, over a given time horizon $t \geq 0$. The output $y_0(t)_{ver} \in \mathbb{Y}$ is labeled by a map $L : \mathbb{Y} \rightarrow \Sigma$, which assigns symbols α in the alphabet Σ of the formulae discussed previously to half spaces on the output, as

$$L(y_0(t)_{ver}) = \alpha \in \Sigma \Leftrightarrow \bigwedge_{p_i \in \alpha} A_{p_i} y_0(t)_{ver} \leq b_{p_i}, \quad (6.5)$$

for given $A_{p_i} \in \mathbb{R}^{1 \times p}$, $b_{p_i} \in \mathbb{R}$. In other words, sets of atomic propositions in AP are associated to polyhedra over $\mathbb{Y} \subset \mathbb{R}^p$. Let us underline that properties are defined over the behaviour $y_0(t)_{ver}$ of the model, and not over the noisy measurements $\tilde{y}(t)_{ex}$ of the model considered within the measurement setup. Additionally, for the verification problem the input signal is modelled as a bounded signal

$u(t) \in \mathbb{U}_{ver}$, and represents external non-determinism from the environment acting on the system.

6.3.b Model set selection

As a first step we need to embed the available a-priori knowledge on the underlying system within a parameterised model set. Note that although the goal of parameter exploration in formal verification has recently attracted quite some attention [Batt et al., 2007, Frehse et al., 2008, Henzinger and Wong-Toi, 1996], there are as of yet no general scalable results for the computation of the satisfaction function for nonlinearly-parameterised, discrete-time LTI models. The use of linearly-parameterised model sets, especially those defined through orthonormal basis functions (as further elaborated next), has been widely used for the modelling of physical systems, such as the thermal dynamics of buildings [Reginato et al., 2009, Virk and Loveday, 1994].

Whilst in general the uncertainty about a model representing a linear time-invariant system does not map onto a linearly-parameterised model set, we argue that a linearly-parameterised model set can encompass a relevant class of models. For instance, any asymptotically stable LTI model can be represented uniquely by its (infinite) impulse response [Heuberger et al., 1995], and the coefficients of the impulse response define a linear parameterisation for this model. Further, for asymptotically stable systems, the coefficients of the impulse response converge to zero, so that a truncated set of impulse coefficients provide a good approximate LTI model set with a finite-dimensional, linear parameterisation. These impulse responses define a finite set of orthonormal basis functions [Heuberger et al., 2005, Chapters 4 and 7],[Van den Hof et al., 1995] and construct a valid model set for a physical system solely based on knowledge of asymptotic stability. Alternative choices for an orthonormal basis such as Laguerre functions and Kautz functions [Heuberger et al., 2005], can incorporate additional and more extensive prior knowledge of the physical system.

We conclude that, as an alternative to the use of a nonlinearly parameterised set of models, structural information (even when not exact) can be used to select a set of orthonormal basis functions, whose finite truncation defines a finite-dimensional linearly-parameterised model set indexed over the coefficients of the basis functions. Thus, in the following we consider a linearly parameterised model set \mathcal{G} that encapsulates system \mathbf{S} , and specifically $\mathcal{G} = \{(A, B, C(\theta), D(\theta)), \theta \in \Theta\}$.

A system satisfies a property if, assuming it can be equivalently represented by a mathematical model $\mathbf{M}(\theta^0)$, all the words generated by the model satisfy that property. Since properties are encoded over the external (input-output) behaviour of the system \mathbf{S} , which is the behaviour of $\mathbf{M}(\theta^0)$ (where in our case $\theta^0 \in \Theta$), we may equivalently assert that any property ψ is verified by the system, $\mathbf{S} \models \psi$, if and only if it is verified by the unknown model representing the system, namely $\mathbf{M}(\theta^0) \models \psi$. Within the modelling perspective offered in this work, let us introduce Θ_ψ to be the feasible set of parameters, such that for every parameter in that set the property ψ holds, i.e., $\forall \theta \in \Theta_\psi : \mathbf{M}(\theta) \models \psi$. As such Θ_ψ is characterised

as the level set of the satisfaction function f_ψ , $\Theta_\psi = \{\theta \in \Theta : f_\psi(\theta) = 1\}$. The quantification of Θ_ψ is of key importance in our work.

6.3.c Safety verification of bounded-time properties

Models \mathbf{M} in the class \mathcal{G} have the following representation $(A, B, C(\theta), 0)$:

$$\mathbf{M}(\theta) : \begin{cases} x(t+1) &= Ax(t) + Bu(t), \\ \hat{y}(t, \theta) &= C(\theta)x(t), \end{cases} \quad (6.6)$$

and are parameterised by $\theta \in \Theta \subset \mathbb{R}^{pm}$, $\theta = \text{vec}(C)$ and $C(\theta) \in \mathbb{R}^{p \times n}$. We assume a prior probability distribution $p(\theta)$, which structures the knowledge of the uncertainty in θ . In addition to this *strictly proper* model class we will also allow for a *proper* model class $(A, B, C(\theta), D(\theta))$, where both the C and the D -matrices are parameterised, so that $\theta = \text{vec}([C \ D])$. For a given initial condition $x(0)$ and input sequence, the output of the “true” model $\hat{y}(t, \theta^0)$ is equal to the system output $y_0(t)$.

Consider a measurement setup as in Figure 6.1, related to an unknown parameter θ^0 . Signals $u(t)_{ex}$ and $\tilde{y}(t)_{ex}$ represent the input and the measured output, respectively, and $e(t)$ is an additive zero-mean, white, Gaussian-distributed measurement noise with covariance Σ_e that is uncorrelated from the input. From this setup N_s samples are collected in a data set $Z^{N_s} = \{u(t)_{ex}, \tilde{y}(t)_{ex}\}_{t=1}^{N_s}$. Given the operating conditions of the experiment setup, the measured signal $\tilde{y}(t)_{ex}$ can be fully characterised: its probability density, conditional on the parameters θ , is

$$\begin{aligned} p(Z^{N_s} | \theta) &= \prod_{t=1}^{N_s} p(\tilde{y}(t)_{ex} | \theta) \\ &= \frac{1}{\sqrt{|\Sigma_e|^{N_s} (2\pi)^{pN_s}}} \exp \left[-\frac{1}{2} \sum_{t=1}^{N_s} (\hat{y}(t, \theta) - \tilde{y}(t)_{ex})^T \Sigma_e^{-1} (\hat{y}(t, \theta) - \tilde{y}(t)_{ex}) \right], \end{aligned}$$

and can be directly used in Proposition 6.2. This conditional density $p(Z^{N_s} | \theta)$ depends implicitly on the given initial state $x(0)_{ex}$ and, in the case of a given uncertainty distribution over $x(0)_{ex}$, $p(Z^{N_s} | \theta)$ should be marginalised over $x(0)_{ex}$ [Peterka, 1981]. The a-posteriori uncertainty distribution is obtained as the analytical solution of the parametric inference in (6.3) [Peterka, 1981].

Recall now that for a given specification ψ , we seek to determine a feasible set of parameters Θ_ψ , which is such that the corresponding models admit property ψ , namely $\mathbf{M}(\theta) \models \psi, \forall \theta \in \Theta_\psi$. Since models $\mathbf{M}(\theta)$ have a linearly-parameterised state space realisation as per (6.6), it follows that when the set of initial states \mathbb{X}_{ver} and of inputs \mathbb{U}_{ver} are bounded polyhedra, the verification of a class of safety properties expressed by formulae with labels as in (6.5) leads to a set of feasible parameters Θ_ψ that is a polyhedron, which can be easily computed. More precisely, the following result can be derived.

Theorem 6.3 Consider properties ψ composed within the LTL fragment $\psi ::= \alpha | \bigcirc \psi | \psi_1 \wedge \psi_2$, with $\alpha \in \Sigma$. Given a bounded polyhedral set (a polytope) of initial states $x(0) \in \mathbb{X}_{ver}$ and of inputs $u(t) \in \mathbb{U}_{ver}$ for $0 \leq t < \infty$, and considering a labelling map as in (6.5), then the feasible set Θ_ψ of the parameterised model set (6.6) is a polyhedron.

Proof: [of Theorem 6.3] Let \otimes denote the Kronecker product. Consider the input set \mathbb{U}_{ver} to be the convex hull of U , i.e. $\text{conv}(U) = \mathbb{U}_{ver}$. Similarly let the set of initial states be $\text{conv}(X_{ver}) = \mathbb{X}_{ver}$. Let the model set be given as $\mathbf{M}(\theta) = (A, B, C(\theta), D)$. We will temporarily assume that D is set to be equal to zero, and afterwards (cf. Point 3) we will show how to work with a parameterised D . As can be deduced from the *and* operations in (6.5), note that for simplicity the syntax fragment $\psi ::= \alpha | \bigcirc \psi | \psi_1 \wedge \psi_2$ with $\alpha \in \Sigma = 2^{AP}$ is equivalent to $\psi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2$ with $p \in AP$. We structure the proof in three parts.

1. We claim that for every specification ψ composed from the syntax fragment $\psi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2$ and $\theta \in \Theta$, the words generated by a model $\mathbf{M}(\theta) = (A, B, C(\theta), 0)$ with state $x(t)$ satisfy the specification ψ , denoted $\langle \mathbf{M}(\theta), x(t) \rangle \models \psi$, if and only if

$$\left((I_{n_\psi} \otimes x(t))^T N_\psi + K_\psi \right) \theta \leq B_\psi. \quad (6.7)$$

The matrices $N_\psi \in \mathbb{R}^{n_\psi \times np}$, $K_\psi \in \mathbb{R}^{n_\psi \times np}$, $B_\psi \in \mathbb{R}^{n_\psi}$ in the above satisfaction relation have dimensions that are functions of the parametrisation and of a property-dependent “dimension” n_ψ , which will be obtained inductively over the syntax of the specification. Next we focus the study of fragments of the LTL syntax.

For any *atomic proposition* the model starting from state $x(t)$ satisfies a property p_i , i.e., $\langle \mathbf{M}(\theta), x(t) \rangle \models p_i \Leftrightarrow A_{p_i} y \leq b_{p_i}$, with $A_{p_i} \in \mathbb{R}^{1 \times p}$ and $b_{p_i} \in \mathbb{R}$ we construct the matrices N_{p_i} , K_{p_i} and B_{p_i} as follows. Consider $y(t)$ for a given $x(t)$ then

$$A_{p_i} y(t) = A_{p_i} C(\theta) x(t) = x(t)^T (I_n \otimes A_{p_i}) \theta.$$

This yields $N_{p_i} = (I_n \otimes A_{p_i}) \in \mathbb{R}^{n \times np}$, $K_{p_i} = O_{1 \times np} \in \mathbb{R}^{1 \times np}$, and $B_{p_i} = b_{p_i} \in \mathbb{R}^{1 \times 1}$.

The *next* operation $\bigcirc \psi_1$ with matrices $(N_{\psi_1}, K_{\psi_1}, b_{\psi_1})$ yields matrices

$$\begin{aligned} N_{\bigcirc \psi_1} &= \mathbf{1}_{|U|} \otimes (I_{n_{\psi_1}} \otimes A^T) N_{\psi_1}, \\ K_{\bigcirc \psi_1} &= \mathcal{U} (I_{n_{\psi_1}} \otimes B)^T N_{\psi_1} + \mathbf{1}_{|U|} \otimes K_{\psi_1}, \\ B_{\bigcirc \psi_1} &= \mathbf{1}_{|U|} \otimes B_{\psi_1}, \end{aligned}$$

where the i -th set of n_{ψ_1} rows of $\mathcal{U} \in \mathbb{R}^{|U| n_{\psi_1} \times m}$ is defined as

$$(I_{n_{\psi_1}} \otimes u_i^T) \text{ with } u_i \in U$$

and where $n_{\bigcirc \psi_1} = |U| n_{\psi_1}$. This can be derived as

$$\begin{aligned} \langle \mathbf{M}(\theta), x(t) \rangle \models \bigcirc \psi &\Leftrightarrow \forall u(t) \in \mathbb{U}_{ver} : \left((I_{n_{\psi_1}} \otimes x(t+1))^T N_{\psi_1} + K_{\psi_1} \right) \theta \leq B_{\psi_1}, \\ &\Leftrightarrow \forall u(t) \in \mathbb{U}_{ver} : \left((I_{n_{\psi_1}} \otimes Ax(t))^T N_{\psi_1} + (I_{n_{\psi_1}} \otimes Bu(t))^T N_{\psi_1} + K_{\psi_1} \right) \theta \leq B_{\psi_1}. \end{aligned}$$

Since the above is an affine function in $u(t)$, the image of every $u(t) \in \text{conv}(U) = \mathbb{U}_{ver}$ can be expressed as a convex combination of the values at the vertices $u_i \in U$, c.f. [Belta et al., 2002]. Then an equivalent expression is

$$\Leftrightarrow \forall u_i \in U : \left((I_{n_{\psi_1}} \otimes Ax(t))^T N_{\psi_1} + (I_{n_{\psi_1}} \otimes u_i)^T (I_{n_{\psi_1}} \otimes B)^T N_{\psi_1} + K_{\psi_1} \right) \theta \leq B_{\psi_1},$$

which can be rewritten as

$$\Leftrightarrow \left(\mathbf{1}_{|U|} \otimes (I_{n_{\psi_1}} \otimes Ax(t))^T N_{\psi_1} + \mathcal{U} (I_{n_{\psi_1}} \otimes B)^T N_{\psi_1} + \mathbf{1}_{|U|} \otimes K_{\psi_1} \right) \theta \leq \mathbf{1}_{|U|} \otimes B_{\psi_1}.$$

Matrices $K_{\circ\psi}$, and $B_{\circ\psi}$ can be obtained directly. To obtain $N_{\circ\psi}$ now rewrite the first term:

$$\begin{aligned} & \mathbf{1}_{|U|} \otimes (I_{n_{\psi_1}} \otimes x^T(t)) (I_{n_{\psi_1}} \otimes A^T) N_{\psi_1} \\ &= (I_{|U|} \mathbf{1}_{|U|}) \otimes (I_{n_{\psi_1}} \otimes x^T(t)) (I_{n_{\psi_1}} \otimes A^T) N_{\psi_1} \\ &= \left(I_{|U|n_{\psi_1}} \otimes x^T(t) \right) \left(\mathbf{1}_{|U|} \otimes (I_{n_{\psi_1}} \otimes A^T) N_{\psi_1} \right). \end{aligned}$$

The *and* operation $\psi_1 \wedge \psi_2$ for $(N_{\psi_1}, K_{\psi_1}, B_{\psi_1})$ and $(N_{\psi_2}, K_{\psi_2}, B_{\psi_2})$ with $n_{\psi_1 \wedge \psi_2} = (n_{\psi_1} + n_{\psi_2})$ gives

$$N_{\psi_1 \wedge \psi_2} = \begin{bmatrix} N_{\psi_1} \\ N_{\psi_2} \end{bmatrix}, K_{\psi_1 \wedge \psi_2} = \begin{bmatrix} K_{\psi_1} \\ K_{\psi_2} \end{bmatrix}, B_{\psi_1 \wedge \psi_2} = \begin{bmatrix} B_{\psi_1} \\ B_{\psi_2} \end{bmatrix}.$$

This can be derived from

$$\begin{aligned} \langle \mathbf{M}(\theta), x(t) \rangle \models \psi_1 \wedge \psi_2 &\Leftrightarrow \bigwedge_{i \in \{1,2\}} \left((I_{n_{\psi_i}} \otimes x(t))^T N_{\psi_i} + K_{\psi_i} \right) \theta \leq B_{\psi_i} \\ &\Leftrightarrow \left((I_{n_{\psi_1 \wedge \psi_2}} \otimes x(t))^T \begin{bmatrix} N_{\psi_1} \\ N_{\psi_2} \end{bmatrix} + \begin{bmatrix} K_{\psi_1} \\ K_{\psi_2} \end{bmatrix} \right) \theta \leq \begin{bmatrix} B_{\psi_1} \\ B_{\psi_2} \end{bmatrix}. \end{aligned}$$

2. The matrix-valued function

$$\left((I_{n_\psi} \otimes x(0))^T N_\psi + K_\psi \right) \theta$$

is affine in $x(0)$ (for a fixed θ), therefore its value at the initial condition $x(0) \in \mathbb{X}_{ver}$ is a convex combination of the function values at the vertices X_{ver} of \mathbb{X}_{ver} . Thus the satisfaction relation $\langle \mathbf{M}(\theta), x(0) \rangle \models \psi$ represented by the multi-affine inequality holds uniformly over $x(0) \in \mathbb{X}_{ver}$ if and only if it holds for the vertices of \mathbb{X}_{ver} .

This gives a set of affine inequalities in θ , thus the feasible set Θ_ψ is a polyhedron and is given as

$$\left\{ \theta \in \Theta : \bigwedge_{x_i \in X_{ver}} \left((I_{n_\psi} \otimes x_i)^T N_\psi + K_\psi \right) \theta \leq B_\psi \right\}.$$

Let us remark that set Θ_ψ is a polyhedron because it is formed by a finite set of

half spaces.

3. To complete the proof of Theorem 6.3 we need to extend the results to models with parameterised D . The dynamics of model (A, B, C, D) with both C and D fully parameterised can be reformulated as

$$\begin{aligned} \begin{bmatrix} x(t+1) \\ u(t+1) \end{bmatrix} &= \begin{bmatrix} A & B \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} + \begin{bmatrix} 0 \\ I \end{bmatrix} u(t+1) \\ y(t) &= [C \quad D] x(t). \end{aligned}$$

Using the new matrices $(\tilde{A}, \tilde{B}, \tilde{C}(\theta), 0)$ the obtained results still hold. For part 2. set of vertices X_{ver} needs to be extended with the vertices of U as $X_{ver} \times U$. \square

In the computation of the feasible set, the faces of the polyhedron Θ_ψ are shown to be a function of the vertices (recall that a polytope can be written as the convex hull of a *finite* set of vertices) of the bounded set of initial states \mathbb{X}_{ver} and of the set of inputs \mathbb{U}_{ver} , and are also expected to grow in number as a function of the time horizon of the property.

The result in Theorem 6.3 is valid for any finite composition of the LTL fragment $\psi ::= \alpha | \bigcirc \psi | \psi_1 \wedge \psi_2$, as such it only holds for finite horizon properties. Properties defined over the infinite horizon will be the objective of Section 6.3.e.

Remark 6.2 *The feasible set Θ_ψ obtained in Theorem 6.3 is a Borel-measurable set as it defines (if not empty) a closed set in the parameter space.*

6.3.d Case study: bounded-time safety verification

6.3.d.1 Single-input single output

Consider a system \mathbf{S} and verify whether the output $y_0(t)_{ver}$ remains within the interval $\mathcal{I} = [-0.5, 0.5]$, labeled as ι , for the next 5 time steps, under $u(t)_{ver} \in \mathbb{U}_{ver} = [-0.2, 0.2]$ and $x(0)_{ver} \in \{0_2\} = \mathbb{X}_{ver}$. Introduce accordingly the alphabet $\Sigma = \{\iota, \tau\}$ and the labelling map $L : L(y) = \iota, \forall y \in \mathcal{I}, L(y) = \tau, \forall y \in \mathbb{Y} \setminus \mathcal{I}$. Now check whether the following LTL property holds: $\mathbf{S} \models \bigwedge_{i=1}^5 (\bigcirc)^i \iota$.

We assume that system \mathbf{S} can be represented as an element of a model set \mathcal{G} , with models expressed via transfer functions characterised by second-order Laguerre-basis functions [Heuberger et al., 1995] (a special case of orthonormal basis functions). This translates to the following parameterised state-space representation:

$$\begin{aligned} x(t+1) &= \begin{bmatrix} a & 0 \\ 1-a^2 & a \end{bmatrix} x(t) + \begin{bmatrix} \sqrt{1-a^2} \\ (-a)\sqrt{1-a^2} \end{bmatrix} u(t), \\ \hat{y}(t, \theta) &= \theta^T x(t). \end{aligned} \tag{6.8}$$

The parameter set is chosen as $\theta \in \Theta = [-10, 10]^2$, whereas the coefficient a is chosen to be equal to 0.4. We select, as prior available knowledge on the system, a uniform distribution $p(\theta)$ on the model class, and pick a known variance $\sigma_\epsilon^2 = 0.5$

for the white additive noise on the measurement. The set of feasible parameters $\Theta_\psi \subset \Theta$ is represented in Figure 6.2 and is computed according to Theorem 6.3. Based on the prior available knowledge, the confidence associated to $\theta^0 \in \Theta_\psi$ amounts to 0.0165: this quantity is obtained by numerical computation of (6.2) with probability distribution² $p(\theta)$. Thereafter, we have set up an experiment on the system with “true parameter” $\theta^0 = [1 \ 0]^T$ (Figure 6.2) and with input signal $u(t)_{ex}$, a realisation of a white noise with a uniform distribution over $[-0.2, 0.2]$, and measured $\tilde{y}(t)_{ex}$ for 200 consecutive time instances. In comparison to the confidence obtained with the prior $p(\theta)$, the uncertainty distribution is now refined as $p(\theta|Z^{N_s})$, and the resulting confidence in the property is increased to 0.779, as per (6.2).

Along this line of experiments, we have repeated the test 100 times, for several instances of the parameter θ^0 characterising the underlying system S . In all instances, after obtaining 200 measurements the a-posteriori probability is used to assess the confidence in the safety of the system, as displayed in Table 6.1 via mean and variance terms. Observe that θ^0 is just outside of Θ_ψ for $[-1, -1]^T$ and for $[1, 1]^T$. For $\theta^0 = [-1, 1]^T$ and $\theta^0 = [1, -1]^T$, the true parameter lies in the feasible set but very close to the edges. This is reflected in the results in 6.1. For the points clearly inside the feasible set the confidence generally becomes high with low variance. Whereas for the points closest to the edge $\theta^0 = [-1, 1]^T$ and $\theta^0 = [1, -1]^T$, the variance is higher and the confidence has only increased up to around .49. In comparison, the points just outside the feasible edge give a lower confidence than the former two. The observed initial increase from .0165 to around .34 is expected and can be explained by the closeness to the feasible set; additional measurements will make the confidence converge to zero again. In conclusion, the experiments show that the measurements can be used to quantify the confidence level.

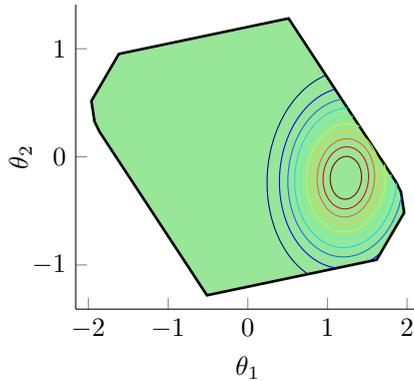


Figure 6.2: Feasible set of parameters $\Theta_\psi \subset \Theta$, and contour lines of the posterior $p(\theta|Z^{N_s})$, obtained for $\theta^0 = [1 \ 0]^T$ after 200 measurements.

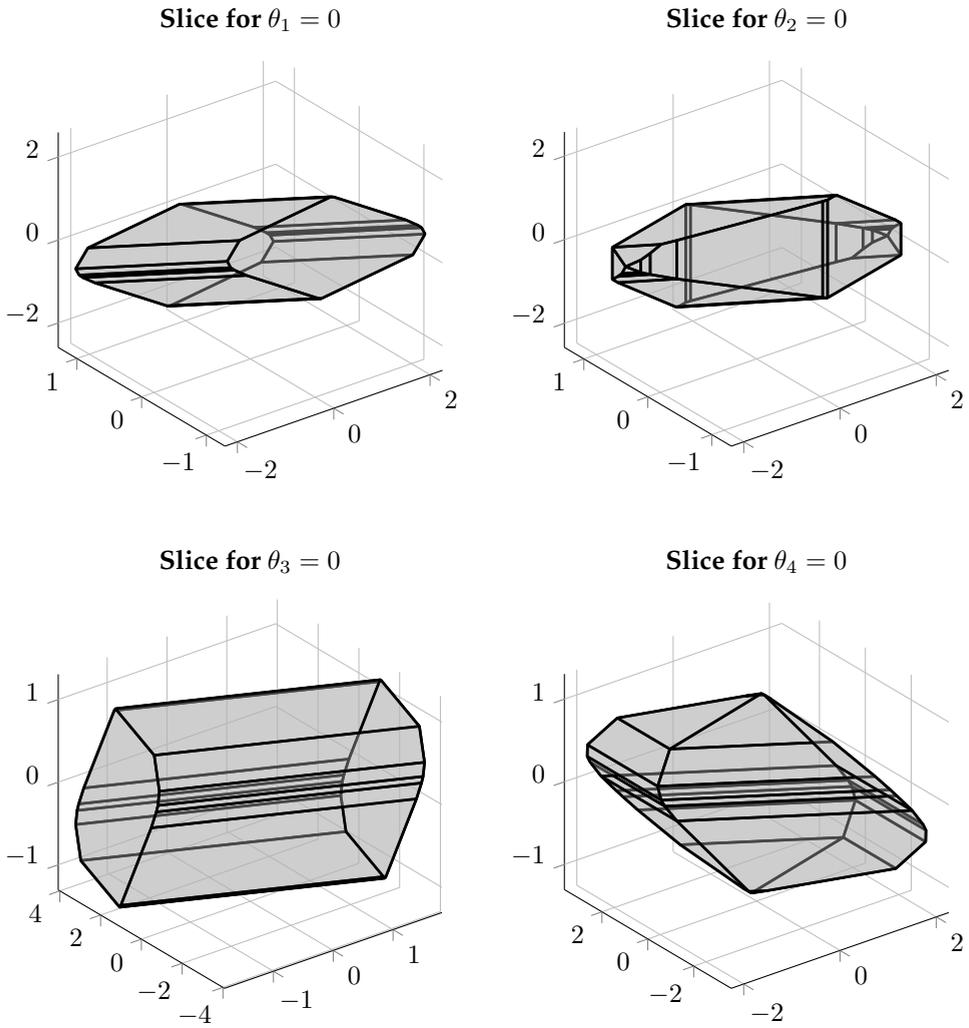


Figure 6.3: 3 dimensional plots of the 4 dimensional feasible set.

Table 6.1: Mean (μ) and variance (σ^2) of the confidence obtained from 100 experiments with 200 measurements each.

θ^0	μ	σ^2	θ^0	μ	σ^2
$[-1 \ -1]^T$	0.348	0.073	$[1 \ -1]^T$	0.491	0.085
$[-1 \ 0]^T$	0.705	0.060	$[1 \ 0]^T$	0.730	0.056
$[-1 \ 1]^T$	0.492	0.086	$[1 \ 1]^T$	0.339	0.065

6.3.d.2 Multiple-input multiple-output system: feasible set

To show case the workings of the feasible set computations for multiple-input multiple-output (MIMO) models we consider a straight forward extension of (6.8), that is

$$\begin{aligned} x(t+1) &= \begin{bmatrix} a & 0 \\ 1-a^2 & a \end{bmatrix} x(t) + \begin{bmatrix} \sqrt{1-a^2} & 0 \\ 0 & (-a)\sqrt{1-a^2} \end{bmatrix} u(t), \\ \hat{y}(t, \theta) &= \begin{bmatrix} \theta_1 & \theta_3 \\ \theta_2 & \theta_4 \end{bmatrix} x(t). \end{aligned} \quad (6.9)$$

We verify whether the output $y_0(t)$ remains within the polytope

$$\begin{bmatrix} -1 & 0 \\ -1 & -1 \\ 1 & 0 \\ 1 & 1 \end{bmatrix} y_0 \leq \begin{bmatrix} 0.5 \\ 0.5 \\ 0.5 \\ 0.5 \end{bmatrix}$$

for the next 5 time steps under $\mathbb{U}_{\text{ver}} = \text{conv} \left(\begin{bmatrix} -0.2 \\ 0 \end{bmatrix}, \begin{bmatrix} 0.2 \\ 0.1 \end{bmatrix}, \begin{bmatrix} 0.1 \\ -0.1 \end{bmatrix} \right)$ and $x(0)_{\text{ver}} \in \{0_2\}$. Using the results in the previous subsection can now compute the feasible set. In Figure 6.3 we display slices of the feasible set. For every slice 1 parameter θ_i is fixed to a value within the feasible set and the resulting set is plotted.

6.3.e Verifying unbounded-time properties using invariant sets

In this section we extend the approach of Section 6.3.c, to hold on the LTL fragment $\psi ::= \alpha | \bigcirc \psi | \psi_1 \wedge \psi_2$ with additionally the *unbounded* invariance (safety) operator. The subsection is built up as follows:

- first we connect the notion of positive invariance with that of feasible set;
- then we discuss how to practically compute a feasible set for invariance properties, with the set of initial states limited to be the origin;
- this is then extended to computing feasible sets under initial states in a polytope that includes the origin;

²Integrals are solved via the numerical integration tool in `Matlab`.

- finally, we interpret these results and complete the section with results on the verification of unbounded-time properties.

Recall the form of the k -bounded and of the unbounded invariance operators, namely $\Box^k \psi = \bigwedge_{i=0}^k \bigcirc^i \psi$ and $\Box \psi = \neg(\text{true} \cup \neg \psi)$, respectively. The extension from a k -bounded operator, covered by the result in Theorem 6.3, to the unbounded invariance one, is based on the concept of robust positive invariance [Blanchini and Miani, 2007, Def. 4.3], recalled next.

Definition 6.4 For the system $x(t+1) = Ax(t) + Bu(t)$, the set $S \subseteq \mathbb{X}$ is said to be robustly positively invariant if, for all $x(0) \in S$ and $u(t) \in \mathbb{U}_{ver}$, the condition $x(t) \in S$ holds for all $t \geq 0$.

Recall that the feasible set Θ_ψ is defined as the set of parameters for which property ψ holds, namely $\forall \theta \in \Theta_\psi : \mathbf{M}(\theta) \models \psi$. The satisfaction relation $\mathbf{M}(\theta) \models \psi$ depends implicitly on the set of initial states $x(0) \in \mathbb{X}_{ver}$ and on the set of inputs \mathbb{U}_{ver} . Let us extend the definition of the feasible set to explicitly account for its dependence on the set of initial conditions: given a bounded and convex set $S \subset \mathbb{X}$, let $\Theta_\psi(S)$ be defined as the set of parameters in Θ for which the parameterised models $\mathbf{M}(\theta)$ initialised with $x(0) \in S$ satisfy ψ over input signals $u(t) \in \mathbb{U}_{ver}$ $t \geq 0$. Hence the feasible set Θ_ψ can be written as a function of the set of initial states \mathbb{X}_{ver} , that is $\Theta_\psi(\mathbb{X}_{ver})$. Thus the extended map $\Theta_\psi(\cdot)$ takes subsets of the state space into subsets of the parameter space. Note that if S is a robustly positively invariant set that includes the set of initial states $\mathbb{X}_{ver} \subseteq S$, then for all $\theta \in \Theta_\psi(S)$ the models $\mathbf{M}(\theta)$ satisfy ψ over all infinite-time model traces $x(t)$: this allows to state that $\mathbf{M}(\theta) \models \Box \psi$. We can show that the following holds.

Lemma 6.5 The function $\Theta_\psi(\cdot) : 2^{\mathbb{X}} \rightarrow 2^\Theta$, for specifications obtained as $\psi ::= \alpha \mid \bigcirc \psi \mid \psi_1 \wedge \psi_2$, is monotonically decreasing: that is if $S_1 \subseteq S_2 \subseteq \mathbb{X}$, then $\Theta_\psi(S_2) \subseteq \Theta_\psi(S_1)$.

Proof: We leverage the notation used in the proof of Theorem 6.3. Provided that the parameterised model is given as $(A, B, C(\theta), 0)$, we show that any $\theta \in \Theta_\psi(S_2)$ is also an element of $\theta \in \Theta_\psi(S_1)$. Suppose S_2 has a finite number of vertices $x_i \in \mathcal{V}(S_2)$, then for any $\theta \in \Theta_\psi(S_2)$:

$$\bigwedge_{x_i \in \mathcal{V}(S_2)} ((I_{n_\psi} \otimes x_i)^T N_\psi + K_\psi) \theta \leq B_\psi,$$

and for every $x \in S_2$

$$((I_{n_\psi} \otimes x)^T N_\psi + K_\psi) \theta \leq B_\psi.$$

Since the vertices $x_j \in \mathcal{V}(S_1)$ are also elements of S_2 , then

$$\bigwedge_{x_j \in \mathcal{V}(S_1)} ((I_{n_\psi} \otimes x_j)^T N_\psi + K_\psi) \theta \leq B_\psi$$

and $\theta \in \Theta_\psi(S_1)$. This reasoning can be trivially extended to include models with parameterised D matrices. Increasing the number of vertices of S_1 and S_2 , does

not change the result, hence the same holds if S_1 and S_2 are convex sets. \square

Based on the result in Lemma 6.5, we conclude that the maximal feasible set $\Theta_{\square\psi}$ is obtained as a mapping from the minimal robustly positively invariant set S that includes \mathbb{X}_{ver} : $\Theta_{\square\psi} = \Theta_{\psi}(S)$. This leads next to consider under which conditions such minimal robustly positively invariant set S can be exactly computed or approximated.

Feasible set for invariance properties with $\mathbb{X}_{ver} = \{0_n\}$

For $\mathbb{X}_{ver} = \{0_n\}$, assuming a bounded interval \mathbb{U}_{ver} with the origin in its interior, and under some basic assumptions on the dynamics (to be shortly discussed), the minimal robustly positively invariant set can be shown to be a bounded and convex set that includes the origin [Blanchini and Miani, 2007]. Maintaining the condition of \mathbb{U}_{ver} being bounded and having the origin in its interior, we first consider the case that $\mathbb{X}_{ver} = \{0_n\}$ and characterise S via tools available from set theory in systems and control; thereafter we look at extensions to more general sets of initial states \mathbb{X}_{ver} .

Assume that \mathbb{U}_{ver} includes the origin, and denote the forward reachability mappings initialised with $\mathcal{R}^{(0)} := \{0_n\} \subset \mathbb{X}$ as

$$\mathcal{R}^{(i)} := \text{Post}(\mathcal{R}^{(i-1)}), \quad (6.10)$$

with set operation $\text{Post}(X) := \{x' = Ax + Bu, x \in X, u \in \mathbb{U}_{ver}\}$. Denote the limit reachable set as $\mathcal{R}^{\infty} = \lim_{i \rightarrow \infty} \mathcal{R}^{(i)}$. From literature we recall that properties of these i -step reachable sets, as given in [Blanchini and Miani, 2007] include the following: for a reachable pair (A, B) and an asymptotically stable matrix A , the ∞ -reachable set \mathcal{R}^{∞} is bounded and convex [Blanchini and Miani, 2007, Proposition 6.9]. Specifically, the k -step reachable set converges to the ∞ -reachable set via (6.10), since it is monotonically increasing $\mathcal{R}^{(i)} \subseteq \mathcal{R}^{(i+1)}$. Moreover, \mathcal{R}^{∞} is the minimal robustly positively invariant set for the system, so that any positively invariant set includes \mathcal{R}^{∞} [Blanchini and Miani, 2007, Proposition 6.13]. Thus, starting from $x(0) = 0_n$, all $x(t) \in \mathcal{R}^{\infty}$, and furthermore of interest to this work we conclude that $\Theta_{\square^k\psi} = \Theta_{\psi}(\mathcal{R}^{(k)})$ and that $\Theta_{\square\psi} = \Theta_{\psi}(\mathcal{R}^{\infty})$.

Feasible set for invariance properties under polytopic sets of initial states

More generally, if $\mathbb{X}_{ver} \subseteq \mathcal{R}^{\infty}$ and under the same assumptions on matrices A, B and $0 \in \mathbb{U}_{ver}$, then \mathcal{R}^{∞} is the minimal robustly positively invariant set that includes \mathbb{X}_{ver} , and $\Theta_{\psi}(\mathcal{R}^{\infty}) = \Theta_{\square\psi}$. For finite iterations the reachable sets $\mathcal{R}^{(i)}$ are polytopes, and if $\mathcal{R}^{(i)} = \mathcal{R}^{(i+1)}$, then $\mathcal{R}^{(i)} = \mathcal{R}^{\infty}$. Though the iterations can stop in finite time, in general the number of iterations to obtain \mathcal{R}^{∞} can be infinite. Whilst the minimal robustly positively invariant set is not necessarily closed or a polytope, there exist methods to approximate \mathcal{R}^{∞} as detailed in [Blanchini and Miani, 2007]. For instance, for stable systems, $\mathcal{R}^{(k)}$ is shown to converge to \mathcal{R}^{∞} , in the

sense that for all $\epsilon > 0$ there exists \bar{k} such that for $k \geq \bar{k}$, $\mathcal{R}^{(k)} \subseteq \mathcal{R}^\infty \subseteq (1 + \epsilon)\mathcal{R}^{(k)}$ [Blanchini and Miani, 2007, Proposition 6.9].

Recall that the maximal feasible set $\Theta_{\square\psi}$ is obtained as a mapping from the minimal robustly positively invariant set \mathcal{S} including \mathbb{X}_{ver} , so that $\Theta_{\square\psi} = \Theta_\psi(\mathcal{S})$. Let us extend the study to the case where the conditions $\mathbb{X}_{ver} = \{0_n\}$ or its extension $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty$ do not apply, while the condition on the bounded set \mathbb{U}_{ver} is maintained, that is $0 \in \mathbb{U}_{ver}$. Consider the more general case where the set of initial states is a polytope but not necessarily a subset of \mathcal{R}^∞ . Denote the union of the forward reachability mappings initialised with $\mathcal{R}_{\mathbb{X}_{ver}}^{(0)} := \mathbb{X}_{ver} \subseteq \mathbb{X}$ as

$$\mathcal{R}_{\mathbb{X}_{ver}}^{(i)} := \mathcal{R}_{\mathbb{X}_{ver}}^{(i-1)} \cup \text{Post}(\mathcal{R}_{\mathbb{X}_{ver}}^{(i-1)}). \quad (6.11)$$

This set is also known in the literature as the *reach tube*. The corresponding set for infinite time is denoted as $\mathcal{R}_{\mathbb{X}_{ver}}^\infty = \lim_{i \rightarrow \infty} \mathcal{R}_{\mathbb{X}_{ver}}^{(i)}$. Notice that in the earlier case when $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty$, then $\mathcal{R}^\infty = \mathcal{R}_{\mathbb{X}_{ver}}^\infty$. The iteration is monotonically increasing $\mathcal{R}_{\mathbb{X}_{ver}}^{(i)} \subseteq \mathcal{R}_{\mathbb{X}_{ver}}^{(i+1)}$, and whenever $\mathcal{R}_{\mathbb{X}_{ver}}^{(i)} = \mathcal{R}_{\mathbb{X}_{ver}}^{(i+1)}$ it stops after a finite number of iterations with $\mathcal{R}_{\mathbb{X}_{ver}}^\infty = \mathcal{R}_{\mathbb{X}_{ver}}^{(i)}$. Of course, also in this more general case, the number of iterations can be unbounded, however the convergence properties of $\mathcal{R}^{(i)}$ extend directly to the case of sets $\mathcal{R}_{\mathbb{X}_{ver}}^{(i)}$. Since $\mathcal{R}_{\mathbb{X}_{ver}}^{(i)}$ is a union of polytopes, it is not guaranteed to be a convex set. Still, it can be shown via arguments as in the proof of Theorem 6.3 that the computation of the feasible set $\Theta_\psi(\mathcal{S})$ boils down to that of $\Theta_\psi(\text{conv}(\mathcal{S}))$.

Remark 6.3 Let us illustrate the convergence property for sets $\mathcal{R}_{\mathbb{X}_{ver}}^{(i)}$ as follows. For every vertex $x^i(0) \in \mathbb{X}_{ver}$, select a decomposition $x_r^i + x_s^i$ with $x_r^i \in \mathcal{R}^\infty$, which minimises $\|x_s^i\|$ for a chosen vector norm $\|\cdot\|$. Since every element $x(0) \in \mathbb{X}_{ver}$ is a convex combination of the vertices $x^i(0)$, it follows that for all $x(0) \in \mathbb{X}_{ver}$:

$$\begin{aligned} x(0) &= \sum_i a_i x^i(0) = \sum_i a_i x_r^i(0) + \sum_i a_i x_s^i(0) \\ &\in \text{conv}(x_r^i(0)) + \text{conv}(x_s^i(0)) \subseteq \mathcal{R}^\infty + \bar{\mathbb{X}}_{ver}, \end{aligned}$$

with $\sum_i a_i = 1$ for $a_i \geq 0$, where $\bar{\mathbb{X}}_{ver} = \text{conv}(x_s^i(0))$, and where we have employed the standard operation of set addition. We obtain that $\mathbb{X}_{ver} \subseteq \mathcal{R}^\infty + \bar{\mathbb{X}}_{ver}$, and that the minimal positively invariant set $\mathcal{R}_{\mathbb{X}_{ver}}^\infty$ can be bounded by $\mathcal{R}^\infty + \lim_{k \rightarrow \infty} \bigcup_{i=0}^k A^i \bar{\mathbb{X}}_{ver}$. Under the discussed conditions on \mathbb{U}_{ver} and (A, B) , previously necessary for \mathcal{R}^∞ to be a bounded and convex polytope, $A^i \bar{\mathbb{X}}_{ver}$ will converge to $\{0_n\}$. Thus, the iteration $\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}$ is monotonically increasing and bounded, hence it converges. If $\bar{\mathbb{X}}_{ver}$ includes the origin in its interior, then there exists a finite iteration step k , such that $\bigcup_{i=0}^k A^i \bar{\mathbb{X}}_{ver} = \bigcup_{i=0}^{k+1} A^i \bar{\mathbb{X}}_{ver}$. Moreover, for any reachable pair (A, B) and asymptotically stable A , the closure of the minimal robustly positively invariant set $\mathcal{R}_{\mathbb{X}_{ver}}^\infty$ includes the origin.

Robust approximations of the feasible set via $\Theta_\psi(\cdot)$

In order to exploit the convergence in the computation of the feasible set for invariance properties, we need to bound the error incurred in the use of approximations of the sets $\mathcal{R}_{\text{ver}}^\infty$ or \mathcal{R}^∞ . Let \mathcal{B} denote a unit ball centred at the origin and let the Hausdorff distance between sets \mathcal{R}_1 and \mathcal{R}_2 be defined as

$$\delta_H(\mathcal{R}_1, \mathcal{R}_2) = \inf\{\epsilon \geq 0 \mid \mathcal{R}_1 \subseteq \mathcal{R}_2 + \epsilon\mathcal{B}, \mathcal{R}_2 \subseteq \mathcal{R}_1 + \epsilon\mathcal{B}\}.$$

We can show that the following holds.

Lemma 6.6 *Let us consider a model set under a reachable pair (A, B) , an asymptotically stable A , and let the input set \mathbb{U}_{ver} include the origin. Consider a polytope \mathcal{R} , and a property ψ comprised of $\psi ::= \alpha \mid \bigcirc \psi \mid \psi_1 \wedge \psi_2$, with $\alpha \in \Sigma$, for which $\Theta_\psi(\mathcal{R})$ is a non-empty polytope with vertices v_i and the origin in its interior. Let A be bounded as $\|A\|_2 \leq 1$. Then for any $\epsilon_x \geq 0$,*

$$\begin{aligned} \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) &\subseteq \Theta_\psi(\mathcal{R}) \subseteq \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) + \epsilon_\theta\mathcal{B} & (6.12) \\ \text{if } \epsilon_\theta &\geq \frac{\epsilon_x \epsilon_p \max_i (\|v_i\|)^2}{1 + \epsilon_x \epsilon_p \max_i (\|v_i\|)}, \text{ for } \epsilon_p := \max_{p \in AP} \frac{\|A_p\|_2}{|b_p|}. \end{aligned}$$

Proof: 1. $\Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) \subseteq \Theta_\psi(\mathcal{R})$

Based on the definition of this set (c.f. the proof of Theorem 6.3), the set operation $\Theta_\psi(\cdot)$ is monotonically decreasing as in Lemma 6.5. Therefore $\Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) \subseteq \Theta_\psi(\mathcal{R})$ holds.

2. $\Theta_\psi(\mathcal{R}) \subseteq \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) + \epsilon_\theta\mathcal{B}$

Consider the case where the model is parameterised as $(A, B, C(\theta), 0)$. To prove (6.12), we first find a ϵ_θ as a function of ϵ_x such that

$$\Theta_\psi(\mathcal{R}) \subseteq \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) + \epsilon_\theta\mathcal{B}. \quad (6.13)$$

Let v_i be the vertices of the polytope $v_i \in \mathcal{V}(\Theta_\psi(\mathcal{R}))$ (as used in Lemma 6.5), then (6.13) holds if and only if $v_i \in \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B}) + \epsilon_\theta\mathcal{B}$. Equivalently, this means that there exists an $r_\theta \in \epsilon_\theta\mathcal{B}$ such that $v_i - r_\theta \in \Theta_\psi(\mathcal{R} + \epsilon_x\mathcal{B})$. This is equivalent to demanding that for every $x_j \in \mathcal{V}(\mathcal{R})$, $v_i \in \mathcal{V}(\Theta_\psi(\mathcal{R}))$ and $r_x \in \epsilon_x\mathcal{B}$, there exists a vector $r_\theta \in \epsilon_\theta\mathcal{B}$:

$$\begin{aligned} &((I_{n_\psi} \otimes (x_j^T + r_x^T))N_\psi + K_\psi)(v_i - r_\theta) \leq B_\psi \\ \Leftrightarrow &((I_{n_\psi} \otimes x_j^T)N_\psi + K_\psi)(v_i - r_\theta) + ((I_{n_\psi} \otimes r_x^T)N_\psi)(v_i - r_\theta) \leq B_\psi. \end{aligned}$$

Take $(v_i - r_\theta) = (1 - \alpha_i)v_i$ with $\alpha_i \in [0, 1]$, then

$$\begin{aligned} &((I_{n_\psi} \otimes x_j^T)N_\psi + K_\psi)(1 - \alpha_i)v_i + ((I_{n_\psi} \otimes r_x^T)N_\psi)(1 - \alpha_i)v_i \leq B_\psi \\ \Leftrightarrow &(1 - \alpha_i)(I_{n_\psi} \otimes r_x^T)N_\psi v_i \leq \alpha_i B_\psi. \end{aligned} \quad (6.14)$$

Separate the matrix N_ψ and B_ψ into its block matrices $N_\psi^j = [N_\psi]_{\{1+(j-1)n:nj\} \times \{1:n\}}$

and $B^j = [B_\psi]_j$, such that inequality (6.14) is equivalent to the set of inequalities

$$\begin{aligned} (1 - \alpha_i) r_x^T N_\psi^j v'_i &\leq \alpha_i B^j, \text{ for } j = 1, \dots, n_\psi \\ \Leftrightarrow r_x^T N_\psi^j v'_i &\leq \frac{\alpha_i}{(1 - \alpha_i)} B^j. \end{aligned}$$

Given that 0 is in the interior of $\Theta_\psi(\mathcal{R})$, it follows that $B_j > 0$ for $j = 1, \dots, n_\psi$

$$\max_j \left(r_x^T N_\psi^j v'_i \right) (B^j)^{-1} \leq \frac{\alpha_i}{(1 - \alpha_i)}.$$

The term on the left can be upper bounded based on the Cauchy-Schwarz inequality

$$\begin{aligned} \max_j \left(r_x^T N_\psi^j v'_i \right) (b^j)^{-1} &\leq \max_j \|(N_\psi^j)^T r_x\|_2 \|v'_i\|_2 (b^j)^{-1} \\ &\leq \max_j \|(N_\psi^j)^T\|_2 \|r_x\|_2 \|v'_i\|_2 (b^j)^{-1} \text{ and } \|r_x\|_2 \leq \epsilon_x \\ &\leq \epsilon_x \epsilon_p \|v'_i\|_2. \end{aligned}$$

The last inequality follows from the introduction of the precision of the labelling, denoted as ϵ_p , and defined as

$$\epsilon_p = \max_{p \in AP} \frac{\|A_p\|_2}{|b_p|}. \quad (6.15)$$

Remember that $\|L \otimes K\|_2 = \|L\|_2 \|K\|_2$. Then based on Theorem 6.3 and on the condition $\|A\|_2 \leq 1$, it can be shown that

$$\max_j \|(N_\psi^j)^T\|_2 |B^j|^{-1} \leq \max_{p \in AP} \frac{\|A_p\|_2}{|b_p|}.$$

Note that $\frac{\alpha_i}{(1 - \alpha_i)}$ monotonically increases with α_i for $\alpha_i \in [0, 1)$. Therefore a bound on α_i can be found as

$$\alpha_i = (\epsilon_x \epsilon_p \|v_i\|) / (1 + \epsilon_x \epsilon_p \|v_i\|) \text{ for } j = 1, \dots, n_\psi. \quad (6.16)$$

It follows that (6.13) holds if

$$\epsilon_\theta = \max(\|v_i\|_2) \frac{\epsilon_x \epsilon_p \max(\|v_i\|_2)}{1 + \epsilon_x \epsilon_p \max(\|v_i\|_2)}. \quad (6.17)$$

For the case that the model is parameterised in both C and D , i.e., $(A, B, C(\theta), D(\theta))$ the derivation is a bit more cumbersome (cf. proof of Theorem 6.3), but can be repeated with no change to the end result. \square

Let us briefly discuss the conditions under which Lemma 6.6 is applicable. The requirement that $\Theta_\psi(\mathcal{R})$ is not empty is raised to avoid the trivial case where $\Theta_\psi(\mathcal{R}) = \emptyset$ in (6.12) holds for all ϵ_θ . The condition that $\Theta_\psi(\mathcal{R})$ is a polytope (and

hence bounded) is necessary to obtain a bounded Hausdorff distance. This distance quantifies the difference between two sets, and is a necessary step to bound the approximation error. The requirement that $\Theta_\psi(\mathcal{R})$ includes the origin is a sufficient condition and relates to well-posedness for bounded input sets including the origin. When considering invariance properties defined for $0 \in \mathbb{U}_{ver}$ and for any polytope \mathbb{X}_{ver} , the requirement that $0_n \in \Theta_\psi(\cdot)$ is necessary for $\Theta_{\square\psi}$ to be non-empty: this can be intuitively illustrated by noting that under an assumption of asymptotic stability for A , for any θ and for $u(\cdot) = 0$, the output $\hat{y}(t, \theta)$ of the model in (6.6) converges to 0. Hence for a property to be satisfied under these conditions it should at least hold for the zero output, which is equivalent to demanding that it holds for the parameter $\theta = 0_n$. For any atomic proposition $p_i \in AP$ (see Equation (6.5)) it can be shown that there is an invertible mapping between the row vectors, proportional to the normals of the faces of the polyhedral set $\Theta_{p_i}(x(0))$, and the initial state $x(0)$. Therefore, if $\mathcal{R}^{(k)}$ has the origin in its interior, then $\Theta_{p_i}(\mathcal{R}^{(k)})$ has to be bounded, and as a consequence so does any feasible set comprising this atomic proposition. This holds for $k \geq n$ if (A, B) is a reachable pair and if \mathbb{U}_{ver} has 0 in its interior. Under the same conditions there exists a k such that $\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}$ has 0_n in its interior. The generalisation to the case dealing with a Hausdorff distance of the feasible set for invariance properties with a set of inputs $0 \notin \mathbb{U}_{ver}$ is outside of the scope of this work.

Convergence properties of robust approximations

We can employ Lemma 6.6 to bound the Hausdorff distance between $\Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)})$ and $\Theta_{\square\psi}$. If $\mathbb{X}_{ver} = \{0_n\}$ and the spectral radius of A is strictly less than 1 (that is $\rho(A) < 1$ or equivalently A is asymptotically stable), then the Hausdorff distance can be bounded as

$$\delta_H(\mathcal{R}^{(k)}, \mathcal{R}^\infty) \leq \epsilon(k) := \|A^k\|_2 \max_{u \in \mathbb{U}} (|u|) c_1, \quad (6.18)$$

with c_1 a bound on $\sum_{i=0}^{\infty} \|A^i B\|$, which is the peak-to-peak performance of the dynamical system formed by (A, B) . The derivation of the inequality above, and of the subsequent results, can be found in the Appendix. Stronger results can be obtained via dedicated software for these computations [Frehse et al., 2011]. In the case that $\mathbb{X}_{ver} \not\subseteq \mathcal{R}^\infty$ then the forward reachable iteration can be rewritten as

$$\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} = \left(\bigcup_{i=0}^k A^i \mathbb{X}_{ver} \right) + \mathcal{R}^{(k)}.$$

The Hausdorff norm can be bounded as

$$\delta_H(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}, \mathcal{R}_{\mathbb{X}_{ver}}^\infty) \leq \epsilon(k) + \|A^{k+1}\|_2 \delta_H(\mathbb{X}_{ver}, \{0_n\}).$$

Note that for $\rho(A) < 1$ the norm $\|A^k\|_2 \rightarrow 0$ for $k \rightarrow \infty$. In case the conditions of Lemma 6.6 on $\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} \subseteq \mathbb{X}$ and $\Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)})$ hold, the Hausdorff distance

$\delta_H(\Theta_{\square^k\psi}, \Theta_{\square\psi})$ can be bounded by

$$\|A^k\|_2 \max_i (\|v_i\|)^2 \epsilon_p \left(\max_{u \in \mathbb{U}} (|u|) c_1 + \|A\| \delta_H(\mathbb{X}_{ver}, \{0_n\}) \right). \quad (6.19)$$

Verification of unbounded-time properties

Based on the convergence properties of the feasible set, the asymptotic behaviour of the confidence computed in Proposition 6.2 can be stated as follows.

Corollary 6.7 (Convergence) *Under the conditions of Lemma 6.6, the feasible sets $\Theta_{\square^k\psi}$ and $\Theta_{\square\psi}$ are measurable; further, for a Gaussian distribution $p(\theta) \sim \mathcal{N}(\mu_\theta, R_\theta)$ with a covariance $R_\theta \succ 0$, $\mathbb{P}(\theta \in \Theta_{\square^k\psi}) \rightarrow \mathbb{P}(\theta \in \Theta_{\square\psi})$ for $k \rightarrow \infty$.*

Proof: For a strictly positive R_θ , the Gaussian density distribution takes finite values over the parameter space, therefore the convergence of a monotonically-decreasing polytope over the parameter space induces the convergence of the associated probability measure. \square

Theorem 6.3 can now be generalised to include unbounded-time invariance properties as follows.

Theorem 6.8 *Consider a polytopic set of initial states $x(0) \in \mathbb{X}_{ver}$, inputs $u(t) \in \mathbb{U}_{ver}$ for $t \geq 0$, and a labelling map as in (6.5). Let $\hat{\mathcal{R}}_{\mathbb{X}_{ver}}^\infty$ be a polytopic superset of the minimal robustly positively invariant set that includes \mathbb{X}_{ver} , denoted as $\mathcal{R}_{\mathbb{X}_{ver}}^\infty$. Then the feasible set admits a polyhedral subset $\hat{\Theta}_\psi \subset \Theta_\psi$ for every specification ψ expressed within the LTL fragment $\psi ::= \alpha | \bigcirc \psi | \psi_1 \wedge \psi_2 | \square \psi$, and if $\hat{\mathcal{R}}_{\mathbb{X}_{ver}}^\infty = \mathcal{R}_{\mathbb{X}_{ver}}^\infty$ then $\hat{\Theta}_\psi = \Theta_\psi$.*

Proof: Every property $\phi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2 | \square \psi$ with $p \in AP$ can be rewritten as $\psi_1 \wedge \square \psi_2$ where ψ_1 and ψ_2 have syntax $\psi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2$. Consider a property $\psi_* = \psi_1 \wedge \square \psi_2$, and let us leverage equivalences among LTL formulae [Baier and Katoen, 2008]. For $\bar{\psi}_1$ and $\bar{\psi}_2$ in $\psi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2$ the properties $\psi_{1*} := \psi_1 \wedge \bar{\psi}_1$ and $\psi_{2*} := \psi_2 \wedge \bar{\psi}_2$ are such that $\psi_* \wedge (\bar{\psi}_1 \wedge \square \bar{\psi}_2) \equiv \psi_{1*} \wedge \square \psi_{2*}$. Now consider $\bigcirc \psi_* \equiv \bigcirc(\psi_1 \wedge \square \psi_2) \equiv \bigcirc(\psi_1 \wedge \square \psi_2) \equiv (\bigcirc \psi_1) \wedge (\bigcirc \square \psi_2)$, from the distributive law of \bigcirc . Using the semantics of \bigcirc and \square , it follows that $\bigcirc \square \psi_2$ is equivalent to $\square \bigcirc \psi_2$. Thus for $\psi_{1*} := \bigcirc \psi_1$ and $\psi_{2*} := \bigcirc \psi_2$ it holds that $\bigcirc \psi_* \equiv \psi_{1*} \wedge \square \psi_{2*}$. Take $\square \psi_* \equiv \square(\psi_1 \wedge \square \psi_2) \equiv (\square \psi_1) \wedge (\square \square \psi_2)$ based on the distributive law (c.f. [Baier and Katoen, 2008, p.248]), which is subsequently equal to $(\square \psi_1) \wedge (\square \psi_2) \equiv \square(\psi_1 \wedge \psi_2)$ by applying the idempotency law and the distributive law. Hence $\square \psi_* \equiv \square(\psi_1 \wedge \psi_2)$.

In conclusion, every property $\phi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2 | \square \psi$ can be written as $\psi_1 \wedge \square \psi_2$ where ψ_1 and ψ_2 have syntax $\psi ::= p | \bigcirc \psi | \psi_1 \wedge \psi_2$: this is because we have shown that every operation (\bigcirc , \wedge , \square) preserves this rewriting.

For the set of initial states \mathbb{X}_{ver} , a property ψ is invariant

$$\langle \mathbf{M}(\theta), x(0) \rangle \models \square \psi, \forall x(0) \in \mathbb{X}_{ver}$$

if and only if $\forall x \in \mathcal{R}_{\mathbb{X}_{ver}}^\infty : \langle \mathbf{M}(\theta), x \rangle \models \psi$. Let $\hat{\mathcal{R}}_{\mathbb{X}_{ver}}^\infty$ be a polytopic superset of $\mathcal{R}_{\mathbb{X}_{ver}}^\infty$ with a finite set of vertices $v_{\mathcal{R}} \in V_{\mathcal{R}}$. Then the subset approximation of the feasible set $\Theta_{\square\psi}$ follows as $\Theta_{\square\psi} \supseteq \hat{\Theta}_{\square\psi} =$

$$\left\{ \theta \in \Theta : \bigwedge_{v_{\mathcal{R}} \in V_{\mathcal{R}}} ((I_{n_\psi} \otimes v_{\mathcal{R}}^T) N_\psi + K_\psi) \theta \leq B_\psi \right\},$$

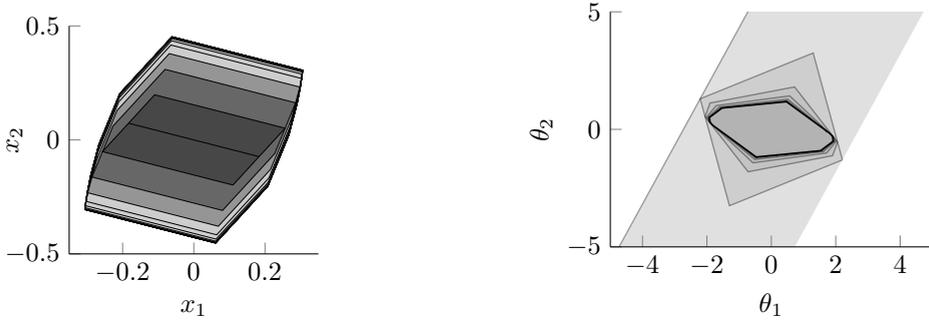
where $\hat{\Theta}_{\square\psi} \subseteq \Theta_{\square\psi}$. Note that if $\hat{\mathcal{R}}_{\mathbb{X}_{ver}}^\infty = \mathcal{R}_{\mathbb{X}_{ver}}^\infty$ then $\hat{\Theta}_{\square\psi} = \Theta_{\square\psi}$. The feasible set of $\psi_1 \wedge \square\psi_2$ is equal to $\Theta_{\psi_1 \wedge \square\psi_2} = \Theta_{\psi_1} \cap \Theta_{\square\psi_2}$. And $\Theta_{\psi_1 \wedge \square\psi_2}$ can be upper and lower bounded as $\Theta_{\psi_1} \cap \hat{\Theta}_{\square\psi_2} \subseteq \Theta_{\psi_1 \wedge \square\psi_2} \subseteq \Theta_{\psi_1} \cap \Theta_{\square^k\psi_2}$ with $k \in \mathbb{N}$. This proves Theorem 6.8 for the case where the model is $(A, B, C(\theta), 0)$. The proof for the model class with additional parameterisation of D can be derived similarly. \square

The extension beyond the LTL fragment discussed above may lead to feasible sets that are in general not convex, and is therefore beyond the scope of this work.

6.3.f Case study (continuation): unbounded-time safety verification

We study convergence properties for the safety specification ι considered in the case study in Section 6.3.d, maintaining the same operating conditions as before for the verification step and for the experiments. In Figure 6.4a the forward reachability sets $\mathcal{R}^{(k)}$ with $k = 1, \dots, 20$ are obtained for the model dynamics in (6.8). Figure 6.5 (upper plot) displays bounds $\epsilon(k)$ on the Hausdorff distances $\delta_H(\mathcal{R}^{(k)}, \mathcal{R}^\infty)$ computed with (6.18): starting from a slanted line segment for $\mathcal{R}^{(1)}$ as in Figure 6.4a, it can be observed that the forward reachable sets $\mathcal{R}^{(k)}$ converge rapidly, as confirmed with the error bound displayed in Figure 6.5 (upper plot).

Based on $\mathcal{R}^{(k)}$, the feasible set for the k -bounded invariance $\square^k \iota$ can be computed as $\Theta_{\square^k \iota} = \Theta_\iota(\mathcal{R}^{(k)})$. The feasible sets $\Theta_{\square^k \iota}$ with $k = 1, \dots, 20$ are plotted in Figure 6.4b. Observe that the feasible set $\Theta_{\square^1 \iota}$ is not bounded, but for $k \geq 2$ the feasible sets are bounded and, as expected, decrease in size with time. In Figure 6.5 (middle plot) bounds on the Hausdorff distances $\delta_H(\Theta_{\square^k \iota}, \Theta_{\square^{k+1} \iota})$ are given for $k = 2, \dots, 20$ (no finite bound is computed for the index $k = 1$, since for that instance the feasible set is not bounded). Let us conclude this case study looking at confidence quantification, as a function of the time horizon. Figure 6.5 (lower plot) represents the confidence over the property $\mathbb{P}(\theta \in \Theta_{\square^k \iota} \mid Z^{N_s})$, for indices $k = 1, \dots, 20$. Unlike the case discussed in Section 6.3.d, which focused on looking at statistics of the confidence via mean and variance drawn over multiple experiments, we zoom in on asymptotic properties by considering a data set Z^{N_s} comprising a single trace made up of 200 measurements, simulated under the same conditions as in Section 6.3.d, and with $\theta_0 = [1 \ 0]^T$. From the resulting probability density distribution $p(\theta \mid Z^{N_s})$, it may be observed that the confidence converges rapidly to a nonzero value.



(a) The first 20 iterations of the forward reachable set $\mathcal{R}^{(k)}$, $k = 1, \dots, 20$ for the case study. The reachable sets grow in size from dark grey ($k = 1$) to light grey ($k = 20$), so that $\mathcal{R}^{(k-1)} \subseteq \mathcal{R}^{(k)}$.

(b) The feasible sets for the k -bounded invariance property $\square^k \iota$, with $k = 1, \dots, 20$, obtained for the case study.

Figure 6.4: Reachable and feasible sets for the unbounded-time verification problem.

6.4 Discussion on the generalisation of the results

The discussed approach based on polytopes allows for analytical expressions of the feasible set, however the implementation may not scale to models with very large dimension: in particular, the number of half-planes characterising the feasible set may increase with the time bound of the LTL formula ψ (that is, with the repeated application of the \bigcirc operator), and with the cardinality of the set of atomic propositions in the alphabet Σ . Still, these computations are essentially equivalent to those of known reachability algorithms, therefore the method is extensible well beyond the 2-dimensional case study, especially when applying sophisticated reachability analysis tools in the literature [Frehse et al., 2011, Cattaruzza et al., 2015]. Therefore the discussed limitations related to the current implementation of the approach, ought to be dealt with in the future by the use of tailored and less naïve computational approaches. In the discussion of model selection, we elaborated possible generalisations beyond linearly-parameterised model sets. Future extension will in particular deal with hybrid models, since when systems are not linear, their (local) behaviour is often well approximated with piecewise-linear dynamical models.

We are presently working to extensions of the considered set of logic formulae of interest, and plan to employ experiment design to optimise the input-output signal interaction for efficient data usage over general classes of models, as initially attempted in [Haesaert et al., 2015b]. Additionally, the design of control policies that optimise properties of interest over partly unknown systems is topic of current work.

Finally, current work targets the applicability of tractable solutions to model-based

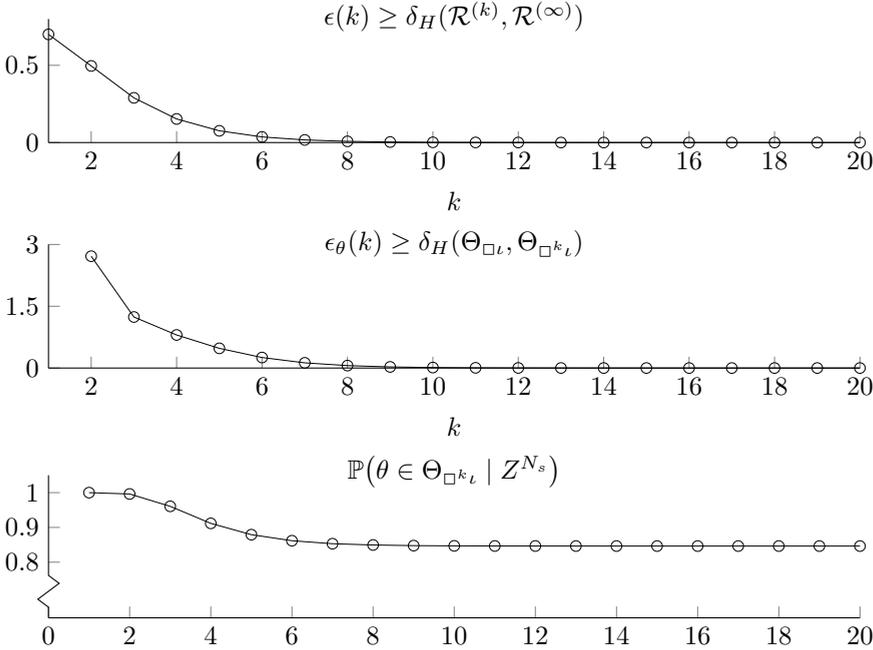


Figure 6.5: (Upper plot) Error bound on the approximation level of the k -th forward reachable sets, which is such that $\mathcal{R}^{(\infty)} \subseteq \mathcal{R}^{(k)} + \epsilon(k)$ for $k = 1, \dots, 20$. (Middle plot) The Hausdorff distance $\epsilon_\theta(k)$ between Θ_{\square^k} and $\Theta_{\square^{k_\ell}}$ with $k = 2, \dots, 20$, obtained for the case study. (Lower plot) Confidence that $\mathbf{S} \models \square^k$ for $k = 1, \dots, 20$ for the case in Section 6.3.d, with a new experiment consisting of 200 samples collected as Z^{N_s} .

and data-driven verification over complex physical systems.

6.5 Conclusions

This chapter has introduced a new framework for the integrated formal verification and modelling of physical systems with partly unknown dynamics. A Bayesian framework allowing for the efficient incorporation of measurement data and prior information has been combined with a verification procedure. The new approach allows for the computation of the confidence level over the validity of a property of interest on the unknown system. The method has been applied to the verification of LTI models of systems over bounded and unbounded safety properties (a fragment of LTL logics), and its computational overhead has been focus of discussion.

6.A Derivation of the bounds in Section 6.3.e

We derive the Hausdorff distance used in the subsection “Convergence properties of robust approximations”.

1. Hausdorff distance of forward reachable mappings. We sketch the method to bound the Hausdorff distance, whereas a more formal derivation can be found in the literature on robustly positively invariant sets [Blanchini and Miani, 2007].

The k -step forward reachable set equals to

$$\mathcal{R}^{(k)} = \bigcup_{i=1}^k \left\{ \sum_{j=1}^i A^{j-1} B u(i-j), \text{ for } u(j) \in \mathbb{U}_{ver} \right\}.$$

For $0 \in \mathbb{U}_{ver}$, the minimal invariant set \mathcal{R}^∞ can be written as

$$\mathcal{R}^{(\infty)} = \left\{ \sum_{j=0}^{i-1} A^j B u(j) + A^i \sum_{k=0}^{\infty} A^k B u(k), \text{ for } u(\cdot) \in \mathbb{U}_{ver} \right\}.$$

If the spectral radius of a A is strictly smaller than 1, $\rho(A) < 1$, then

$$\mathcal{R}^{(\infty)} \subseteq \mathcal{R}^{(k)} + \epsilon(k)\mathcal{B},$$

with

$$A^k \sum_{i=0}^{\infty} A^i B u(k) \subseteq \epsilon(k)\mathcal{B}, \text{ for } u(\cdot) \in \mathbb{U}_{ver}.$$

Note that $\epsilon(k)$ is bounded for $\rho(A) < 1$. For a matrix A without defective eigenvalues, i.e. where the eigenvectors form a complete basis, this L_1 norm (the peak-to-peak performance) can be easily bounded using the spectral radius of A , by selecting

$$\epsilon(k) = \frac{|\rho(A)|^k}{1 - |\rho(A)|} \|B\|_2 \max_{u \in \mathbb{U}_{ver}} (|u|) \geq \|A^k\|_2 \sum_{i=0}^{\infty} \|A^i B\|_2 |u(k)|.$$

In case that the matrix A is defective, we opt to bound the L_1 -norm by exploiting the absolute sum of the L_2 induced norm for A^i , $i \rightarrow \infty$: $\sum_{i=0}^{\infty} \|A^i\|_2$. Note that $\|A^i\|_2$ converges to 0 for $i \rightarrow \infty$ since $\rho(A) < 1$, therefore there exists a finite l such that $\|A^l\|_2 < 1$ and we can upper bound the absolute sum as

$$\begin{aligned} \sum_{i=0}^{\infty} \|A^i\|_2 &\leq \left(\sum_{i_1=0}^{l-1} \|A^{i_1}\|_2 \right) \left(\sum_{i_2=0}^{\infty} \|A^{i_2}\|_2 \right) \\ &= \left(\sum_{i_1=0}^{l-1} \|A^{i_1}\|_2 \right) \frac{1}{1 - \|A^l\|_2}. \end{aligned}$$

Thus in general, the Hausdorff distance can be bounded as

$$\delta_H(\mathcal{R}^{(k)}, \mathcal{R}^{(\infty)}) \leq \epsilon(k) = \|A^k\|_2 \max_{u \in \bigcup_{ver}} (|u|) c_1,$$

with $c_1 = \frac{(\sum_{i_1=0}^l \|A^{i_1}\|_2)}{1 - \|A^l\|_2} \|B\|_2$ for l such that $\|A^l\|_2 < 1$. Note that c_1 can be replaced by any bound on the L_1 norm (the peak-to-peak performance) of the dynamical system formed by (A, B) .

In case that $\mathbb{X}_{ver} \not\subseteq \mathcal{R}^\infty$ then the forward reachable iteration can be rewritten as

$$\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} = \left(\bigcup_{i=0}^k A^i \mathbb{X}_{ver} \right) + \mathcal{R}^{(k)},$$

for which we know that

$$\mathcal{R}_{\mathbb{X}_{ver}}^{(\infty)} \subseteq \mathcal{R}_{\mathbb{X}_{ver}}^{(k)} + \epsilon(k) + \|A\|^{k+1} \delta_H(\mathbb{X}_{ver}, \{0\}).$$

Thus the Hausdorff norm is upper bounded as

$$\delta_H(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}, \mathcal{R}_{\mathbb{X}_{ver}}^{(\infty)}) \leq \epsilon(k) + \|A\|^{k+1} \delta_H(\mathbb{X}_{ver}, \{0\}).$$

2. Hausdorff distance on feasible sets. Suppose that the conditions in Lemma 6.6 hold for $\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}$, then we can compute a value for ϵ_θ such that $\Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}) \subseteq \Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} + \epsilon_x \mathcal{B}) + \epsilon_\theta \mathcal{B}$, where ϵ_x is a bound on the Hausdorff distance $\delta_H(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}, \mathcal{R}_{\mathbb{X}_{ver}}^{(\infty)})$.

The set operation $\Theta_\psi(\cdot)$ is monotonically decreasing, therefore $\Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} + \epsilon(k) \mathcal{B}) \subseteq \Theta_{\square\psi} = \Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(\infty)}) \subseteq \Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)}) = \Theta_{\square^k\psi}$, and $\Theta_{\square^k\psi} \subseteq \Theta_\psi(\mathcal{R}_{\mathbb{X}_{ver}}^{(k)} + \epsilon(k) \mathcal{B}) + \epsilon_\theta \mathcal{B} \subseteq \Theta_{\square\psi} + \epsilon_\theta \mathcal{B}$, and

$$\Theta_{\square\psi} \subseteq \Theta_{\square^k\psi} \subseteq \Theta_{\square\psi} + \epsilon_\theta \mathcal{B}.$$

Based on Lemma 6.6, with $\epsilon_p = \max_{p_i} \frac{|A_{p_i}|}{|b_{p_i}|}$, we obtain

$$\epsilon_\theta = \frac{\epsilon_x \epsilon_p \max_i (\|v_i\|)^2}{1 + \epsilon_x \epsilon_p \max_i (\|v_i\|)} \leq \epsilon_x \epsilon_p \max_i (\|v_i\|)^2.$$

Note that since $\|A^k\|_2$ converges to 0 for $k \rightarrow \infty$ for $\rho(A) < 1$, and since $\max_i (\|v_i\|)^2$ is not increasing, the error ϵ_θ also converges to 0.

*Inside every non-Bayesian there is a
Bayesian struggling to get out*

Dennis V. Lindley

7

Bayesian experiment design for formal verification

The formal verification of a system with partly unknown dynamics hinges on the use of data that is obtained from experiments performed on this system. We question how to design experiments to optimise the data-collection for this data-driven verification of formal system properties. More precisely, we employ a Bayesian framework both for the computation of the confidence level in the formal properties of interest and for the design of optimal experiments. After considering the available experiment design methods in literature, we develop two new design methodologies. The first gives an empirical solution to the classical Bayesian experiment design problem, which we formulate for the formal verification problem. Its solution defines offline a sequence of inputs for the experiment. In the second approach, we show that by formulating the experiment design problem as a stochastic optimal control problem, solvable via dynamic programming, we obtain a policy for the online selection of optimal experiment inputs. A number of numerical case studies are used to elucidate the approaches and their differences.

7.1 Introduction

In this chapter, we are interested in the optimal design of input signals, applied to partly unknown systems during an identification experiment, for the specific case that the data obtained from the experiment is to be used to verify or falsify system properties (cf. Chapter 6). Properties of interest include safety or reachability requirements on the dynamics that can be naturally formulated as specifications in a given temporal logic [Baier and Katoen, 2008]. Towards the goal of efficient data-acquisition, the input signals exciting the system should be chosen to maximise the amount of information gained.

The design of experiments applied to dynamic systems is not a new topic. Early work by *inter alia* Goodwin et al. [1973] considered experiments used for the identification of a parameterised model based on data, as such they focused on the design of experiments that reduce the variance of the parameter estimates. In this scope, optimality criteria for the experiment design can be expressed as a function of the parameter variance. When used in model-based controller design [Hjalmarsson, 2009, 2002], the quality of the employed model estimated from data influences the obtained performance. Later, application-oriented research also included this target-application of the estimated model into the experiment design problem. Solutions to the aforementioned experiment design problems typically depend on knowledge of the true model of the system, and the literature distinguishes three approaches: an iterative approach, where knowledge on an estimate of the nominal system is used to design the experiment at each stage; a min-max design that is robust to worst-case scenarios; and a Bayesian design that uses knowledge of a prior uncertainty distribution over the model [Lindley, 2000]. In recent research on experiment design, where identification for control prevails, the first approach is predominant both for the direct optimisation of an application oriented optimality criterion Hjalmarsson [2005, 2009] and for the dual experiment time optimisation Bombois et al. [2006, Least-costly]. Still, some work has been done on the robust experiment design using the min-max approach [Rojas et al., 2007]. On the other hand, the third approach, well known from Bayesian statistics [Lindley, 2000], is not yet widely employed.

The use of data to make a decision on a property of a partly unknown system, has recently been studied in the design and interpretation of experiments for closed-loop performance diagnoses Mesbah et al. [2012a,b, 2015] and hypothesis testing [Bombois et al., 2011].

We first formulate our data-driven verification problem as introduced in Chapter 6. For this problem, we qualitatively evaluate the available work on hypothesis testing methods of Mesbah et al. [2012a,b, 2015], Bombois et al. [2011] and the methods from identification for control. We show that, since the experiment outcome is the confidence in a property, experiment design criteria formulated on this confidence distinguish themselves from typical problems related to the quality of an estimated model. This difference follows from the fact that for our verification problem there is only an indirect need to minimise the accuracy (and especially the variance) of a parameter estimate.

To analyse the design of experiments for data-driven verification, we consider the case that the system can be modelled by a given, linearly parameterised, model set. Embedded within a Bayesian framework, where prior knowledge is defined by a probability distribution over the possible set of models, we consider two optimality criteria of interest. For the first, we consider a classical, utility-based, interpretation of Bayesian experiment design [Lindley, 2000] in which the optimality criterion is the expected confidence at the end of the experiment. As a close alternative, we formulate as a second criterion the reaching of a certain threshold confidence in accepting or rejection the property of interest evaluated over the prior knowledge. We note, and will highlight later on, that the latter objective is related to research on sequential Bayesian experiment design [Huan and Marzouk, 2016] and to the large body of work on sequential hypothesis testing [Wald, 1945].

We analyse and solve the classical Bayesian experiment design problem first. For this, we detail the Bayesian experiment design formulation for the design of input sequences, and show that it can be solved offline before the experiment by using Monte-Carlo methods.

Then we consider the second optimisation criterion and we show that for this case, the problem of experiment design can be reformulated as a stochastic reach-avoid problem over a Markov decision process (MDP) [Bertsekas and Shreve, 1996]. More precisely, we formulate this optimal control problem to synthesise the optimal experiment design policies using dynamic programming. For this, the idea of reformulating the experiment design problem as an MDP optimisation originates from [Larsson, 2014], and is newly enhanced by embedding the posterior distribution of the model parameters into the state of the MDP. This extended MDP allows for an *online input design* that depends on the collection of the available data, since the state of the MDP encompasses the collected data via the updated posterior distribution. A similar idea has also been developed for Bayesian experiment design with respect to the data-collection on static functions [Huan and Marzouk, 2016], where an MDP is defined with as state the current posterior distribution. Note that as they consider a static system without a state, the MDP modelling does not need to include the dynamics of the systems. The online experiment design framework can be seen as an extension to the offline design, introduced first, as it synthesises a state-based policy, rather than a state-independent and offline-computed input sequence.

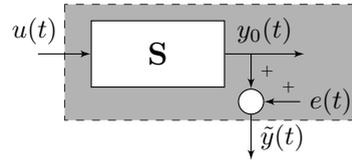
Structure. In the next section, we detail the experiment design objective for data-driven verification. The subsequent section (cf. Section 7.3) covers the literature on experiment design in system identification and shows that the posed problem distinguishes itself from the standard experiment design problems for estimation. Sections 7.4 and 7.5 give the respective offline and online experiment design formulations and solutions.

7.2 Bayesian experiment design for data-driven verification

7.2.a Modelling framework and specifications

The system, denoted by S as given in its measurement set-up in Figure 7.1, is a physical system (introduced already in Chapter 2) and is measured and controlled in discrete time. An input signal $u(t)$, $t \in \mathbb{N}$, captures how the environment acts on the system. Similarly, the output $y_0(t)$ indicates how the system interacts with the environment (namely, how it can be measured). The measurements $\tilde{y}(t)$ at $t \in \mathbb{N}$ of $y_0(t)$ are disturbed by the measurement noise $e(t)$.

Figure 7.1: System S has input $u(t)$ and output $y_0(t)$. In the measurement setup, the measured output $\tilde{y}(t)$ includes the system output $y_0(t)$ and the measurement noise $e(t)$.



In most cases the knowledge of the behaviour of a system is only partial, making it impossible to represent the system with a “true” model. In such cases, a-priori available knowledge allows us to construct a model set \mathcal{G} , with elements $M \in \mathcal{G}$ representing possible mathematical models of S . The model set \mathcal{G} is defined to be a collection of state-space models in this work, but other modelling formalism such as transfer functions can also be used. Let us denote the parameterisation of the model set \mathcal{G} as the mapping $M(\cdot) : \Theta \rightarrow \mathcal{G}$. That is a mapping from the parameters $\theta \in \Theta$ in the parameter set, which is a subset of a Euclidean space $\Theta \subset \mathbb{R}^n$, to the models M in \mathcal{G} . This allows for a parametrised expression of the

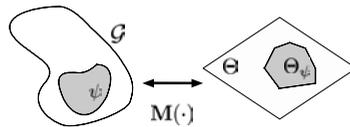


Figure 7.2: Model set indexed by the parameterisation depicted together with the set of feasible parameters Θ_ψ for which the property of interest is satisfied.

model set as $\mathcal{G} = \{M(\theta) | \theta \in \Theta\}$.

Within the scope of the developed work, we will make a sequence of assumptions that simplify the Bayesian calculations and that allow us to reason more easily about the problem at hand. They are listed at the end of this subsection. Extensions beyond this rather simple assumptions are possible but would require more involved computations, therefore they are left to future work. In the sequel, we presume to be given a linearly parameterised model set, such as

$\mathcal{G} := \{\mathbf{M}(\theta) \mid \theta \in \Theta\}$ with

$$\mathbf{M}(\theta) : \begin{cases} x(t+1) &= Ax(t) + Bu(t), \\ \hat{y}(t, \theta) &= \theta^T x(t). \end{cases} \quad (7.1)$$

For this, we assume that the signals $u(t)$ and $\hat{y}(t, \theta)$ take values in, respectively, $\mathbb{U} := \mathbb{R}$ and $\mathbb{Y} := \mathbb{R}$; and that the state evolves over \mathbb{R}^n . The chosen parameterised model set is assumed to contain the “true” model denoted as $\mathbf{M}(\theta^0)$, $\theta^0 \in \Theta$ that exactly represents the behaviour of the system \mathbf{S} . It is then the (unknown) model denoted by $\mathbf{M}(\theta^0) = \mathbf{S}$ for which $y_0(t) = \hat{y}(t, \theta^0)$ that we would ideally like to formally verify. The uncertainty about $\mathbf{M}(\theta^0)$ is structured as a distribution over the parameter set Θ . In the sequel we will assume knowledge of the initialisation of $\mathbf{M}(\theta)$. More precisely, we assume to be given the value of $x(0)$, the initialising state of the true model $\mathbf{M}(\theta^0)$.

Consider Figure 7.2, it depicts the model set \mathcal{G} and the mapping of its elements to the parameter space. As portrayed in the figure, the model class \mathcal{G} has a subset of models that satisfy the specification or property that we want to verify. We will presume that we can map this set of models to a set in the parameter space, which we can describe or approximate analytically. Let us define $\Theta_\psi \subseteq \Theta$ to be the maximal feasible set of parameters, such that for every parameter in that set the property ψ holds, i.e. $\forall \theta \in \Theta_\psi : \mathbf{M}(\theta) \models \psi$ and $\forall \theta \notin \Theta_\psi : \mathbf{M}(\theta) \not\models \psi$. The formula $\mathbf{M}(\theta) \models \psi$ reads as “the model indexed with θ satisfies property ψ ”. In the remainder of this chapter we consider temporal properties that, given a set of parameterised models $\mathbf{M}(\theta)$, translate to polytopic sets of feasible parameters: this is the case for specifications expressed in a fragment of linear temporal logic (as explained in the previous chapter) such as those for safety requirements. Still the developed theory holds for any subset Θ_ψ , it specifically also holds for non-convex sets and/or unbounded set Θ_ψ .

Beyond the standard assumptions on the system, which includes, among others, that it can be modelled by a linear time-invariant model within a known model set, we have also assumed that

- we do the experiment on a single-input single-output discrete-time system, that
- the model set can be linearly parameterised, and that
- we have knowledge on the initialisation of the models.

Further, on the experiment set-up, introduced in detail next, we will assume that

- we know and can choose the input signal (in a deterministic fashion) to the system at every time instant, and that
- the output measurements are exact up to an error $e(t)$, which in turn can be modelled as a white noise disturbance $e(t)$, whose probability density distribution is known to be a zero-mean Gaussian distribution with a given variance σ_e^2 .

Additionally, as is standard in the Bayesian framework, we assume that our knowledge or uncertainty on the system can be captured by a prior distribution over the

parameter space.

7.2.b Data-driven verification

Whenever the lack of full knowledge on the system behaviour hinders the use of formal verification procedures, it is possible to collect data of the system via an experiment. By combining data gathered from an experiment set-up together with model knowledge of the system we gain knowledge to verify whether the system does or does not satisfy a specification. The number of data necessary for this model-based and data-driven verification method does not only depend on the property and the available prior information but also on the design of the experiment. More precisely it depends on the input sequence applied to the system during the experiment.

Experiment. Data is collected in an *experiment* by exciting the system with an input sequence $\mathbf{u}_{N_s} = [u(0) \ u(1) \ \dots \ u(N_s - 1)]^T$, where N_s denotes the length of the input sequence. Noisy observations $\tilde{y}(t)$ of the output $y_0(t)$ are modelled as signals perturbed by Gaussian¹ white noise $e(t)$ that is additive to $y_0(t)$, i.e. $\tilde{y}(t) = y_0(t) + e(t)$. Let us denote the output samples obtained by exciting the system with the input \mathbf{u}_{N_s} as $\tilde{\mathbf{y}}_{N_s} = [\tilde{y}(1) \ \tilde{y}(2) \ \dots \ \tilde{y}(N_s)]^T$. The uncertainty about $\mathbf{M}(\theta^0)$ is expressed by the prior probability density distribution, denoted $p(\theta)$, and by the posterior probability density distribution $p(\theta|\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$. The collected input-output data contains statistical information on the behaviour of the system, and allows us to refine the uncertainty by updating the probability density distribution $p(\theta|\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ over the parameter space.

Confidence computation. According to Bayesian probability calculus [Lindley, 2000], the confidence (or equivalently the credibility) of a property ψ is measured with the probability distribution over the parameter space quantifying the uncertainty. Thus the confidence is computed over Θ_ψ based on the prior density distribution $p(\theta)$ (before acquiring data) or the posterior density distribution (after data acquisition), respectively:

$$\mathbb{P}(\Theta_\psi) = \int_{\Theta_\psi} p(\theta) d\theta, \quad \mathbb{P}(\Theta_\psi|\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \int_{\Theta_\psi} p(\theta|\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) d\theta. \quad (7.2)$$

In the latter case, the (additional) experiment data and parametric inference has refined the a-posteriori uncertainty distribution.

Bayesian inference. Given a prior distribution $p(\theta)$ and a data set $\tilde{\mathbf{y}}_{N_s}$ obtained by taking N_s measurements of $\tilde{y}(t)$, the a-posteriori uncertainty distribution $p(\theta|\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ is based on parametric inference [Lindley, 2000, Peterka, 1981] structured over the parameter set Θ as

$$p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \frac{p(\tilde{\mathbf{y}}_{N_s}, \theta | \mathbf{u}_{N_s})}{p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s})} = \frac{p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s}) p(\theta)}{\int_{\Theta} p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s}) p(\theta) d\theta}. \quad (7.3)$$

¹ Under this assumption we have conjugacy and hence the Bayesian inference introduced next can be solved analytically [Peterka, 1981].

Consider the model class (7.1) with the respective parameterisation and prior $p(\theta) = \mathcal{N}(\mu, R)$, then the posterior based on measurements up to time N_s is given as

$$p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \mathcal{N}(\mu^+, R^+) \text{ with } \begin{cases} R^+ &= [R^{-1} + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s})]^{-1}, \\ \mu^+ &= R^+ [R^{-1} \mu + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \tilde{\mathbf{y}}_{N_s}], \end{cases}$$

for $\Phi(\mathbf{u}_{N_s}) = [x(1) \dots x(N_s)] \in \mathbb{R}^{n \times N_s}$ (cf. Appendix 7.A). Notice that since $x(0)$ is known, $\Phi(\mathbf{u}_{N_s})$ can be constructed based on (7.1).

Note that, both the resulting posterior probability distribution $p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \mathcal{N}(\mu^+, R^+)$, the data realisation $\tilde{\mathbf{y}}_{N_s}$, and the unknown true parameter θ can be described by random variables with Gaussian distributions given as

$$p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s}) = \mathcal{N}(\Phi^T(\mathbf{u}_{N_s})\theta, I\sigma_e^2), \quad (7.4a)$$

$$p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s}) = \mathcal{N}(\Phi^T(\mathbf{u}_{N_s})\mu, R_{\tilde{\mathbf{y}}_{N_s}}), \quad (7.4b)$$

$$R_{\tilde{\mathbf{y}}_{N_s}} = [\sigma_e^2 I + \Phi^T(\mathbf{u}_{N_s}) R \Phi(\mathbf{u}_{N_s})],$$

$$p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \mathcal{N}(\mu^+, R^+), \quad (7.4c)$$

$$R^+ = [R^{-1} + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s})]^{-1}, \quad (7.4d)$$

$$\mu^+ = R^+ [R^{-1} \mu + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \tilde{\mathbf{y}}_{N_s}],$$

$$p(\mu^+ | \mathbf{u}_{N_s}) = \mathcal{N}(\mu, R - R^+), \quad (7.4e)$$

of which the derivation is given for completeness in Appendix 7.A. In (7.4a), the distribution of the output-data $\tilde{\mathbf{y}}_{N_s} = [\tilde{y}(1) \dots \tilde{y}(N_s)]^T$ is a function of the parameterisation.

7.2.c Bayesian experiment design

The overall goal of this chapter can be stated as follows: *starting from available a-priori knowledge on system \mathbf{S} , optimally design an experiment and implement it to gather measurements, which allow for the refinement of the uncertainty about the system and for the quantification of the confidence in a specification ψ defined over the system as in (7.2).*

More precisely, we study whether we can excite the system during the experiment to gain a maximal amount of knowledge on the property, along either one of the following two reasonings

- A. the expected confidence of verifying a property of the system correctly based on the experiment is optimised, or
- B. the probability that a given threshold confidence is reached during the experiment is optimised.

The former question reflects the classical, utility-based, Bayesian experiment design as described by Lindley in [Lindley, 2000] and has first been used for this data-driven verification in [Haesaert et al., 2015b]. We reiterate the proposed formulation of the Bayesian experiment design problem and solve problem **A** by design-

ing an input sequence offline as explained in Section 7.4. Then we go to the latter question, subtly different from the former. This dual problem formulation is beneficial. It can be written as a stochastic optimal control problem which can be decoupled into sub-problems as a (time-independent) dynamic programming problem. As a consequence, the problem **B** naturally allows for the online selection of the optimal input. We choose to pursue this online input selection by modelling the system together with the inference updates as an extended MDP and by designing a policy for the dynamic programming problem on that MDP. Hence the solution of the optimal control problem gives an optimal policy for the online selection of input actions to the system during the experiment.

In general the set of allowed experiments is limited, that is, we can only choose an input sequence \mathbf{u}_{N_s} from a specific set $\mathcal{E} \subset \mathbb{R}^{N_s}$. Examples include $\mathcal{E} := \{\mathbf{u}_{N_s} : u(t) \in \mathbb{U}_{ex}, \forall t = 0, \dots, N_s - 1\}$, with \mathbb{U}_{ex} a bounded set, such as for instance $[-u_{\max}, u_{\max}]$, $u_{\max} \in \mathbb{R}_0^+$.

7.3 Experiment design solutions in literature

7.3.a Frequentist setting

Consider a system **S** with the underlying model $\mathbf{M}(\theta^0)$. Then identification or parameter estimation amounts to the selection of a parameter $\hat{\theta}$ based on experiment data $(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$. Implicitly, $\hat{\theta}$ is a function of the noise realisation during the experiment. As such the stochastic properties of the random variable $\hat{\theta}$ can be described in relation to θ^0 . Estimator properties define the quality of the identification procedure and are a function of these stochastic properties of $\hat{\theta}$, including;

- bias, that is, the expected value $\mathbb{E}(\hat{\theta})$ in comparison to θ^0 ;
- variance of $\hat{\theta}$.

For linear time invariant models, one generally chooses estimators that are consistent. That is, estimators that are asymptotically unbiased and have a variance converging to 0 for increasing data. Additionally, when the true model is in the model set it is of interest to consider “optimal” estimators, i.e., estimators with minimal variance. For the combination of the model set (7.1) and the system, the prediction error method Ljung [1999] gives an unbiased, minimal variance² and consistent estimator. In this method, the estimate $\hat{\theta}$ of θ^0 is obtained by solving the following least squares problem

$$\hat{\theta}_N = \arg \min_{\theta} \frac{1}{N_s} \sum_{t=1}^{N_s} (y(t, \theta) - \tilde{y}(t))^2.$$

²This holds when measurement noise can be assumed to be white and Gaussian.

Given knowledge of θ^0 the resulting distribution of the estimate $\hat{\theta}$ is given by

$$\hat{\theta} \sim \mathcal{N}(\theta^0, R(\mathbf{u}_{N_s})), \quad R(\mathbf{u}_{N_s})^{-1} := \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}). \quad (7.5)$$

Because of the chosen (linear-in-parameter) model structure it holds that $R(\mathbf{u}_{N_s})$ is independent of θ^0 .

Of course, the value of the variance does not only depend on the chosen estimator, it also depends on the information content of the measurements. Generally, experiment design starts from knowledge of the true system and has as objective the increase of this information content to optimise the distribution of the random variable $\hat{\theta}$ in some way. As such, given that the Gaussian distribution of $\hat{\theta}$ is fully characterised its mean θ^0 and variance R , the optimisation objective can be expressed as a function of θ^0 and R . Thus the first step in the experiment design problem is an optimisation with respect to θ^0 and R . In a secondary step, the lack of knowledge of the true model θ^0 has to be tackled. This phased approach is a frequentist approach. Since the domain of system identification is largely dominated by frequentist approaches, so are the classical experiment design and identification for control parts, as detailed next. These frequentist based reasonings have in common that optimality is first defined conditional on θ^0 and then exported to the actual case where the parameter θ^0 is unknown. When this gives troubles, this is referred to as the ‘chicken and egg’ problem. Notice that though the ‘chicken and egg’ problem is most apparent for nonlinear parameterisations for which also R depends on both θ^0 and \mathbf{u}_{N_s} , we will show that when considering hypothesis testing [Mesbah et al., 2012b] cases that are similar to our verification problem, the resulting ‘chicken and egg’ problem can cause serious issues.

7.3.b Classical experiment design

In classical experiment design, the minimisation of the parameter estimate variance is of importance. More precisely, one considers the optimisation of a function of the expected variance matrix; different types of optimality can be discerned by the used function, as is surveyed by [Cooley and Lee, 2001, Gevers et al., 2011]. Well known types are **A**-optimality, minimising the trace of the variance matrix; **C**-optimality, minimising the weighted trace of the variance matrix; **D**-optimality, maximising the determinant of the information matrix (related to the inverse of the variance); **E**-optimality, maximising the minimum eigenvalue of the information matrix; and lastly **T**-optimality, maximising the trace of the information matrix.

7.3.c Identification for control

Part of the area of system identification [Ljung, 1999] and control engineering, is the domain of identification for control, which investigates measurement-based model construction of physical systems and this is particularly directed towards the use of these models for model-based control design. This research direction [Van den Hof and Schrama, 1995, Hjalmarsson, 2005, Bombois et al., 2006,

Hjalmarsson, 2009, Gevers, 2005] has focused on the development of approaches towards data-driven modelling for control, where the quality of the estimator is evaluated with respect to the objective of the model-based controller. Using the wrong model in control causes a degradation of the controller performance, the objective is hence to design the input of an experiment minimising the expected performance degradation [Hjalmarsson, 2009, 2005, Gevers, 2005]. Let us make this more concrete. For any model $\mathbf{M}(\theta)$ the relative performance degradation is denoted as $V_{rel}(\mathbf{M}(\theta), \mathbf{M}(\theta^0))$ and is based on the performance $\mathcal{P}(\mathbf{M}(\theta), \mathbf{M}(\theta^0))$. The latter quantifies the performance of the controller, when it is applied to the true model $\mathbf{M}(\theta^0)$ and designed based on the model $\mathbf{M}(\theta)$. The use of a model $\mathbf{M}(\theta)$ different from $\mathbf{M}(\theta^0)$, $\theta \neq \theta^0$, causes a performance degradation, which is quantified by the *relative performance degradation* as

$$V_{rel}(\mathbf{M}(\theta), \mathbf{M}(\theta^0)) = \frac{1}{2} \left\| \frac{\mathcal{P}(\mathbf{M}(\theta^0), \mathbf{M}(\theta^0)) - \mathcal{P}(\mathbf{M}(\theta), \mathbf{M}(\theta^0))}{\mathcal{P}(\mathbf{M}(\theta^0), \mathbf{M}(\theta^0))} \right\|^2.$$

Note that, through its Hessian, $V_{rel}(\mathbf{M}(\theta), \mathbf{M}(\theta^0))$, can be approximated with a quadratic function

$$V_{rel}(\mathbf{M}(\theta), \mathbf{M}(\theta^0)) \approx \frac{1}{2} (\theta - \theta^0)^T V_{rel}''(\theta^0) (\theta - \theta^0).$$

We note that this defines ellipsoidal level sets around θ^0 . The accuracy of an identified model $\mathbf{M}(\hat{\theta})$ can now be evaluated with $V_{rel}(\mathbf{M}(\hat{\theta}), \mathbf{M}(\theta^0))$. Within our linear parameterised model set, the variance of this estimate $\hat{\theta}$ obtained during the identification from data is equal to $R(\mathbf{u}_{N_s})$ (similar to the covariance of the posterior uncertainty) and depends on the inputs applied to the model. More precisely a sequence of inputs defines the model quality of an estimate $\mathbf{M}(\hat{\theta})$ as $(\theta^0 - \hat{\theta}) \sim \mathcal{N}(0, R(\mathbf{u}_{N_s}))$. This allows us to define an ellipsoidal set around θ^0 , that contains the nominal estimate $\hat{\theta}$ with a prescribed probability. Experiment design in identification for control now targets the design of $R(\mathbf{u}_{N_s})$, shaping the ellipsoidal set, such that this ellipse is a subset of an application oriented level set of $(\theta - \theta^0)^T V_{rel}''(\theta^0) (\theta - \theta^0)$. In the former, the level set is constructed based on the allowed performance degradation.

For general model sets (especially those beyond our linear parameterisation) optimal experiment depends on the *unknown* value θ_0 . This dependence on θ^0 of the optimal experiment to find θ^0 reflects the ‘chicken and egg’ problem. For a practical implementation there are several ways to deal with the lack of knowing θ_0 exactly, see [Hjalmarsson, 2009]. Starting from the standard assumption that the function V_{rel} is continuously differentiable, the computation of level sets can be done approximately. Therefore the most common option is an adaptive/sequential experiment design where θ_0 is approximated with $\hat{\theta}$. This can be done either at each time sample [Pronzato, 2000, 2008, Pronzato and Walter, 1985, Gerencsér et al., 2009, Hjalmarsson et al., 1996, Gerencsér and Hjalmarsson, 2005] or iteratively for batches of samples [Barentin et al., 2008]. Consider as an example, the iterative design of experiments for parameter estimation based on available knowledge as explored in [Stigter et al., 2006], where the input is designed via a

receding horizon optimisation based on the nominal estimate of the system, and is updated in time during the experiment.

A secondary method to deal with lack of knowledge of θ_0 is to design the experiment *robustly* with respect to the possible set of θ_0 ; this is also referred to as min-max experiment design and is used in [Rojas et al., 2007, Pronzato and Walter, 1988].

An alternative approach would be based on a Bayesian framework. Let us note that this is often used in statistics literature, as surveyed in [Chaloner and Verdinelli, 1995]. This avoids the creation of the ‘chicken and egg’ problem. There has been little need for work on the Bayesian experiment design problem for controller synthesis. The cause is that in general, for identification for control problems, the approximate solutions have provided good results.

Most of the experiment design problems, dealing with dynamical systems, are as presented above and come down to minimising stochastic properties of the estimator, which represents the expected model quality. Instead we want to verify properties over the system. This is not directly related to the model quality and hence also does not enjoy the beneficial characteristics of it.

7.3.d Experiment design for hypothesis testing

The use of measurements for the detection of failures and performance degradations has recently led to research on the corresponding experiment design. These design methodologies, focussing on hypothesis testing and failure detection [Bombois et al., 2011, Mesbah et al., 2012a], use the same two-step approach, where an optimal design is first formulated for the true model, after which a solution is sought to deal with the lack of knowledge on the true model.

The goal of the experiment design for hypothesis testing is making the right decision between two separate hypotheses. For a given bound on the false rate, the detection rate (that is, accepting the right hypothesis) is a function of the information content of the measurement data. By associating a single model to each hypothesis, Bombois et al. [2011] formulates the hypothesis testing problem as a parameter estimation problem with a single unknown parameter. This parameter spans the set of models between the null hypothesis model and the alternative hypothesis model. Thus, referred to as the detection parameter, this parameter discerns between the models belonging to the different hypotheses. As a consequence, in Bombois et al. [2011], the optimal input for the experiment design problem formulated on the properties of the hypothesis testing problem is shown to be equivalent to the input minimising the parameter variance.

Mesbah, Bombois, Ludlage, and Van den Hof [2012b] use a fully parameterised model set. A subset of this set represents the models associated to the null hypothesis \mathcal{H}_0 , and its complement defines the alternative hypothesis \mathcal{H}_1 . Starting from a known model $\mathbf{M}(\theta^0)$ the optimal experiment is defined as the experiment that optimises the probability that the right hypothesis is selected. In [Mesbah et al., 2012b] the decision rule is based on taking the hypothesis associated to the current estimate $\hat{\theta}$.

More precisely, let us characterise the hypotheses by the (feasible) set Θ_ψ , such that the null hypothesis \mathcal{H}_0 holds if $\theta^0 \in \Theta_\psi$. In contrast, the alternative hypothesis \mathcal{H}_1 is $\theta^0 \notin \Theta_\psi$. At the end of the experiment, the hypothesis \mathcal{H}_0 is accepted if $\hat{\theta} \in \Theta_\psi$, and rejected otherwise. Therefore Mesbah et al. [2012b] claims that if $\theta^0 \in \Theta_\psi$ the optimal experiment is such that with high confidence $\hat{\theta} \in \Theta_\psi$. But to actually design the experiment one now again needs to replace the unknown value θ^0 (cf. ‘chicken and egg’ problem). Suppose that the estimated θ_{init} , used to replace θ^0 , satisfies \mathcal{H}_1 , whereas θ^0 validates \mathcal{H}_0 , that is $\theta_{init} \in \Theta \setminus \Theta_\psi$ and $\theta^0 \in \Theta_\psi$. Let Θ_ψ be a bounded set and consider the distribution of the estimate $\hat{\theta}$ as

$$\hat{\theta} \sim \mathcal{N} \left(\theta^0, (\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right) \quad (7.6)$$

approximated with $\mathcal{N} \left(\theta_{init}, (\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right)$. Since $\theta_{init} \in \Theta \setminus \Theta_\psi$ we should design \mathbf{u}_{N_s} by solving

$$\mathbf{u}_{N_s, opt} \in \arg \sup \mathbb{P} \left(\hat{\theta} \in \Theta \setminus \Theta_\psi \mid \hat{\theta} \sim \mathcal{N} \left(\theta_{init}, (\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right) \right)$$

subject to some constraints on \mathbf{u}_{N_s} (e.g., bounded input power or amplitude). Let us analyse the behaviour of (7.6) for decreasing levels of excitation. For this, we consider that for

$$(\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s})) \rightarrow 0, \quad (7.7)$$

the distribution of $\hat{\theta}$ converges to a uniform distribution over \mathbb{R}^n . Given that $\Theta = \mathbb{R}^n$, the measure of the distribution, induced by (7.7), over any bounded set will converge to zero. Thus if Θ_ψ is bounded, then this also implies that

$$\mathbb{P} \left(\hat{\theta} \in \Theta \setminus \Theta_\psi \mid \hat{\theta} \sim \mathcal{N} \left(\theta_{init}, (\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right) \right) \rightarrow 1.$$

More precisely, by definition it holds that

$$\sup_{\mathbf{u}_{N_s}} \mathbb{P} \left(\hat{\theta} \in \Theta \setminus \Theta_\psi \mid \hat{\theta} \sim \mathcal{N} \left(\theta_{init}, (\sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right) \right) \leq 1;$$

hence if there exists an input such that $\Phi(\mathbf{u}_{N_s}) = 0$, then this is an optimal experiment. In other words, the optimal experiment for \mathcal{H}_1 would be the one that keeps $x(t)$ at zero, or equivalently that minimises the information content. Note that for a bounded experiment time and a bounded input power, an input \mathbf{u}_{N_s} keeping $x(t) = 0 \forall t$ is a global optimum. Since the goal is to improve the hypothesis testing procedure by improving the information content in the data, the fact that a trivial zero solution is optimal implies that the optimality criterion is ill-posed.

As suggested by [Mesbah et al., 2012b] the actual experiment design can be preceded by taking measurements with a white noise signal as input. This on its own can avoid ill-posedness issues as in this case the distribution (7.6) changes to

$$\hat{\theta} \sim \mathcal{N} \left(\theta_{init}, (\sigma_e^{-2} \Phi_0 \Phi_0^T + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}))^{-1} \right), \quad (7.8)$$

where θ_{init} is the outcome of the white noise experiment and where we denote with $\sigma_e^{-2}\Phi_0\Phi_0^T$ the contribution to this experiment. For this case, denote for the trivial zero input case the confidence α as

$$\alpha := \mathbb{P}\left(\hat{\theta} \in \Theta_\psi \mid \hat{\theta} \sim \mathcal{N}\left(\theta_{init}, (\sigma_e^{-2}\Phi_0\Phi_0^T)^{-1}\right)\right),$$

then

$$\mathbb{P}(\hat{\theta} \in \Theta \setminus \Theta_\psi \mid \hat{\theta} \sim \mathcal{N}\left(\theta_{init}, (\sigma_e^{-2}\Phi_0\Phi_0^T)^{-1}\right)) = 1 - \alpha.$$

Only, if there exists an input sequence \mathbf{u}_{N_s} such that

$$\mathbb{P}(\hat{\theta} \in \Theta \setminus \Theta_\psi \mid \hat{\theta} \sim \mathcal{N}\left(\theta_{init}, (\sigma_e^{-2}\Phi_0\Phi_0^T + \sigma_e^{-2}\Phi(\mathbf{u}_{N_s})\Phi^T(\mathbf{u}_{N_s}))^{-1}\right)) > 1 - \alpha,$$

then a nontrivial solution to the experiment design problem will be selected.

In [Mesbah et al., 2012b] this formulation of the experiment design problem is applied to a specific type of failure detection, thereby avoiding this ill-posedness issues.

Instead of taking a frequentist approach and trying to resolve the ‘chicken and egg’ problem afterwards, we will consider a full Bayesian framework for the experiment design problem. A lot of the ideas discussed in the sequel would also be applicable to the type of failure detection or performance degradation questions where measurement data is used for standard hypothesis testing as in [Mesbah et al., 2012b].

7.4 Maximising a-posteriori confidence (offline)

7.4.a Bayesian experiment design: problem formulation

Targeting the formulation of an experiment design objective that does not hinge on a satisfactory solution to the ‘chicken and egg’ problem, we formulate the optimisation criterion over the prior knowledge and based on the expected utility of the experiment. As mentioned before, the end result will be the optimisation of the expected confidence of verifying a property (i.e., Problem A in Section 7.2.c). Remark that this criterion is a utility-based criterion within Bayesian experiment design.

Within the Bayesian utility-based experiment design reasoning, data is used to make a decision (as e.g. verifying a property) that, depending on the true system, has a certain quality or utility. As such one makes the decision that maximises this expected utility. For us, this is the decision to either accept that ψ holds or to reject it. The utility of this decision is zero when wrong, whereas it is valued to be equal to one when it correctly verifies the behaviour of the system.

The *expected utility*, denoted J , is related to the acceptance or rejection of a specification based on the identification experiment. In order to optimise it, consider that system \mathbf{S} can be represented as $\mathbf{M}(\theta^0)$, with a nominal parameter θ^0 . Although

θ^0 is in general unknown, it can be perceived as a realisation of the uncertainty distribution over the parameters space, i.e. $\theta^0 \sim p(\theta)$, $\theta \in \Theta$. The acceptance or rejection of $\mathbf{S} \models \psi$ can be equivalently cast as the choice between hypothesis \mathcal{H}_0 : $\mathbf{M}(\theta^0) \models \psi$ and hypothesis \mathcal{H}_1 : $\mathbf{M}(\theta^0) \not\models \psi$. This entails a decision which is valued with 1 when correct, and with 0 when incorrect. For a given choice of \mathcal{H}_0 or \mathcal{H}_1 , and a nominal parameter θ^0 , the utility is then a binary-valued function

$$\text{ut}(\mathcal{H}_i, \theta^0) := \begin{cases} 1 & \text{if } \begin{cases} \mathcal{H}_0 = \mathcal{H}_i & \text{and } \mathbf{M}(\theta^0) \models \psi \\ \mathcal{H}_1 = \mathcal{H}_i & \text{and } \mathbf{M}(\theta^0) \not\models \psi, \end{cases} \\ 0 & \text{else.} \end{cases} \quad (7.9)$$

Note that ut has a 0 value when the chosen hypothesis is wrong, which is related in statistics to type I and type II error, respectively [Shanmugan and Breipohl, 1988, page 514]. Thus type I and II errors are valued equally.

Conditional on a data set $(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$, the nominal parameter is distributed over the parameter space as $\theta^0 \sim p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$. The *expected utility* of a decision \mathcal{H}_i conditional on the data set is thus $\mathbb{E}[\text{ut}(\mathcal{H}_i, \theta^0) | \theta^0 \sim p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})]$. Note that the *expected utility* represents the confidence that $\mathbf{M}(\theta^0) \models \psi$ or $\mathbf{M}(\theta^0) \not\models \psi$, and is a function of both the decision and the experiment $(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$. Thus when deciding on \mathcal{H}_0 or \mathcal{H}_1 , the expected utility is either

$$\begin{aligned} \overline{\text{ut}}(\mathcal{H}_0; (\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})) &:= \mathbb{E}[\text{ut}(\mathcal{H}_0, \theta^0) | \theta^0 \sim p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})] \\ &= \int_{\theta \in \Theta_\psi} p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) d\theta = \mathbb{P}(\Theta_\psi | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}), \text{ or} \\ \overline{\text{ut}}(\mathcal{H}_1; (\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})) &:= \mathbb{E}[\text{ut}(\mathcal{H}_1, \theta^0) | \theta^0 \sim p(\theta | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})] = 1 - \mathbb{P}(\Theta_\psi | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}). \end{aligned}$$

As a criterion to be optimised, we then choose the *expected utility* maximised over the decision \mathcal{H}_0 or \mathcal{H}_1 , namely

$$\begin{aligned} J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) &:= \max_{\mathcal{H}_i} \overline{\text{ut}}(\mathcal{H}_i; (\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})) \\ &= \max \{ \mathbb{P}(\Theta_\psi | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}), \mathbb{P}(\Theta \setminus \Theta_\psi | \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) \}. \end{aligned} \quad (7.10)$$

Every experiment contains statistical information on the behaviour of the system. The objective is to design an experiment \mathbf{u}_{N_s} that optimally exploits the dynamic behaviour of the system and thus optimises the expected value of the criterion J . Note that the criterion of interest $J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ is a function of the measured output $\tilde{\mathbf{y}}_{N_s}$ and input \mathbf{u}_{N_s} data, as $J : \mathbb{U}^{N_s} \times \mathbb{Y}^{N_s}$. The data realisation $\tilde{\mathbf{y}}_{N_s} \sim p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s})$ has a probability density function that is conditioned on the input signal \mathbf{u}_{N_s} . Given a prior $p(\theta)$, the probability density distribution of the data can be expressed as

$$p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s}) = \int_{\Theta} p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s}) p(\theta) d\theta, \quad (7.11)$$

where $p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s})$ is the data distribution conditioned on the input \mathbf{u}_{N_s} and on the parameter θ . The Bayesian experiment design problem optimises the expected value of the criterion J over the input signal \mathbf{u}_{N_s} for a given prior $p(\theta)$, and is

formulated as:

$$\max_{\mathbf{u}_{N_s} \in \mathcal{E}} \mathbb{E} [J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) \mid \tilde{\mathbf{y}}_{N_s} \sim p(\tilde{\mathbf{y}}_{N_s} \mid \mathbf{u}_{N_s})], \quad (7.12)$$

where the set of allowed experiments \mathcal{E} is defined as $\mathcal{E} := \{\mathbf{u}_{N_s} : u(t) \in \mathbb{U}_{ex}, \forall t = 0, \dots, N_s - 1\}$, with \mathbb{U}_{ex} a bounded set, such as for instance $[-u_{\max}, u_{\max}]$, $u_{\max} \in \mathbb{R}$.

Remark 7.1 (Interpretation of the Bayesian experiment design problem) *We observe that based on the prior $p(\theta)$ we can compute a distribution over the data $\tilde{\mathbf{y}}_{N_s}$. The criterion J expresses the expected utility of a realised data set $\tilde{\mathbf{y}}_{N_s}$, interpreted as the confidence in the optimal decision, that is, accepting or rejecting \mathcal{H}_0 . This is computed from the posterior probability distribution $p(\theta \mid \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$. Note that both J and the distribution of the output data $\tilde{\mathbf{y}}_{N_s}$ are implicitly a function of the prior $p(\theta)$.*

7.4.b Computing the Bayesian optimal experiment: an empirical approach

In general, as we cannot compute the expected value in (7.12) analytically, we need develop an empirical approach to approximately solve the problem. More precisely, we give an empirical approximation of the objective function and an input parameterisation. Notice that an imprecise solution to the experiment design, obtained before the start of the experiment, does not influence the correctness or soundness of the use of the obtained data during the experiment for the data-driven formal verification.

Consider the experiment design problem

$$\max_{\mathbf{u}_{N_s} \in \mathcal{E}} \mathbb{E}[J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) \mid \tilde{\mathbf{y}}_{N_s} \sim p(\tilde{\mathbf{y}}_{N_s} \mid \mathbf{u}_{N_s})]$$

with $J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ the expected utility as given in (7.10). Note that $J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ depends on the posterior distribution, which can be expressed as $p(\theta \mid \tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) = \mathcal{N}(\mu^+, R^+)$. Hence $J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})$ depends on the measurements $\tilde{\mathbf{y}}_{N_s}$ only through μ^+ , as in (7.4c). It follows that the optimisation problem can be written as an expected value over μ^+ (instead of $\tilde{\mathbf{y}}_{N_s}$). Thus this reduces the dimensionality of the expected value in (7.12) from the horizon of the data to the dimensionality of the parameterisation,

$$\max_{\mathbf{u}_{N_s} \in \mathcal{E}} \int_{\Theta} \max_{\Theta} \{\mathbb{P}(\Theta_\psi \mid \mu^+, \mathbf{u}_{N_s}), \mathbb{P}(\bar{\Theta}_\psi \mid \mu^+, \mathbf{u}_{N_s})\} p(\mu^+ \mid \mathbf{u}_{N_s}) d\mu^+, \quad (7.13)$$

with $\bar{\Theta}_\psi = \Theta \setminus \Theta_\psi$ and subject to $p(\mu^+ \mid \mathbf{u}_{N_s}) = \mathcal{N}(\mu, R - R^+)$.

As an affine transformation of the measurements $\tilde{\mathbf{y}}_{N_s}$, the posterior mean μ^+ is a random variable with a Gaussian distribution as $\mu^+ = R^+ [R^{-1} \mu + \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \tilde{\mathbf{y}}_{N_s}]$. Using this mean, a practical lower approximation of the maximisation inside the

integral is found as $\mathbb{P}(\Theta_\psi|\mu^+, \mathbf{u}_{N_s}) = \int_{\Theta_\psi} p(\theta|\mu^+, \mathbf{u}_{N_s})d\theta$ for $\mu^+ \in \Theta_\psi$, and $1 - \mathbb{P}(\Theta_\psi|\mu^+, \mathbf{u}_{N_s})$ else. This provides a relaxed version of (7.13), expressed as

$$\max_{\mathbf{u}_{N_s} \in \mathcal{E}} \int_{\Theta_\psi} \int_{\Theta_\psi} p(\theta|\mu^+, \mathbf{u}_{N_s})p(\mu^+|\mathbf{u}_{N_s})d\theta d\mu^+ + \int_{\bar{\Theta}_\psi} \int_{\bar{\Theta}_\psi} p(\theta|\mu^+, \mathbf{u}_{N_s})p(\mu^+|\mathbf{u}_{N_s})d\theta d\mu^+. \quad (7.14)$$

Notice that this lower bound is obtained by replacing the optimal decision rule, represented by the maximisation in (7.13), by a decision based on the set-membership of μ^+ . This alternative (and not optimal) decision rule is the same as the one taken in [Mesbah et al., 2012b] and it allows to represent the optimisation problem in a computationally simpler expression. By accepting the hypothesis based on what the current posterior mean satisfies, we get an integral over two disjoint sets in the $\Theta \times \Theta$ space. This integral replaces the one in (7.13), which is the integral over Θ of the maximum of two integrals representing $\mathbb{P}(\bar{\Theta}_\psi|\mu^+, \mathbf{u}_{N_s})$ and $\mathbb{P}(\Theta_\psi|\mu^+, \mathbf{u}_{N_s})$.

The combined distribution of θ and μ^+ , denoted by variable $\underline{\theta}^T = [\theta^T \ \mu^{+T}]$, has a Gaussian distribution $p(\underline{\theta} | \mathbf{u}_{N_s}) = \mathcal{N}(\mu_{\underline{\theta}}, R_{\underline{\theta}})$, with mean $\mu_{\underline{\theta}}^T = [\mu^T \ \mu^T]$ and covariance matrix

$$R_{\underline{\theta}} = \begin{bmatrix} R & (R - R^+) \\ (R - R^+) & (R - R^+) \end{bmatrix}.$$

Since the integral in (7.14) can in general not be computed analytically, we can either compute it with an efficient numerical method or we can empirically approximate it (statistically) via a Monte Carlo method. Within the scope of this work, we choose to work with the latter option, as explained next.

Let $\epsilon \in \mathbb{R}^{2n}$ be a dummy random variable with density distribution $\mathcal{N}(0, I)$, which is independent of the decision variable \mathbf{u}_{N_s} . For a given set A let $\mathbf{1}_A$ be the indicator function, that is, $\mathbf{1}_A(x) = 1$ if $x \in A$ and 0 otherwise. The value of the relaxed optimisation problem (7.14) can be approximated as

$$\hat{\mathbb{E}}J \approx \frac{1}{M} \sum_{i=1}^M \mathbf{1}_{(\Theta_\psi \times \Theta_\psi) \cup (\bar{\Theta}_\psi \times \bar{\Theta}_\psi)}(\mu_{\underline{\theta}} + \Lambda \epsilon_i), \quad (7.15)$$

with M realisations of $\epsilon_i \sim \mathcal{N}(0, I)$ and $\Lambda \Lambda^T = R_{\underline{\theta}}$. Note that $\mu_{\underline{\theta}} + \Lambda \epsilon_i$ are realisations of $\underline{\theta} \sim \mathcal{N}(\mu_{\underline{\theta}}, R_{\underline{\theta}})$. Hence, for a given input \mathbf{u}_{N_s} , (7.15) is an unbiased estimate of the objective in (7.14) and it is also consistent, i.e., for $M \rightarrow \infty$ the estimated objective converges to the optimisation objective in (7.14) with probability 1. Via $R_{\underline{\theta}}$ the transformation of the relations ϵ_i depends on the choice of input data \mathbf{u}_{N_s} . As such, the maximisation of (7.15) over the input data \mathbf{u}_{N_s} induces a bias (converging to zero for an increasing number of data) and is hence not an unbiased estimate of the maximised value of (7.14).

Next, we give two alternatives to solve the empirical optimisation formulated with (7.15). The first one uses a parameterisation of the input signal over time. The second, applicable to experiments with longer time horizons, exploits the Fisher information matrix and properties of stationarity of the input signal.

7.4.b.1 Parameterisation of finite-time input sequences

Even for short experiments the optimisation over \mathbf{u}_{N_s} is complex due to the high number of design variables. Therefore in this section, we reduce the complexity of the optimisation problem, by reducing the number of free design variables. For the following case study, we choose a multi-sine parameterisation given as $u(t) = \sum_{k=1}^{n_p} \beta_k \sin(\omega_k t + \alpha_k)$, with parameters $\alpha_k \in [0, 2\pi]$ and $\beta_k \in [0, \infty)$ for $k = 1, \dots, n_p$ at predefined frequencies ω_k . As such, the N_s decision variables of \mathbf{u}_{N_s} are reduced by a parameterisation of the input signal \mathbf{u}_{N_s} to $2n_p$. Note that whereas β_k defines the amplitude of the individual sinusoids, the overall amplitude of the input signal depends also on the phases α_k .

Case study – Bounded-time safety verification.

Let us reiterate the case study in Section 6.3.d. Consider a system \mathbf{S} with input signals taking values in a bounded set $u(t) \in \mathbb{U} = [-0.2, 0.2]$. We assume that system \mathbf{S} can be represented as an element of a model set \mathcal{G} with transfer functions characterised by second-order Laguerre-basis functions [Heuberger et al., 2005] (a special case of orthonormal basis functions), which translates to the following parameterised state-space representation:

$$\begin{aligned} x(t+1) &= \begin{bmatrix} a & 0 \\ 1-a^2 & a \end{bmatrix} x(t) + \begin{bmatrix} \sqrt{1-a^2} \\ (-a)\sqrt{1-a^2} \end{bmatrix} u(t), \\ \hat{y}(t, \theta) &= \theta^T x(t). \end{aligned}$$

The coefficient a is chosen to be $a = 0.4$. We further consider, as a-priori available knowledge on the system, a distribution $p(\theta) = \mathcal{N}(\mu, R)$ on the parameter space, and a given initial state $x_0 = [0 \ 0]^T$.

We want to do an experiment that will allow us to verify whether the output $y_0(t)$ remains within the interval $\mathbf{I} = [-0.5, 0.5]$, labeled as ι , for the first 5 time steps for any input sequence taking values in \mathbb{U} . Introduce accordingly the alphabet $\Sigma = \{\iota, \tau\}$ and the labelling map $L : L(y) = \iota, \forall y \in \mathbf{I}, L(y) = \tau, \forall y \in \mathbb{Y} \setminus \mathbf{I}$. This means that we want to know whether the corresponding finite-horizon LTL property holds over the system \mathbf{S} , that is $\mathbf{S} \models \bigwedge_{i=1}^5 (\circ)^i \iota$. To this end we first determine the feasible Θ_ψ for which the models satisfy $\bigwedge_{i=1}^5 (\circ)^i \iota$. That is $\mathbf{M}(\theta) \models \bigwedge_{i=1}^5 (\circ)^i \iota$ if and only if $\theta \in \Theta_\psi$. A graphic representation of the obtained feasible set is given in Figure 7.3. For more information on the computation of these type of sets, we refer to Chapter 6. Remark that for a given prior $p(\theta) = \mathcal{N}(\mu, R)$, we can now immediately compute our (a-priori) confidence in $\mathbf{S} \models \bigwedge_{i=1}^5 (\circ)^i \iota$, that is,

$$\mathbb{P}(\Theta_\psi | \theta \sim \mathcal{N}(\mu, R)).$$

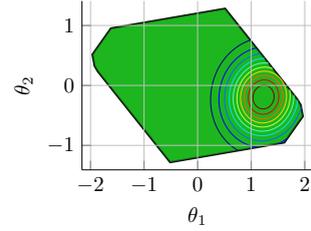
Similarly, the maximal confidence in either accepting or rejecting the system property is now given as³

$$\max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}.$$

³For brevity we dropped in the notation the dependence on the prior.

Remember that $\bar{\Theta}_\psi$ is the complement of Θ_ψ . To increase the maximal confidence,

Figure 7.3: The green region in the parameter space $[\theta_1 \ \theta_2]^T$ is the feasible set Θ_ψ of the case study. The contour lines give the density function of an a-posteriori distribution over the parameter space obtained based on a single experiment.



we will do an experiment on the system. For this, we know that when we measure the output $y_0(t) = \hat{y}(t, \theta^0)$, it will be disturbed by measurement noise $e(t)$. This noise is assumed to be a white additive measurement noise with known variance $\sigma_e^2 = 0.5$. Within an experiment we apply an input signal $u(t)$, and measure the output signal $\tilde{y}(t) = y_0(t) + e(t)$ over a given time horizon $N_s = 100$. The collection of input-output data $\mathbf{u}_{N_s}, \tilde{\mathbf{y}}$ can then be used to refine the maximal confidence

$$\max\{\mathbb{P}(\Theta_\psi | \mathbf{u}_{N_s}, \tilde{\mathbf{y}}), \mathbb{P}(\bar{\Theta}_\psi | \mathbf{u}_{N_s}, \tilde{\mathbf{y}})\}.$$

Of interest is the design of an experiment which optimises the expected maximal confidence. We will solve the design problem for an input $u(t)$ parameterised as a multi-sine with frequencies $(\omega_0, 2\omega_0, \dots, 5\omega_0)$ and fundamental frequency $\omega_0 = 2\pi/10$. Additionally, we consider an experiment set $\mathcal{E} = \{\mathbf{u} : u(t) \in [-0.2, 0.2] = \mathbb{U}\}$.

To analyse the impact of the experiment design method developed in this section, we compare it with the classical **D**-optimal experiment design [Gevers et al., 2011]. For **D**-optimality we minimise the determinant of R^+ subject to $\mathbf{u}_{N_s} \in \mathcal{E}$. We refer to the design method minimising the empirical approximation (7.15) of the experiment design problem (7.13) proposed in this section as Θ_ψ -optimal.

Both the Θ_ψ - and **D**-optimal experiment designs have been performed for priors with several different mean values μ , and with a fixed variance of $R = 0.2I_{2 \times 2}$. The empirical evaluations of the integral in (7.15) has been done with 1000 data points.

Note that the **D**-optimal experiment is independent of the prior mean; it simply optimises the determinant of R^+ based on R , which does not depend on the parameterisation due to its linearity. Of course, even though the **D**-optimal experiment inputs are equal for the different prior means, it is expected that the statistical evidence they provide for the verification of the property is still varying with the prior means.

After designing an experiment \mathbf{u}_{N_s} both for the Θ_ψ - and **D**-optimal experiment, the optimisation objective

$$\mathbb{E} [J(\tilde{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s}) | \tilde{\mathbf{y}}_{N_s} \sim p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s})] \quad (7.16)$$

is evaluated empirically. For this, 400 data samples $\tilde{\mathbf{y}}_{N_s}$ are drawn from the distribution $p(\tilde{\mathbf{y}}_{N_s} | \mathbf{u}_{N_s})$. This is done by first drawing a parameter value from the prior

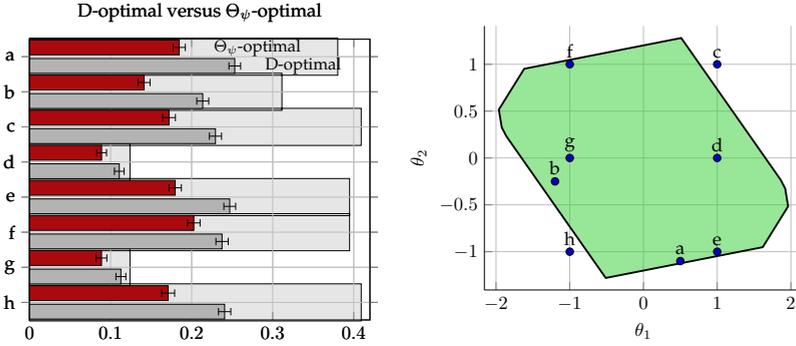


Figure 7.4: The left plot gives the empirical evaluation of $1 - \mathbb{E}[J]$ (as in (7.16)) on the horizontal axis for the safety verification case study for both the Θ_ψ -optimal (red bar, top) and the **D**-optimal experiment design (grey bar, below). The wider light grey bar gives $1 - \max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$. The mean values of the priors in the experiment design labeled on the vertical axis are depicted in the right plot. On the bars for the empirical evaluation of $1 - \mathbb{E}[J]$, their standard deviation is also drawn on the bars by the symbol \vdash .

distribution $\theta \sim p(\theta)$, and subsequently performing an identification experiment. In Figure 7.4 the empirical evaluation of $\mathbb{E}[J(\hat{\mathbf{y}}_{N_s}, \mathbf{u}_{N_s})]$ for both the Θ_ψ - and **D**-optimal experiment designs are plotted, together with the result obtained without performing additional experiments, i.e., $\max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$. Note that the figure displays the values $1 - \mathbb{E}[J]$ and $1 - \max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$ for convenience, and also provides the standard deviation of the empirical evaluations. Note that $1 - \max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$ and similarly $1 - \mathbb{E}[J]$ represents the lack of confidence in the optimal decision, that is how unsure we are about either accepting or rejection the system property.

In Figure 7.4, the result shows that the empirical value of $1 - \mathbb{E}[J]$ is lower for the Θ_ψ -optimal experiment design than for the **D**-optimal experiment design for all the chosen prior mean values. It can be observed that this is especially significant when $1 - \max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$ is large. Additionally, in this case the posterior variances of the Θ_ψ -optimal experiment design tend to align with the closest faces of the feasible region, as will be shown in more detail in the next case study. For mean values that are further from the boundaries of the feasible region such as $[1 \ 0]^T$ (**d**) and $[-1 \ 0]^T$ (**g**), the $\max\{\mathbb{P}(\Theta_\psi), \mathbb{P}(\bar{\Theta}_\psi)\}$ is already quite large and the difference between the Θ_ψ - and **D**-optimal design is less significant. Based on the empirical evaluations in this case study, we have elucidated that indeed an empirical optimisation of the classical experiment design problem can be performed. It has also been shown that in most cases this gives a significant improvement with respect to the **D**-optimal case.

7.4.b.2 Quasi-stationary solutions

By parameterising the input sequence, as done in the previous example, we can design experiments that are performed over short time horizons and take into

account all the transition effects. Next, we provide an alternative parameterisation of the signal $u(t)$, which is less suitable for short experiments as it does not take these transition effects into account. Still, for experiments performed over longer time horizons, it provides a computational beneficial approach. We consider the input signal \mathbf{u}_{N_s} to be the realisation of a quasi-stationary stochastic process [Ljung, 1999]. Under this assumption we can also model $x(t)$ to be a quasi-stationary process. This allows us to parameterise the update of the covariance matrix from R to R^+ directly, and hence reduce the dimensionality of the optimisation problem. More precisely, we consider the case where $\mathbf{u}_{N_s} = [u(0) \ u(1) \ \dots \ u(N_s - 1)]$ is a realisation of a quasi-stationary stochastic process (cf. Chapter 2) with power spectrum $\Phi_u(\omega)$. A direct link between this power spectrum and the information content of the data can be used to better parameterise the updated covariance matrix R^+ and thus also R_θ .

For a given parameterised model $\mathbf{M}(\theta)$, the information content of a set of measurements can be expressed by the Fisher information matrix [Cramér, 1946]. This information matrix quantifies how much information the data $\tilde{\mathbf{y}}_{N_s}$ carries about θ as follows

$$\mathcal{I} := \mathbb{E} \left[\left(\frac{\partial \ln p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s})}{\partial \theta} \right) \left(\frac{\partial \ln p(\tilde{\mathbf{y}}_{N_s} | \theta, \mathbf{u}_{N_s})}{\partial \theta} \right)^T \right]_\theta. \quad (7.17)$$

For the used linear parameterisation given in (7.1), \mathcal{I} is independent of the parameter and simplifies to

$$\mathcal{I} = \sigma_e^{-2} \Phi(\mathbf{u}_{N_s}) \Phi^T(\mathbf{u}_{N_s}) = \sum_{t=1}^{N_s} \sigma_e^{-2} x(t) x^T(t). \quad (7.18)$$

Using \mathcal{I} we can express the update of the covariance matrix as

$$(R^+)^{-1} = R^{-1} + \mathcal{I}. \quad (7.19)$$

Define $F(q^{-1})$ as the filter representation from $u(t)$ to $x(t)$ of (7.1), for which q^{-1} denotes the backward shift operator $q^{-1}x(t) = x(t-1)$. More precisely, let $x(t) = F(q^{-1})u(t)$ be a (quasi-)stationary signal, then the Fisher information matrix can be approximated with its expected value as a function of the distribution of \mathbf{u}_{N_s}

$$\mathcal{I} \approx \hat{\mathcal{I}} = \frac{N_s}{\sigma_e^2} \mathbb{E}[x(t) x^T(t)], \quad (7.20)$$

or equivalently, as a function of the input signal spectrum

$$\hat{\mathcal{I}} = \frac{N_s}{\sigma_e^2} \frac{1}{2\pi} \int_{-\pi}^{\pi} F(e^{j\omega}) F(e^{j\omega})^* \Phi_u(\omega) d\omega. \quad (7.21)$$

This allows us to analyse the properties of the experiment set-up averaged over time, neglecting transient effects. If the set of allowed experiments \mathcal{E} can be expressed based on stationary properties of \mathbf{u} , then we can replace the design of the

input sequence \mathbf{u}_{N_s} with the design of the desired power spectrum $\Phi_u(\omega)$. Examples of such sets include experiments with bounded input power $\mathcal{P}_u \leq c$ (see Section 2.4.a), that is $\mathcal{E} := \{\mathbf{u} \mid \mathcal{P}_u \leq c\}$.

The parameterisation of updates of the matrix $\hat{\mathcal{I}}$ has some beneficial properties. These will be detailed after the case study. To gain insight, we consider a very basic example also used in [Jansson, 2004]. With this example we can analyse solutions of the experiment design problem and show the applicability of the presented quasi-stationary framework.

Example 7.2 (Finite Impulse Response (FIR) model) Consider a system \mathbf{S} that can be modelled as

$$\hat{y}(t, \theta) = \theta_1 u(t) + \theta_2 u(t-1) + e(t)$$

or equivalently

$$\begin{aligned} x(t+1) &= \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} x(t) + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u(t) \\ \hat{y}(t, \theta) &= \theta^T x(t) \\ \tilde{y}(t, \theta) &= \hat{y}(t, \theta) + e(t) \end{aligned} \tag{7.22}$$

with $e(t)$ a white noise sequence. Note that there exists a θ^0 such that $\hat{y}(t, \theta^0)$ is the output of the true model of \mathbf{S} and $\tilde{y}(t, \theta^0)$ the measured output.

For this case study, we have selected a simple feasible set given as

$$\Theta_\psi := \text{conv}\{(-1, 0), (0, 1), (1, 0), (0, -1)\}.$$

This diamond shaped set based on the convex hull of $\{(-1, 0), (0, 1), (1, 0), (0, -1)\}$ is given as the grey area in Figure 7.5. We assume that some property of interest ψ is verified by models in this set Θ_ψ . We consider several different priors $\mathcal{N}(\mu, R)$ based on the mean values

$$\mu \in \left\{ (0, 0), \left(\frac{1}{2}, \frac{1}{2}\right), \left(\frac{1}{2}, -\frac{1}{2}\right), \left(-\frac{1}{2}, \frac{1}{2}\right), \left(-\frac{1}{2}, -\frac{1}{2}\right) \right\}.$$

and covariances $R \in \{0.15I, 0.05I\}$.

For each of the priors, we will design a stationary input optimising the empirical approximation (7.15) of the experiment design problem given in (7.13). We consider the design of a power spectrum Φ_u subject to a bounded signal power, that is $\mathcal{P}_u \leq c$, or equivalently $R_u(0) \leq c$.

For the stationary input, we compute the update (7.19) with the expected Fisher information matrix (7.21). We know that for a 2nd order FIR model, the first 2 autocorrelation coefficients of the input parameterise all information matrices in (7.21), as follows

$$\hat{\mathcal{I}} = \frac{N_s}{\sigma_e^2} \begin{bmatrix} R_u(0) & R_u(1) \\ R_u(1) & R_u(0) \end{bmatrix} \tag{7.23}$$

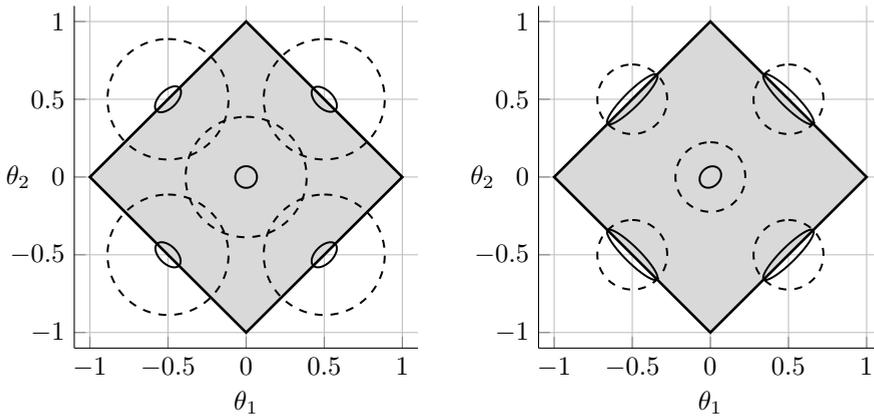
with $R_u(\tau)$ the auto-correlation of $u(t)$ as defined in Chapter 2.

As such the optimisation of the empirical objective (7.15), solved next, can be rewritten via the Fisher information matrix. We consider the case with $\frac{N_\varepsilon}{\sigma_\varepsilon^2} = 1$ and $c = 1$ where the set of possible Fisher information matrices can be represented as

$$\hat{\mathcal{I}} \in \left\{ R_u(0) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + R_u(1) \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \mid 0 \leq R_u(0) \leq 1 \text{ and } |R_u(1)| < R_u(0) \right\}; \quad (7.24)$$

for the given set, the auto-correlations correspond to a spectral density of $\Phi_u(\omega)$ that can be generated by a quasi-stationary signal [Jansson, 2004], this is guaranteed by taking $|R_u(1)| < R_u(0)$.

For an experiment with maximal power, that is $R_u(0) = 1$, we solve the empirical approximation (7.15) (with $M = 10^4$) to find the optimal $R_u(1)$. Results of these experiment design problems are portrayed in Figures 7.5 and 7.6 and explained next. With Figure 7.5 we will show the impact of the prior on the resulting update of the covariance matrix R to R^+ . Figure 7.6 depicts the implications of this input design on the distribution of the



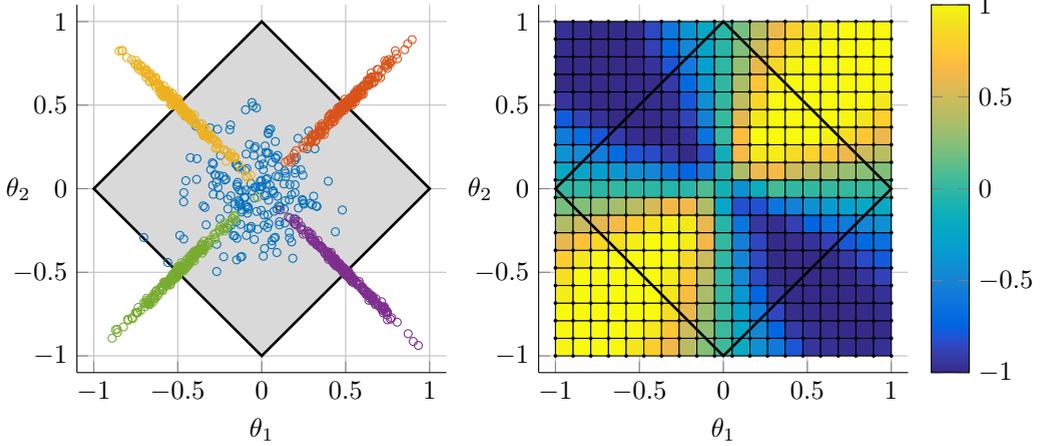
(a) First set of posterior covariance matrices after the optimisation with prior covariance $R = 0.15I$.

(b) Second set of posterior covariance matrices after the optimisation with prior covariance $R = 0.05I$.

Figure 7.5: Depicted are the results of the experiment optimisation for priors defined by $R \in \{0.05I, 0.15I\}$ and mean points $\mu \in \{(0, 0), (\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, -\frac{1}{2}), (-\frac{1}{2}, \frac{1}{2}), (-\frac{1}{2}, -\frac{1}{2})\}$. The figures both give the feasible set of interest in grey and represents the original Gaussian distribution of the five priors via a level set drawn based on the ellipse $\theta = \{\mu + R^{1/2}l \mid \|l\| = 1\}$ with a dashed line. The solid lines represent corresponding posterior probability density distribution with the (experiment dependent) posterior mean μ^+ shifted to the prior mean μ . Thus the resulting optimal R^+ for each of the mean values μ is drawn with a solid line at the original μ .

posterior mean values together with a more nuanced representation of the optimal design for a range of prior mean values.

First we analyse Figure 7.5. The dashed lines represent the different prior probability dis-



(a) Sample realisations μ^+ of $\mathcal{N}(\mu, R - R^+)$ at the different mean values of the prior, that is $\mu \in \{(0, 0), (\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, -\frac{1}{2}), (-\frac{1}{2}, \frac{1}{2}), (-\frac{1}{2}, -\frac{1}{2})\}$ for $R = 0.05I$ for the optimal experiment. The realisations are labelled colour-wise based on the respective mean value.

(b) The figure depicts the optimal experiment input, expressed by the optimal correlation $R_u(1)$, over a (gridded) range of initial mean values. More specifically, the axis θ_1 and θ_2 indicate the initial mean values as $\mu = [\theta_1, \theta_2]^T$.

Figure 7.6: Depicted are the results of the experiment optimisation for priors defined by $R = 0.05I$ and varying mean values. The left plot gives the sample realisations of the updated mean value after applying an optimal experiment. For the same experiment, the right figure gives the solution of the stationary experiment problem for varying prior mean values.

tributions. To show the optimal update of the covariance matrix R^+ relative to the original prior distribution, we have depicted the covariance matrix R^+ with the full lines centered at the original prior mean values. We observe that if the original R is relatively large as in Figure 7.5a, the optimal information matrix is a full rank excitation or equivalently, R^+ is strictly smaller than R . Still it can be noticed that the optimal R^+ is aligned with the edge of the feasible set at which the initial mean μ is located. Instead, when R is smaller as in Figure 7.5b, the initial confidence of μ at the edges of the feasible set is very close to 0.5 and as a consequence the optimal experiment is no longer full rank. This can be observed in Figure 7.5b by noting that the posterior covariance R^+ is not strictly smaller than R . This implies that an excitation that decreases the uncertainty orthogonal to the edge of the feasible set is preferred above a full rank excitation. This can most clearly be seen in the related Figure 7.6a, where the optimal $\Phi_u(\omega)$ (or equivalently $R_u(1)$) is such that the next $\mu^+ \sim \mathcal{N}(\mu, R - R^+)$ has a probability distribution perpendicular to the local edge.

While Figure 7.5 might suggest that the confidence in satisfying the property remains fixed to 0.5 (as all level sets are symmetrically situated on the edges of the feasible set) this is not the case since the update means μ^+ are not indicated yet. Figure 7.6 gives next to the distribution of $\mu^+ \sim \mathcal{N}(\mu, R - R^+)$ in Figure 7.6a also a more nuanced image of the optimal $R_u(1)$ in Figure 7.6b. For this the optimal $R_u(1)$ is computed for 20 times 20

prior mean values and $R = 0.05I$.

Let us analyse the impact of the possible choices for $R_u(1)$. For $R_u(1) \rightarrow 1$ the distribution of μ^+ , that is, $\mathcal{N}(\mu, R - R^+)$ converges to a Gaussian on the $(1, 1)$ line, perpendicular to the $(1, -1)$ direction of the edges at $(-\frac{1}{2}, -\frac{1}{2})$ and $(\frac{1}{2}, \frac{1}{2})$ as depicted in Figure 7.6a. In contrast for $R_u(1) \rightarrow -1$ the distribution becomes perpendicular to the edges at $(-\frac{1}{2}, \frac{1}{2})$ and $(\frac{1}{2}, -\frac{1}{2})$. We note that the \mathbf{D} -optimal experiment is the experiment with $R_u(1) = 0$. Now if we observe Figure 7.6b, we see that these \mathbf{D} -optimal experiments are preferred at the horizontal and vertical cross-section of the feasible set. Starting from the \mathbf{D} -optimal choice at the origin and going towards the edges through $(-\frac{1}{2}, -\frac{1}{2})$ and $(\frac{1}{2}, \frac{1}{2})$ again we get $R_u(1) \rightarrow -1$. Similarly for the other edges, the closeness to the middle of an edge defines $R_u(1)$.

In the next section, we consider the online formulation of objective \mathbf{B} . There we will use a continuation of this example (cf. Example 7.5) to draw a comparison between the design of stationary inputs based on the \mathbf{D} -optimal criterion, the utility-based criterion of this section (cf. objective \mathbf{A}); and online solution to objective \mathbf{B} .

Beyond the simple model structures, such as the FIR model utilised in the example, where the parameterisation of the Fisher information matrix follows naturally, results from inter alia [Hildebrand and Gevers, 2003, Zarrop, 1979] can be used to parameterise the Fisher information matrix for a given set of stationary inputs. We consider some basic theory on the Fisher information matrix, useful for the optimal design of experiments. But first, we define the average per sample information matrix $\bar{\mathcal{I}}$

$$\bar{\mathcal{I}} = \frac{1}{\sigma_e^2} \mathbb{E} [x(t)x^T(t)] \quad (7.25)$$

with $\hat{\mathcal{I}} = \sum_{t=1}^{N_s} \bar{\mathcal{I}}$. Now $\bar{\mathcal{I}}$ can be written in the form

$$\bar{\mathcal{I}} = \sum_{k=1}^n x_k M_k$$

where M_k is a symmetric matrix that depends on the dynamics (7.1) and x_k is a function of the dynamics (7.1) and the used power spectrum. Practically, this means that we only need n parameters to span the design variables. Based on the theory developed and represented in inter alia [Hildebrand and Gevers, 2003, Zarrop, 1979, Gevers et al., 2011] the set of x_k for which there exists at least one input spectrum can be represented efficiently based on convex sets and LMIs. Further we can rewrite bounds on the inputs, that is expressions for \mathcal{E} , as bounds on x_k . This implies that even for model structures that are more complex than the FIR model, we can use the stationary design to first optimise the experiment with respect to the Fisher information matrix based on x_k .

When we find an optimal selection of x_k , we still need to find the corresponding power spectrum and input signal. A complete explanation of how to parameterise the information matrix and how to compute the resulting quasi-stationary input signal $u(t)$ is beyond the scope of this work. A more complete summary of this topic can be found in [Larsson, 2014, Jansson, 2004].

7.5 Maximising probabilistic reachability of confidence threshold (online)

Introduction: towards short experiments

Experiments are often very expensive. Hence in this section, we will formulate the design problem to address the question how to do an – as short as possible – experiment while still attaining the required statistical evidence to make a decision for a given formal verification problem. To this end we will newly allow for the use of measured data to online select the experiment input. The idea to finish with collecting samples when the required evidence has been collected is not new. Especially within the framework of hypothesis testing, this topic has been investigated. Within statistics, common methods include the sequential hypothesis test as introduced by Wald [1945]. In this section, we will consider the objective to reach a certain confidence in a decision to accept or reject a property.

Notice that this objective can be extended to objectives defined based on sequential hypothesis testing like the sequential probability ratio testing of Wald [1945]. The applicability of these kinds of testing, often necessary for bounding of the type I and type II errors, depends on the hypothesis of interest and the distributions. For the moment, we restrict ourselves to the sequential implementation of experiments on dynamical systems for the efficient collection of a certain amount of statistical evidence, expressed via the confidence.

In this section, we formalise the problem statement for the design of short experiments with as goal the formal verification of systems. That is instead of considering the criterion **A** in Section 7.2.c, we now consider objective **B**. We seek for the experiment input that optimises the probability that a given threshold confidence is reached during the experiment. Contrary to the previous section, we now allow for the online choice of input actions based on the available measurements history. For this, we first model the experiment as a Markov decision process (MDP). The idea of reformulating the experiment design problem as an MDP optimisation originates from [Larsson, 2014], and is newly modelled by adding next to the state of the dynamical system also the posterior distribution of the model parameters into the state of the MDP.

The stochastic transitions of this MDP are a combination of the state update and the Bayesian inference over the uncertainty distribution. The inclusion of the posterior distribution as a state of an MDP within an experiment design problem has also been formulated by [Huan and Marzouk, 2016]. The work of Huan and Marzouk [2016] focusses on Bayesian experiment design for learning static functions or variables and considers optimisation criteria based on information metrics more standard in machine learning. Instead, our MDP includes the dynamic nature of the experiment set-up caused by the state transitions of the dynamical system and considers as objective the attainment of a confidence threshold.

In the sequel, we approach the experiment design problem as follows,

- first, we model the experiment as an Markov decision process (cf. Section 7.5.a);

- then in Section 7.5.b, we consider the design criterion quantified as the probability that a given confidence threshold is reached before the end of the experiment and show that it can be computed as a stochastic reach-avoid probability for this MDP;
- in the same subsection, we subsequently give the experiment design problem as a stochastic reach-avoid problem, which is a stochastic optimal control problem solvable via dynamic programming.

In last two parts of this section, we consider the computational issues. In Section 7.5.c, we discuss the computation of solutions to the presented optimal control problem. In the last part, cf. Section 7.5.d, we show that by building on results for quasi-stationary inputs, as presented in Section 7.4.b.2, we can reduce the complexity of the experiment design problem. We use this design to also make a comparison between the results for D-optimal, for the experiment criteria based on maximal confidence (cf. objective A) and for the objective based on the maximal probabilistic reachability of a given confidence threshold (cf. objective B), presented in this section.

7.5.a Online experiment design as a Markov decision process

We consider a Markov decision process defined over a Euclidean state space with Borel measurable stochastic kernels.

Definition 7.1 (Markov decision process) *A discrete-time MDP, denoted as $\Sigma = (\mathbb{X}, \pi_{x(0)}, \mathbb{T}, \mathbb{U})$, is comprised of: a continuous (uncountable) state space $\mathbb{X} \subset \mathbb{R}^n$; an action space \mathbb{U} consisting of a possibly uncountable number of actions; a Borel-measurable stochastic kernel \mathbb{T} , which assigns to each state-action pair $x \in \mathbb{X}$ and $a \in \mathbb{U}$ a probability distribution $\mathbb{T}(\cdot | x, a)$ over \mathbb{X} ; an initial probability distribution is $\pi_{x(0)} : \mathcal{B}(\mathbb{X}) \rightarrow [0, 1]$.*

For a given parameter θ , model $\mathbf{M}(\theta)$ in (7.1) can be regarded as a discrete-time Markov process, expressed at time k with a deterministic transition corresponding to a Dirac distribution (a point distribution), namely $\mathbb{T}(dx' | x, u) = \delta_{Ax+Bu}(dx')$.

We would like to extend the model $\mathbf{M}(\theta)$ to also include the available knowledge on θ . Before the start of the experiment, this available knowledge is structured into a prior distribution $\mathcal{N}(\mu_0, R_0)$ over the parameter space, that is the prior is again a multi-variate Gaussian with mean μ_0 and variance R_0 . Of interest to us is a Markov representation of the Bayesian inference procedure, where the posterior distributions in (7.3) can be computed recursively. More precisely, a prior $p(\theta) = \mathcal{N}(\mu, R)$ is updated via Bayesian inference from system-drawn data $\{\tilde{\mathbf{y}}_t, \mathbf{u}_t\}$ up to time t . Both the resulting posterior probability distribution $p(\theta | \tilde{\mathbf{y}}_t, \mathbf{u}_t) = \mathcal{N}(\mu^+, R^+)$, the data realisation $\tilde{\mathbf{y}}_t$, and the unknown true parameter θ^0 can then be described by random variables with Gaussian distributions as given in the equations in (7.4).

Of interest to us is the iterative updates of the posterior as depicted in Figure 7.7. Employing the Gaussian distributions of (7.4) we can rewrite the data collection

above via data sets of length one ($N_s = 1$) as an MDP. Thus we have a Markov process with stochastic transitions

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t), \\ \mu(t+1) &= \mu(t) + R(t)x(t+1)v(t), & v(t) &\sim \mathcal{N}(0, \Gamma(t)) \\ R(t+1) &= R(t) - R(t)x(t+1)\Gamma(t)x(t+1)^T R(t), \end{aligned} \quad (7.26)$$

driven by the Gaussian realisations $v(t)$ with covariance $\Gamma(t) = (\sigma_e^2 + x(t+1)^T R(t)x(t+1))^{-1}$. We refer to this model as the experiment MDP (XMDP) Σ_θ (as opposed to $\mathbf{M}(\theta)$ in (7.1)) for later use. Note that the covariance matrix is symmetric, that is, $R \in \mathbb{S}^n$, with \mathbb{S}^n the set of real symmetric matrices $R = R^T \in \mathbb{R}^{n \times n}$. Therefore it can be uniquely defined by its upper triangular elements, denoted by $r \in \mathbb{R}^{(n+1)n/2}$. This means that there exists a one-to-one mapping from the variances in matrix $R \in \mathbb{S}^n$ to points $r \in \mathbb{R}^{(n+1)n/2}$, that is $f_R : \mathbb{R}^{(n+1)n/2} \rightarrow \mathbb{S}^n$ and $f_r : \mathbb{S}^n \rightarrow \mathbb{R}^{(n+1)n/2}$, where $f_r = f_R^{-1}$. For a given mapping f_r , the MDP Σ_θ has as state space \mathbb{X}_θ with elements $x_\theta = (x, \mu, r) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^{n(n+1)/2}$ and takes as input signals $u \in \mathbb{U}$.

Given a prior distribution $p(\theta) := \mathcal{N}(\mu(0), R(0))$, $\mu(0) = \mu_0$, $R(0) = R_0$; and given an initial state for $\mathbf{M}(\theta)$, $x(0) = x_0$; the initial state of the MDP is given as $x_\theta(0) = (x_0, \mu, f_r(R))$. More precisely, we consider the initialisation of the MDP with the Dirac distribution $\pi_{x_\theta(0)} := \delta_{x_\theta(0)}$. At every time instant the input can be selected as a function of the state x_θ of the MDP.

Definition 7.2 (Markov policy) *A Markov policy π over the horizon $[0, N_s - 1]$ is a sequence $\pi = (\pi_0, \pi_1, \dots, \pi_{N_s-1})$ of measurable maps, $\pi_k : \mathbb{X}_\theta \rightarrow \mathbb{U}$, $k = 0, 1, \dots, N_s - 1$, from the state space \mathbb{X}_θ to the action space \mathbb{U} . The set of Markov policies is denoted as Π .*

Remark 7.3 *We refer the reader to [Abate et al., 2008] for a complete discussion on measurability issues related to Markov decision processes over continuous state and control spaces, and to corresponding optimal control problems. In this work we refer for simplicity to general measurability requirements, as in the previous statement.*

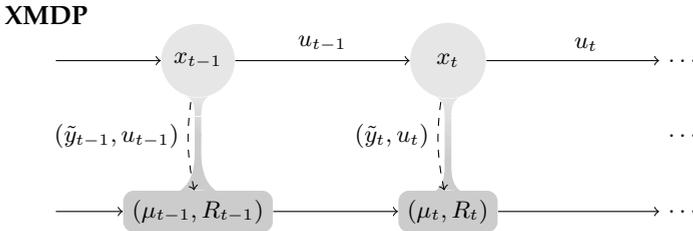


Figure 7.7: Schematic representation of transitions of the experiment, which are modelled as a Markov decision process

Remark 7.4 *We would like to emphasise that we use the notion of Markov decision process, which has also been used in Chapter 3. In that part of the thesis, Markov decision processes have been used to model the stochastic state transitions of a physical system. Instead in this part of the thesis, we have assumed that the physical system has deterministic state transitions and the Markov decisions process is used to model the combination of the physical system and our knowledge on the system. In Chapter 3, we considered abstraction techniques that will allow us to do controller synthesis with quantified accuracy. The properties of interest include the stochastic reach-avoid property tackled in this section and are defined directly on the output space of a generalised MDP. In this section, we will use the experiment MDP for the analysis of the experiment and properties of the MDP are therefore not directly related to properties of the physical system.*

The goal of this work is the design of a Markov policy for Σ_θ , to attain an efficient data collection for the verification (with a given confidence level) of a property defined over the unknown system \mathbf{S} .

7.5.a.1 Experiment setup

Suppose we perform an experiment on \mathbf{S} with a given policy $\pi \in \Pi$. At the start of the experiment, the MDP Σ_θ is initiated as $x_\theta(0) = (x(0), \mu(0), f_r(R_0))$, where $x(0)$ is the initial state of \mathbf{S} and where the Gaussian distribution $\mathcal{N}(\mu_0, R_0)$ represents the prior uncertainty distribution of θ_0 . At $t = 0$ the control input $u(0)$ is selected based on the policy $u(0) = \pi_0(x_\theta(0))$. The subsequent transition from $x_\theta(0)$ to $x_\theta(1)$ is governed by deterministic transitions for $x(0)$ to $x(1)$ and for $r(0)$ to $r(1)$, and by a stochastic transition for $\mu(0)$ to $\mu(1)$ obtained from the measured output $\tilde{y}(1)$ as in (7.4c). At every subsequent time instant the input is chosen based on the current state and the stochastic part of the transition is obtained as a function of the measured output drawn from \mathbf{S} . For the latter, the MDP Σ_θ gives the uncertainty distribution of this transition based on the current state of the MDP $x_\theta(t) = (x(t), \mu(t), r(t))$. Remember that the current state of the system represents the collection of past measurements $\tilde{\mathbf{y}}_t$.

7.5.b Optimal experiment design and stochastic optimal control

The experiment is successfully completed at time $t \in \mathbb{N}$ when $\theta_0 \in \Theta_\psi$ or $\theta_0 \in \Theta \setminus \Theta_\psi$, with associated confidence of at least $1 - \delta$, and where $\theta_0 \sim \mathcal{N}(\mu(t), R(t))$, with $R(t) = f_R(r(t))$. Denote set $K \subset \mathbb{X}_\theta$ as the set of states (x, μ, r) associated with the required confidence on $\theta_0 \sim \mathcal{N}(\mu, f_R(r))$. Hence a given state trajectory $\{x_\theta(t) \mid 0, 1, \dots, N_s\}$ represents a successful experiment if it reaches the target set K . This reachability property can be evaluated over a trajectory as

$$\exists j \in [0, N_s] : x_\theta(j) \in K.$$

If in addition the state of the MDP is required to stay within a given safe set A , then the success of an experiment within the finite horizon becomes equivalent

to a finite-horizon reach-avoid (or constrained reachability) property [Baier and Katoen, 2008] over a safe set A and target set K . This can be expressed as

$$\exists j \in [0, N_s] : x_\theta(j) \in K \wedge \forall i \in [0, j-1] : x_\theta(i) \in A \setminus K.$$

The probability associated to this finite-time reach-avoid event can be characterised as a boolean expression using indicator functions [Abate et al., 2008, Summers and Lygeros, 2010], which leads to an expectation over the state trajectories as

$$r_{x_\theta(0)}^\pi(K, A) = \mathbb{E}_{x_\theta(0)}^\pi \left[\sum_{j \in [0, N_s]} \mathbf{1}_K(x_\theta(j)) \prod_{i=0}^{j-1} \mathbf{1}_{A \setminus K}(x_\theta(i)) \right],$$

where $\mathbf{1}_B(x) = 1$ if $x \in B$, and otherwise it is equal to 0.

We can now state that an experiment is successful if during the experiment the given threshold confidence $1 - \delta$ is reached⁴ and if until then it stays within some safe set of operating conditions. As anticipated in criterion **B** in Section 7.2.c, we take as experiment design criterion the probability of this reach-avoid event. Thus for the experiment MDP and for a given policy π , we can quantify the criterion with the given reach-avoid property $r_{x_\theta(0)}^\pi(K, A)$, where

- K is the reach set or target set, defined by the required confidence threshold for accepting or rejecting a formal property;
- A is a safe set in which the experiment trajectory needs to stay. This set is optional and can be selected as $A := \mathbb{X}_\theta$. Alternatively, requirements on the safe operating range including limits on the state amplitude $x(0)$, can be encoded based on this safe set A .

To enforce reaching a decision before the end of the experiment, let us *penalise* the time it takes to achieve the required confidence. Whilst above we have attached a value of 1 to each trajectory reaching the target set, let us instead consider a discount factor $\gamma \in (0, 1)$ attached to each time step that x_θ stays in $A \setminus K$ before reaching K . Denoting $\mathbf{1}_B^\gamma(x) = \gamma \mathbf{1}_B(x)$, the expression for the experiment design objective becomes

$$s_{x_\theta(0)}^\pi(K, A) = \mathbb{E}_{x_\theta(0)}^\pi \left[\sum_{j \in [0, N_s]} \mathbf{1}_K(x_\theta(j)) \prod_{i=0}^{j-1} \mathbf{1}_{A \setminus K}^\gamma(x_\theta(i)) \right].$$

In the sequel we refer to $s_{x_\theta(0)}^\pi(K, A)$ as the *discounted reach-avoid property*. Notice that unlike the reach-avoid probability $r_{x_\theta(0)}^\pi(K, A)$, the discounted equivalent $s_{x_\theta(0)}^\pi(K, A)$ is not induced by a probability over the trajectories of Σ_θ . Still the discounted quantity $s_{x_\theta(0)}^\pi(K, A)$ over the traces of Σ_θ can be written as a reach-avoid probability $r_{\bar{x}_\theta(0)}^\pi(K, A)$ for an extended MDP $\bar{\Sigma}_\theta$, which includes a transition probability of $(1 - \gamma)$ to model the possibility that the experiment is terminated preemptively. $\bar{\Sigma}_\theta$ extends the state space of Σ_θ with two discrete modes:

⁴Note that there is a difference between the Dirac distribution at a point x δ_x and the $\delta \in [0, 1]$ we use to quantify the confidence threshold.

q_1 for the experiment being *active* and q_2 for the experiment being *interrupted preemptively*. Note that for $\gamma = 1$, $s_{x_\theta(0)}^\pi(K, A) = r_{x_\theta(0)}^\pi(K, A)$.

For a given policy π , the time-dependent value function $\mathbf{W}_k^\pi : \mathbb{X}_\theta \rightarrow [0, 1]$, defined as

$$\mathbf{W}_k^\pi(x_\theta) = \mathbb{E}^\pi \left[\sum_{j \in [k+1, N_s]} \mathbf{1}_K(x_\theta(j)) \prod_{i=k+1}^{j-1} \mathbf{1}_{A \setminus K}^\gamma(x_\theta(i)) \mid x_\theta(k) = x_\theta \right],$$

is the γ -discounted probability that the state trajectory $\{x_\theta(k+1), \dots, x_\theta(N_s)\}$, starting from $x_\theta(k)$, will reach the target set K within the time horizon $[k, N_s]$, while staying within the safe set A . This function allows expressing the discounted reach-avoid probability backward recursively, as follows.

Proposition 7.3 *Given a policy $\pi = (\pi_0, \pi_1, \dots, \pi_{N_s-1})$, define function $\mathbf{W}_k^\pi : \mathbb{X}_\theta \rightarrow [0, 1]$ by backward recursion*

$$\mathbf{W}_k^\pi(x_\theta) = \mathbb{E}_{x_\theta}^{\pi_k} \left[\mathbf{1}_K(x_\theta^{t+1}) + \mathbf{1}_{A \setminus K}^\gamma(x_\theta^{t+1}) \mathbf{W}_{k+1}^\pi(x_\theta^{t+1}) \right],$$

with the compact notation $x_\theta^{t+1} \sim \mathbb{T}(\cdot \mid x_\theta, \pi_k(x))$ for $k = N_s - 1, N_s - 2, \dots, 0$, and initialised with $\mathbf{W}_{N_s}^\pi(x_\theta) = 0$. Then for any initial state $x_\theta(0) \in \mathbb{X}_\theta$, the discounted probabilistic reach-avoid property $s_{x_\theta(0)}^\pi(K, A)$ is

$$s_{x_\theta(0)}^\pi(K, A) = \mathbf{1}_K(x_\theta(0)) + \mathbf{1}_{A \setminus K}^\gamma(x_\theta(0)) \mathbf{W}_0^\pi(x_\theta(0)).$$

Proof: The proof follows [Summers and Lygeros, 2010, Lemma 4], where the above statement is proven for a value function $V_k^\pi(x) = \mathbf{1}_K(x) + \mathbf{1}_{A \setminus K}(x) W_k^\pi(x)$. To allow for the discounting, extend the state space with two discrete modes as described above. Consider an extended safe set $\bar{A} := \{q_1\} \times A$ and a target set $\bar{K} := \{q_1, q_2\} \times K$. Let the probability of going from q_1 to q_2 be $1 - \gamma$ for any continuous state in $A \setminus K$. \square

Rather than selecting and fixing a policy π , as done above, we now focus on the optimal control problem, which seeks the Markov policy π^* that maximises the discounted probabilistic reach-avoid property, and which is such that $s_{x_\theta(0)}^*(K, A) = \sup_\pi s_{x_\theta(0)}^\pi(K, A)$. This optimal policy can be characterised as follows.

Proposition 7.4 *Define functions $\mathbf{W}_k^* : \mathbb{X}_\theta \rightarrow [0, 1]$, by the backward recursions*

$$\mathbf{W}_k^*(x_\theta) = \sup_{u \in \mathbb{U}} \mathbb{E}_{x_\theta}^u \left[\mathbf{1}_K(x_\theta^{t+1}) + \mathbf{1}_{A \setminus K}^\gamma(x_\theta^{t+1}) \mathbf{W}_{k+1}^*(x_\theta^{t+1}) \right],$$

with $x_\theta^{t+1} \sim \mathbb{T}(\cdot \mid x_\theta, u)$ for $k = N_s - 1, N_s - 2, \dots, 0$, and initialised by $\mathbf{W}_{N_s}^*(x_\theta) = 0$. Then for any initial state $x_\theta \in \mathbb{X}$ the optimal probabilistic reach-avoid property $s_{x_\theta(0)}^*(K, A)$ can be expressed as

$$s_{x_\theta(0)}^*(K, A) = \mathbf{1}_K(x_\theta(0)) + \mathbf{1}_{A \setminus K}^\gamma(x_\theta(0)) \mathbf{W}_0^*(x_\theta(0)).$$

Furthermore, $\pi_k^* : \mathbb{X}_\theta \rightarrow \mathbb{U}$ for $k = N_s - 1, N_s - 2, \dots, 0$, is such that $\forall x_\theta \in \mathbb{X}_\theta$:

$$\pi_k^*(x_\theta) = \arg \max_{u \in \mathbb{U}} \mathbb{E}_{x_\theta}^u \left[\mathbf{1}_K(x_\theta^{t+1}) + \gamma \mathbf{1}_{A \setminus K}(x_\theta^{t+1}) \mathbf{W}_{k+1}^*(x_\theta^{t+1}) \right]$$

and $\pi^* = (\pi_0^*, \pi_1^*, \dots, \pi_{N_s-1}^*)$ is the optimal Markov policy for the discounted probabilistic reach-avoid.

With the above proposition, we can now maximise the experiment design criterion. Moreover, we have detailed the optimisation of the experiment inputs with respect to reaching a given confidence threshold as a stochastic optimal control problem. Additionally, this optimal control problem has been formulated recursively, that is, via dynamic programming.

Proof: The proof follows the proof for not discounted probabilistic reach-avoid properties given in [Abate et al., 2008] and [Summers and Lygeros, 2010, Theorem 6]. \square

The above proposition shows that there exists an optimal policy for the discounted reach-avoid problem that is deterministic. This implies that for every state x_θ of Σ_θ the optimal policy delivers a single control input, instead of a probability distribution over the set of possible control actions.

The computation of $s_{x_\theta}^*(K, A)$ is based on N_s recursions given in Proposition 7.4, which can be denoted by a *dynamic programming operator* \mathbf{T} , that is

$$\mathbf{W}_k^* = \mathbf{T} \mathbf{W}_{k+1}^*.$$

Therefore the value of the optimal γ -discounted probabilistic reach-avoid property can be written as the composition of N_s mappings as

$$s_{x_\theta(0)}^*(K, A) = \mathbf{1}_K(x_\theta(0)) + \gamma \mathbf{1}_{A \setminus K}(x_\theta(0)) (\mathbf{T}^{N_s} \mathbf{W}_{N_s}^*)(x_\theta(0)).$$

Let us qualitatively comment on the behaviour of the backwards recursions from \mathbf{W}_{k+1}^* to \mathbf{W}_k^* . Employing a γ -discounting, \mathbf{T} is a contractive mapping [Tkachev and Abate, 2011, Theorem 4], which allows tapping on typical results in stochastic optimal control [Bertsekas and Shreve, 1996]. As a result of this contractivity property, the mapping $\mathbf{T}^{N_s} \mathbf{W}_{N_s}^*$ will converge, for increasing values of N_s , to a unique value function associated to a corresponding infinite horizon problem. We define this value function as

$$\mathbf{W}^* := \lim_{N_s \rightarrow \infty} \mathbf{T}^{N_s} \mathbf{W}_{N_s}^*,$$

and, additionally, we define the *infinite horizon* γ -discounted probabilistic reach-avoid property as

$$\bar{s}_{x_\theta(0)}^*(K, A) = \lim_{N_s \rightarrow \infty} \mathbf{1}_K(x_\theta(0)) + \gamma \mathbf{1}_{A \setminus K}(x_\theta(0)) (\mathbf{T}^{N_s} \mathbf{W}_{N_s}^*)(x_\theta(0)).$$

The probability $\bar{s}_{x_\theta(0)}^*(K, A)$ of the infinite horizon reach-avoid property represents for the actual experiment, where reaching a certain confidence threshold is

the goal, the discounted probability of reaching a decision sometime the future. For problems over a long time horizon N_s , we expect to also obtain a stationary policy [Bertsekas and Shreve, 1996], leading to the practical use of a time-independent and deterministic policy for the experiment design problem.

7.5.c Computing the optimal online experiment

Although we have attained a formal characterisation of the experiment design problem by a stochastic optimal control formulation, the computation of the exact solution is seldom analytical. This is not only because the backwards recursions \mathbf{T} in general cannot be expressed explicitly, but also because the target set K will often not have an analytical expression. Instead, the associated stochastic dynamic programming problem given in Proposition 7.4 ought to be solved approximately.

The soundness of the data-driven verification, performed based on the experiment results, is not related to any approximation error incurred in solving the experiment design problem. More precisely, approximation errors in the computation and optimisation of the probability that a threshold confidence is reached, quantified by $s_{x_\theta(0)}^*(K, A)$, do not influence the processing and interpretation of the obtained data in the data-driven verification. They only change the optimality and hence the time that the experiment will take to complete. Additionally, the numerically heavy computations of π^* can be performed offline before the start of the actual experiment, as such they still allow for a real time online implementation of the experiment input.

Still, even if we allow for approximation errors and compute the solutions with approximate dynamic programming, several factors will still induce high – if not an crippling – computational cost. Within approximate dynamic programming the value functions \mathbf{W}_k^* can either be approximated by a gridding or discretisation over the state-space [Esmaeil Zadeh Soudjani and Abate, 2011] or by smooth functions like neural networks [Busoniu et al., 2010]. The dynamic programming operator, \mathbf{T} , can point-wise be evaluated based on numerical computations or by sample-based algorithms. The inductive computation of value functions based on the dynamic programming operator, \mathbf{T} , is referred to value iteration. Advance application of either value iterations or other methods like Q -iteration or policy iteration [Busoniu et al., 2010] can mitigate the curse of dimensionality and can, hence, also be used to solve this optimal control problem. In this subsection, we will showcase the application of approximate dynamic programming on a simple case study with a single parameter. In the next subsection, we will leave further exploration of approximate solution to the stochastic optimal control problem to the side. Instead, we analyse whether the computational complexity of the problem can be lowered further by first abstracting or reducing the complexity of the experiment MDP. More precisely, we reduce this MDP to one only containing the updates of the uncertainty distribution defined by μ and R . For this, we will use the same reasoning as in Section 7.4.b.2. That is, we again leverage results for stationary input signals. By doing this we can lower the complexity of the problem significantly. We obtain an MDP evolving over a $2n$ dimensional state space controlled by inputs from a n -dimensional space, in contrast to the original MDP

that had $2n + n(n + 1)/2$ states. For this reduced MDP we can again formulate the reaching of a threshold confidence as the reaching of a target set.

We perceive the development of insight into the formulated Bayesian experiment design as key. By reducing the case study, only the most important dynamics are retained, allowing for a distilled image of the problem.

Also within the next subsection, we continue the case study 7.2 with the novel online-implementation. This continuation is used to reflect on the differences between D-optimal design, open-loop designs of the previous section and the introduced online-design methodology.

Analysis of the experiment design problem for one parameter: Case study with approximate dynamic programming

In this section, we will consider a case study with a one-dimensional parameterised model. This allows us to analyse and clarify in depth the stochastic optimal control reformulation of the experiment design problem. We will also use it to show the absence of well-posedness issues present in the frequentist formulation of the experiment design problem for hypothesis testing in Section 7.3.d. This one-dimensional problem leads to an optimisation problem over a three-dimensional MDP.

Consider the parameterised model:

$$\mathbf{M}(\theta) : x(t + 1) = \frac{1}{2}x(t) + u(t), \quad y_0(t) = \theta x(t),$$

with measurements taken as

$$\tilde{y}(t) = y_0(t) + e(t), \quad e(t) \sim \mathcal{N}(0, 0.5).$$

Assume that the feasible set of parameters is given as $\Theta_\psi = [-1, 1]$ (the details of property ψ not being of interest here). The objective is to design an experiment such that the confidence in the decision whether $\mathbf{S} \models \psi$ or $\mathbf{S} \not\models \psi$ based on the posterior probability distribution is at least $1 - \delta$, with $\delta = 0.01$. Consider a given set of allowed inputs for the experiment, $\mathcal{E} = \{-1, 0, 1\}^{N_s}$, and a maximal experiment time of $N_s = 10$. Since the state transitions of $\mathbf{M}(\theta)$ are strictly stable, no additional requirements have to be imposed on the allowed range of $x(t)$. Hence the safe set is chosen as $A = \mathbb{X}_\theta$. At the start of the experiment the prior uncertainty is given as $\mathcal{N}(\mu_0, R_0)$, i.e. $\mu(0) := \mu_0$ and $R(0) := R_0$, and the system is initialised at $x(0) = x_0$ which, with no loss of generality, is assumed to be $x_0 = 0$.

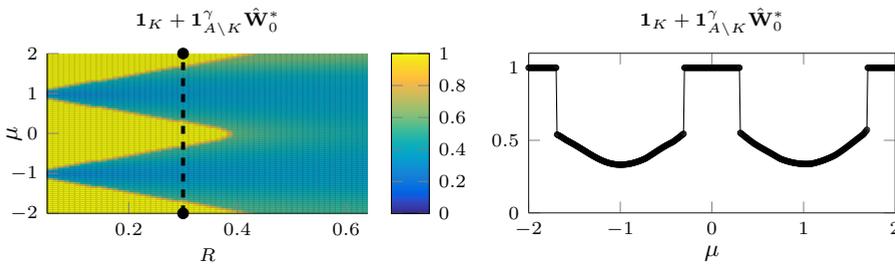
To objective of the experiment input is to maximise the reach-avoid probability defined with A and K , where the latter is defined based on the threshold confidence of $1 - \delta$. We can reformulate this experiment design problem to a discounted probabilistic reach-avoid problem over an experiment MDP. Again we stop the experiment when a confidence of at least $1 - \delta$ is attained. We select a rather high discount factor $\gamma = 0.6$, since we expect a short average experiment time.

Based on the proposed MDP reformulation of the experiment design problem we obtain a three-dimensional MDP with state (x, μ, R) for the discounted reach-avoid problem. We can now solve the stochastic optimal control problem (cf. Proposition 7.4) approximately via fitted value iteration as introduced in [Munos and Szepesvári, 2008] and used for stochastic reach-avoid problems in [Haesaert et al., 2014]. In this algorithm, the integration action – defined by the expected value operation – in the backwards mapping is replaced by samples and the value functions are fitted with a neural network. We have chosen a network with 2 layers, with respectively 15 and 10 neurons, and tansig functions for all neurons.

To be more precise, first a set of trajectories is generated over $A \setminus K$, this gives a collection of $N = 10^4$ states $\{x_\theta^j\}_{j \leq N}$. Then at each recursion for each of these states the value of the backwards mapping is empirically estimated with M samples ($M = 10$), after which the neural network is fitted over the estimates. More details on this simple manner of solving the stochastic reach-avoid problem can be found in [Haesaert et al., 2014].

Within the 10 backwards mappings applied to obtain the outcomes plotted in Figure 7.8b and 7.8a, the approximate value function iteration has already converged. Notice that the approximate function is denoted by $\hat{\mathbf{W}}_0^*$, in contrast to the exact value function \mathbf{W}_0^* . Extending the time horizon beyond $N_s = 10$ will not change the value function $\hat{\mathbf{W}}_0^*$, hence the value functions $\hat{\mathbf{W}}_0^*$ in Figures 7.8b and 7.8a additionally display the approximate solution to an infinite horizon problem.

Figure 7.8a⁵ represents $\mathbf{1}_K + \mathbf{1}_{A \setminus K}^\gamma \hat{\mathbf{W}}_0^*$. As such it gives an approximate evaluation of the optimal reach-avoid probability, where the latter represents the optimisation criterion (cf. Proposition 7.4). Notice here that for decreasing variance an increasing amount of the state space has a value equal to one. These are the states for which $\mathbf{S} \models \psi$ can be accepted or rejected with confidence at least $1 - \delta = 0.99$.



(a) Surface plot of the γ -discounted reach-avoid probability for $x(0) = 0$ evaluated over different values of the prior mean $\mu(0)$ and covariance $R(0)$.

(b) Slice of the discounted optimal reach-avoid probability over the values of μ for the given parameters $R(0) = 0.3$, $x(0) = 0$.

Figure 7.8: Results of the case for the stochastic optimal control formulation: $\mathbf{1}_K + \mathbf{1}_{A \setminus K}^\gamma \hat{\mathbf{W}}_0^*$.

⁵Note that the computation of the backwards mappings has been amended in comparison to the original case study in [Haesaert et al., 2016].

The optimal policy can now be computed as in Proposition 7.3 by selecting at each state $x_\theta = (x, \mu, R) \in A \setminus K$ the action that maximises the expected value function (computed empirically). Notice in Figure 7.8b and 7.8a that when starting the experiment at $x = 0$ and with μ initialised at the edges of the feasible set, i.e. $\mu = 1$ and $\mu = -1$, it can be observed that the function is *locally* convex hence the optimal control action (maximising the expected value function at the next step) is either equal to 1 or to -1 . These are the inputs which yield a maximum decrease in variance R and hence also induce a stochastic transition of μ with large variance. The latter is preferred due to the local convexity. More precisely, not applying an input (i.e. $u(0) = 0$) would mean that μ and R would not change over time. But any input different than zero will excite the system, therefore decrease the variance and create a stochastic transition kernel with a non-singular variance for μ . Hence with probability 1 the value at the next time step will improve. The latter follows from convexity. The fact that a decrease in variance R is optimal, even on the edges of the feasible set, shows that the problem statement is well-posed in that doing an experiment is preferred. Remember that this was not the case in the experiment design proposed by Mesbah et al. [2012b].

7.5.d Quasi-stationary reduced order problem

To approximate the optimal reach-avoid probability and to solve the corresponding experiment design problem, it is necessary to compute and represent the value functions

$$\hat{\mathbf{W}}_k^* : \mathbb{X}_\theta \rightarrow [0, 1]$$

and the corresponding policies

$$\pi_k : \mathbb{X}_\theta \rightarrow \mathbb{U}_{ex}.$$

The state space \mathbb{X}_θ of the experiment MDP (7.26) has an order or dimension equal to $2n + n(n+1)/2$, that is $\mathbb{X}_\theta \subset \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^{(n(n+1)/2)}$. The computation of value functions and policies taking values over a continuous domain, suffers under the curse of dimensionality [Gosavi, 2009]. Recent advances in (approximate) dynamic programming have been able to partially mitigate this explosion of computational complexity caused by increasing dimension that hinder the computation of solutions to large order MDPs. Still, driven by the quadratic increase in dimensionality of the XMDP for a linear increase in dimension of the parametrisation θ , we will, instead of focussing on more advanced techniques to solve the full size problem, *reduce the size or dimensionality of the problem*. More precisely, we will split the high dimensional problem into two lower dimensional optimisation problems.

In the first subproblem, we solve the experiment design over a partial model of the experiment. For this, we develop a reduced order model of the experiment MDP that only models the updates of μ and R . The manner in which the state updates drive the updates of μ and R is now modelled as an input. This is depicted in Figure 7.9 and will be detailed later on. For this reduced MDP, we can again solve a stochastic reach-avoid problem to optimise the outcomes of the ex-

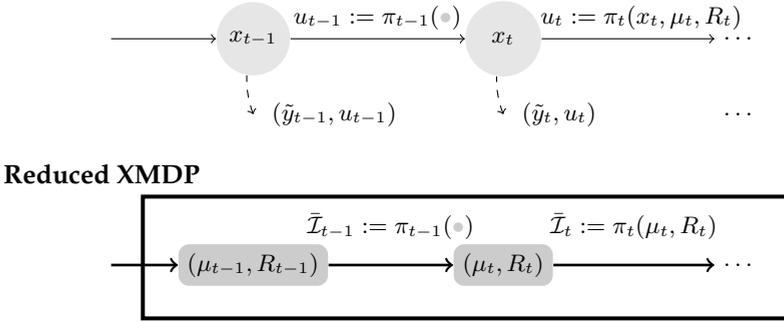


Figure 7.9: Modelling the experiment as an Markov decision process (Reduced XMDP) with a reduce state defined by the mean and covariance matrix.

periment. Same as before, we optimise the probability that a decision is reached within the experiment time.

The result of this reduction will be that to compute the optimal reach-avoid probability, it will only be necessary to compute value functions over a reduced state space, i.e.,

$$\hat{\mathbf{W}}_k^* : \mathbb{R}^{2n} \rightarrow [0, 1].$$

The input to this MDP is no longer $u(t)$, instead the input to the reduced experiment MDP will be shown to be an element of a higher dimensional space (that is of dimension n), and we chose it so that it can represent (approximately) all possible state transitions of the posterior distribution.

The solution of the reduced order experiment design problem requires a *secondary optimisation design step*, that is the design of an input $u(t)$ which generates or tracks the optimal input to the reduced experiment MDP. As such the optimal control problem is solved in a hierarchical fashion.

Reduced XMDP. We reduce the complexity of the experiment design problem by modelling the variance updates $R(t+1)$ of the experiment as a function of a new input. More precisely, we take the average per sample Fisher information matrix, as introduced in Section 7.4.b.2 equation (7.25), as the new input. For this we rewrite the experiment Markov decision process Σ_θ (7.26) as

$$\begin{aligned} x(t+1) &= Ax(t) + Bu(t), \\ \mu(t+1) &= \mu(t) + w(t), \\ \mathcal{I}(t+1) &= \mathcal{I}(t) + \sigma_e^{-2}x(t+1)x(t+1)^T, \end{aligned} \tag{7.27}$$

with $w(t) \sim \mathcal{N}(0, \mathcal{I}(t)^{-1} - \mathcal{I}(t+1)^{-1})$.

If the signal $u(t)$, that is the solution of the experiment design problem, can be modelled as a quasi-stationary signal, then $x(t+1)x(t+1)^T$ acts as an input to the transitions of μ and \mathcal{I} . The expected value of $x(t+1)x(t+1)^T$ can be modelled by the average per sample information matrix (7.25). Thus for the MDP model

given in (7.27), we use the average per sample information matrix as a new input affecting $\mathcal{I}(t+1)$ and $\mu(t+1)$ directly, and thereby neglect the real input signal $u(t)$ and the state updates $x(t)$. Thus the updates $x(t+1)$ are not modelled, instead the space to which $\sigma_e^{-2}x(t+1)x(t+1)^T$ belongs is modelled. As such the updates of $\mathcal{I}(t+1)$ can be governed by inputs in

$$\mathcal{I}(t+1) = \mathcal{I}(t) + \bar{\mathcal{I}}(t). \quad (7.28)$$

Note that a further reduction of the state dimension follows since \mathcal{I} takes only values in an affine subspace spanned by n elements, see Section 7.4.b.2. More precisely $\bar{\mathcal{I}}(t)$ (7.25) takes values in an n -dimensional affine subspace [Larsson, 2014, Jansson, 2004], thus we can again parameterise it as

$$\bar{\mathcal{I}} = \sum_{k=1}^n M_k \bar{i}_k,$$

and consequently, we can also parameterise the subspace to which \mathcal{I} belongs with i where $i(t) := \sum_{l=1}^t \bar{i}(l)$. That is we introduce a mapping to the set of information matrices

$$\mathcal{I}(t) = f_{\mathcal{I}}(i(t)) := \mathcal{I}(0) + \sum_{k=1}^n M_k i_k(t).$$

Let us first point out that we obtained a reduced order model which is of significant lower order, that is, the state is now of dimension $n+n$ instead of $n+n+\frac{1}{2}(n+1) \times n$.

Inputs and policy. The inputs are now defined by $\bar{i} = [\bar{i}_1 \ \dots \ \bar{i}_n]^T$ and the mapping to \mathbb{S}^n as $f_{\bar{\mathcal{I}}} : \bar{i} \mapsto \sum_{k=1}^n M_k \bar{i}_k$. Thus $\bar{\mathcal{I}}(t) = f_{\bar{\mathcal{I}}}(\bar{i}(t))$. Not all values of \bar{i} map to matrices $\bar{\mathcal{I}}$ that can be generated. Still the set of values of \bar{i} for which there exists a input sequence $u(t)$ in the set of allowed experiments can often be represented in an efficient way. This obtained allowed set for \bar{i} over time is denoted as $\mathcal{E}_{\bar{\mathcal{I}}}$ and we assume that there is a set $\mathbb{U}_{\bar{i}}$ such that if $\bar{i} \in \mathbb{U}_{\bar{i}}$ then there exists at least one $u(t) \in \mathcal{E}$ generating it. More precisely, input bounds \mathcal{E} such as power constraints can often be represented by convex constraints on \bar{i} [Larsson, 2014, Jansson, 2004]. Still some of the more complicated bounds on the set of allowed experiments could also be incorporated in the reach-avoid problem defining a safe set representing \mathcal{E} .

For the reduced experiment Markov decision process (rXMDP), a Markov policy π over the horizon $[0, N_s - 1]$ is a sequence $\pi = (\pi_0, \pi_1, \dots, \pi_{N_s-1})$ of measurable maps,

$$\pi_k : \mathbb{X}_\theta^r \rightarrow \mathbb{U}_{\bar{i}}, \quad k = 0, 1, \dots, N_s - 1,$$

from the reduced experiment state space \mathbb{X}_θ^r , with elements (μ, i) , to the $\mathbb{U}_{\bar{i}}$, which can be again mapped with $f_{\bar{\mathcal{I}}}$ to the corresponding average per sample Fisher information matrix $\bar{\mathcal{I}}$. We defined the set of policies for the reduced XMDP as Π_r .

Reduced experiment design problem For a given feasible set Θ_ψ , and a required confidence $1 - \delta$, we again denote with K the set of states (μ, i) associated with the required confidence of the verification problem, that is

$$K := \{(\mu, i) \mid \mathbb{P}(\theta^0 \in \Theta_\psi \mid \theta_0 \sim \mathcal{N}(\mu, (f_{\mathcal{I}}(i))^{-1})) \leq \delta \\ \text{or } \mathbb{P}(\theta^0 \in \Theta_\psi \mid \theta_0 \sim \mathcal{N}(\mu, (f_{\mathcal{I}}(i))^{-1})) \geq 1 - \delta\}.$$

With A we define a subset of safe states, e.g., $A = \mathbb{X}_\theta^r$ or some subset enforcing for example persistence of excitation on i . We say that an experiment is successful if K is reached while staying in A . The value of the experiment is discounted with the time it takes to reach the target set K . As such we again define the objective of the experiment as the discounted probabilistic reach-avoid property, $s_{x_\theta^r(0)}^*(K, A) = \sup_{\pi \in \Pi_r} s_{x_\theta^r(0)}^\pi(K, A)$, which is now defined over the reduced XMDF and whose computation for a finite horizon N_s follows Proposition 7.4. We say that an experiment design (with a given policy $\pi \in \Pi_r$) is optimal if it maximises the (discounted) probability that by the end of the experiment the target set is reached. The corresponding optimal policy, defined as

$$\pi^* = \arg \sup_{\pi \in \Pi_r} s_{x_\theta^r(0)}^\pi(K, A),$$

defines a sequence of policies mapping from the (μ, i) to the average per sample information matrix. Again a time independent policy is found when the discounted reach-avoid problem is extended to the infinite horizon case.

Generating the optimal experiment input sequence $u(t)$. The optimal policy π^* (or its approximation) defines the average per sample information matrix. For the actually implementation of the experiment, we need to design at every time instant $u(t)$ an input, generating $\{u(t)\}_{t \leq N_s}$ with the constraint that the full experiment should belong to \mathcal{E} , i.e., $\{u(t)\}_{t \leq N_s} \in \mathcal{E}$. This design of $u(t)$ is approached in a secondary design phase, and even though the actual design or implementation of this step is beyond the scope of this work, we will still point out two potential shortcomings.

The first issue is the the domain of the policy π . At each time instant π_t defines an input \bar{i} as function of both the current mean $\mu(t)$ and $i(t)$, now π_t is defined on the subset that has a one-to-one mapping with the space spanned by the Fisher information matrix \mathcal{I} for which there exists a stationary power spectrum. Ideally, at time t the covariance matrix $R(t)$ has an inverse that belongs to this space. But the set spanned by the possible Fisher information matrices \mathcal{I} is computed based on stationary assumptions, hence transient effects will generally push $R(t)$ out of this set. As a consequence, it will be necessary to extend the domain of the policy to include these transient matrices too.

The second issue is the design of $u(t)$ based on a given $\bar{\mathcal{I}}(t)$. We design the experiment to be as short as possible. Therefore, the stationarity assumption, which only holds with guarantees over longer time periods is an approximation step. The accuracy loss introduced by this assumption depends on the implementation

of the input signal, that is, how fast and accurate the input generates the required information matrix; and on how fast the input policy changes over time or with \mathcal{I} . Within the identification for control, the generation of an input sequence for a changing $\bar{\mathcal{I}}(t)$ is often also a secondary step for the experiment design problem Larsson [2014]. For them the optimal average per sample Fisher information matrix changes with the nominal estimate of the true parameter. Larsson [2014] introduced the design of an input sequence in an MPC fashion to attenuate for the stationary requirement and track the changing $\bar{\mathcal{I}}(t)$.

In the sequel, we will consider the designed policy π and analyse the expected behaviour of the optimal $\bar{\mathcal{I}}(t)$ during an experiment. Still, as the design of $\{u(t)\}$ and implementation of the experiment in this way is out of the scope of this work, more research will be needed to go from the solution of the reduced online optimal experiment problem to an actual experiment input.

Discussion, comparison and case study Instead of focussing on the implementation of the reduced order solutions into a real experiment set-up, we are interested in the qualitative and quantitative analysis of solutions to the Bayesian experiment problems formulated in this chapter. To enable a comparison we compute solutions to the online problem for the second order FIR model introduced first in Ex. 7.2. As such we can get insight into the properties of the online experiment input, the offline experiment input and D-optimal experiment input obtained by solving the respective experiment design problems.

Example 7.5 (Finite impulse response model: Ex. 7.2 Cont'd) For the model (7.22) we parameterise $\bar{\mathcal{I}}$ with

$$f_{\bar{\mathcal{I}}}(\bar{i}) = \sum_{k=1}^2 M_k \bar{i}_k$$

with

$$M_1 = \frac{1}{\sigma_e^2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \text{ and } M_2 = \frac{1}{\sigma_e^2} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

The inputs $\bar{i}(t)$ represent the autocorrelation of the input signal $u(t)$, c.f. Ex. 7.2. The corresponding input spectrum of $u(t)$ exists for $\bar{i}_1 \in \mathbb{R}^+$ with $|\bar{i}_2| < \bar{i}_1$. As before we only allow for experiments with bounded power, this gives us a bound on the variance of the input signal (auto-correlation without lag), thus $|\bar{i}_1| \leq 1$.

We now formulate the experiment design problem with as objective the verification over the diamond shaped feasible set given in Ex. 7.2. The verification problem, decided upon with a confidence of at least 0.95 (i.e., $\delta = 0.05$), should be verified within the experiment time. Remember that the end of the experiment is not fixed a-priori, but modelled as a probabilistic event to which we now choose to assign a probability of 0.92 (i.e., $\gamma = 0.92$). This in practice discounts the length of longer experiments.

We will first give a way to approximate the optimal value functions $\hat{\mathbf{W}}_k^*$ using approximate dynamic programming. Then we compare the results over the reduced XMDP.

Approximate dynamic programming. We start by approximating the optimal value function \mathbf{W}_k^* and policy π over a time horizon of 20. We limit the set of allowed policies

to those mapping to $(\bar{i}_1, \bar{i}_2) \in \{1\} \times \{-1, 0, 1\}$; that is the collection Π_r with elements $\pi_t^r : \mathbb{R}^2 \times \mathbb{R}^+ \times \mathbb{R} \rightarrow \{1\} \times \{-1, 0, 1\}$. For a given initial variance $R(0)$, or equivalently an initial $\mathcal{I}(0)$, we can quantify the set to which $\mathcal{I}(t)$ belongs for $0 \leq t \leq 20$ as $\{f_{\mathcal{I}}(i) \mid |i_2| \leq i_1, i_1 = t, i_2 \in \mathbb{Z}\}$. Thus the set of policies induces an automatic gridding of the last two dimensions of the state space. In other words, traces $\{x_\theta^r\}_{0 \leq t \leq N_s}$ of length N_s take values in $\prod_{k=0}^{N_s} \mathbb{X}_\theta^r(t)$ with

$$\mathbb{X}_\theta^r(t) := \{\mathbb{X}_\mu \times \{t\} \times \{-t, -t+1, \dots, t-1, t\}\}$$

Note that \mathbb{X}_μ is a bounded subset of the space in which μ evolves, i.e., \mathbb{R}^2 . The value of $\hat{\mathbf{W}}_k^*$ has to be computed over $\mathbb{X}_\theta^r(t)$. By gridding \mathbb{X}_μ , we can approximate following mappings for $k = N_s - 1, N_s - 2, \dots$

$$\hat{\mathbf{W}}_k^*(x_\theta^r) = \max_{\bar{i}_2 \in \{-1, 0, 1\}} \hat{\mathbb{E}} \left[\mathbf{1}_K(x_\theta^{r+}) + \mathbf{1}_{A \setminus K}^\gamma(x_\theta^{r+}) \hat{\mathbf{W}}_{k+1}^*(x_\theta^{r+}) \right] \quad (7.29)$$

for all $x_\theta = (\mu, k, i_2) \in \mathbb{X}_\theta^r(k)$ and with $\hat{\mathbb{E}}$ defined as the empirical mean over transitions

$$\begin{aligned} i_1^+ &= i_1 + 1 = k + 1 \text{ and } i_2^+ = i_2 + \bar{i}_2, \\ \mu^+ &= \mu + w \text{ with } w \sim \mathcal{N}(0, f_{\mathcal{I}}(k, i_2)^{-1} - f_{\mathcal{I}}(k+1, i_2^+)^{-1}), \\ x_\theta^{r+} &= (\mu^+, k+1, i_2^+). \end{aligned}$$

The first two dimensions of the domain of $\hat{\mathbf{W}}_k^*$ are only computed over a gridding of the bounded set defined by \mathbb{X}_μ . As long as $A \setminus K \subset \mathbb{X}_\mu$, it holds that the domain of $\hat{\mathbf{W}}_k^*$ can be trivially extended beyond \mathbb{X}_μ . In this case the expected values in the backwards mappings of Proposition 7.4 are estimated using 3000 realisations of the next state at every grid point in $\mathbb{X}_\theta^r(k)$ and for $\bar{i}_2 \in \{-1, 0, 1\}$. For this case study we additionally picked $\mathcal{I}(0) = \frac{1}{.15} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, and we gridded $\mathbb{X}_\mu := [-1.5, 1.5]^2$ with 50 equidistant cells in each dimension.

The approximate value functions and policies are depicted in Figure 7.10 and Figures 7.11-7.12. Figure 7.10 gives the value function and the policy at the start of the experiment. In the left figure of 7.10, we see that the value function is the lowest around the edges of the feasible set. The highest values can be observed outside of the feasible set, close to the target set K .

If we now look at the actions preferred for the different initial means, we notice that the selection of actions is very similar to those selected for the offline design in Ex. 7.2. Taking a closer look, we know that the roughly speaking the action $(1, 1)$ will induce a stochastic transitions of μ^+ on the $(1, 1)$ line from the origin. Similarly, $(1, -1)$ induces a transition on the $(1, -1)$ line. In contrast, where the stochastic transitions generated by $(1, 0)$ not prefer a single dimensions.

Figure 7.11 includes the approximate value function after 1, 6, 11 and 16 backwards mappings. Thus it gives $\hat{\mathbf{W}}_k^*$ for $k = 19, 14, 9, 4$ over a part of the relevant state space. It can be observed that the value functions increases with the backwards mappings. This validates the expectations. Additionally, Figure 7.12 gives the approximate policies associated to these value functions. For each k , we give for several values of i the value of $\hat{\mathbf{W}}_k^*$

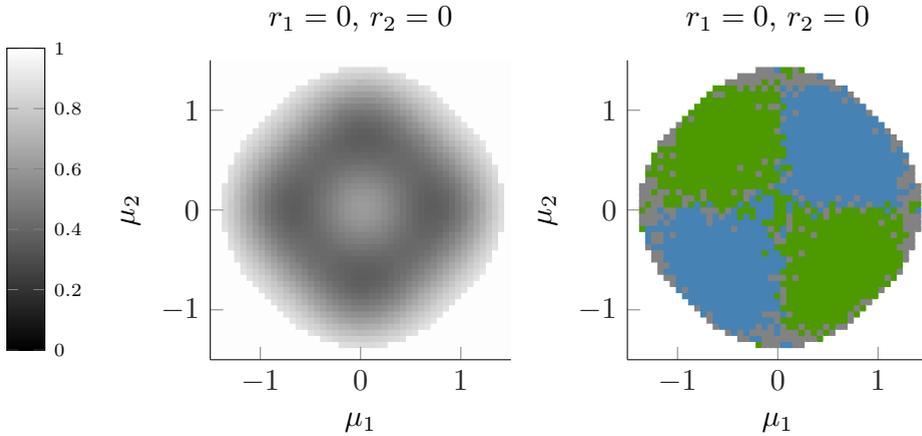


Figure 7.10: The figures represent the computed value function and the policy. The left figure depicts the value of $\hat{\mathbf{W}}_k^*(x_\theta^r)$ for $x_\theta^r(0) = (\mu_1, \mu_2, 0, 0)$. Thus for each mean μ in $A \setminus K$ it gives the discounted probability that within 20 time steps the required confidence threshold of 0.95 is reached. The white area (with value 1) is where for the initial $\mathcal{I}(0)$ and mean is such that $x_\theta^r \in K$. The right picture depicts for $x_\theta^r(0) = (\mu_1, \mu_2, 0, 0)$ the computed policy $\pi_0 \in \Pi_r$. The states are coloured with blue (for $(1, 1)$), green $((1, -1))$, and grey $((1, 0))$ as a function of the action (\bar{i}_1, \bar{i}_2) selected by the policy. When multiple actions gave the same value the (grey) input action $(1, 0)$ has been chosen.

over the different mean values μ . Since for a given k , $i = (i_1, i_2)$ and $i = (i_1, -i_2)$ give mirrored value functions we have only given the evaluations for positive i_2 .

We expect that when the mean is close to an edge of the feasible set, that it gives a lower value (reach-avoid probability) and that the optimal update of the covariance (or equivalently, the action) is orthogonal to that edge. This can be observed in the policies when i_2 is close to zero. Instead for large (absolute) values of i_2 the actions that decrease $|i_2|$ will be preferred. Remember that for $i_2 = 0$, we get a Fisher information matrix, hence also a covariance matrix, that is a scaled identity matrix.

Comparison. We are now able to compare the **D**-optimal experiment design, introduced in Section 7.3; the offline experiment design introduced in Section 7.4; and the online solution introduced in this section. For this we continue the case study 7.5 together with the solution method for the offline problem given in the preceding case study 7.2. For the initial

$$\mathcal{I}(0) = \frac{1}{0.15} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

and for all discretised states in $\mathbb{X}_\mu \cap (A \setminus K)$, we simulate an experiment based on the stochastic updates of \mathcal{I} and μ . For each initial point this is repeated 50 times. The resulting percentages of runs ending in a decision, averaged over all cells in $A \setminus K$, are presented in Table 7.1a.

We observe that only 70 percent of the experiment ended in a decision for the **D**-optimal design. We note that in this case the offline design is doing well in the beginning but

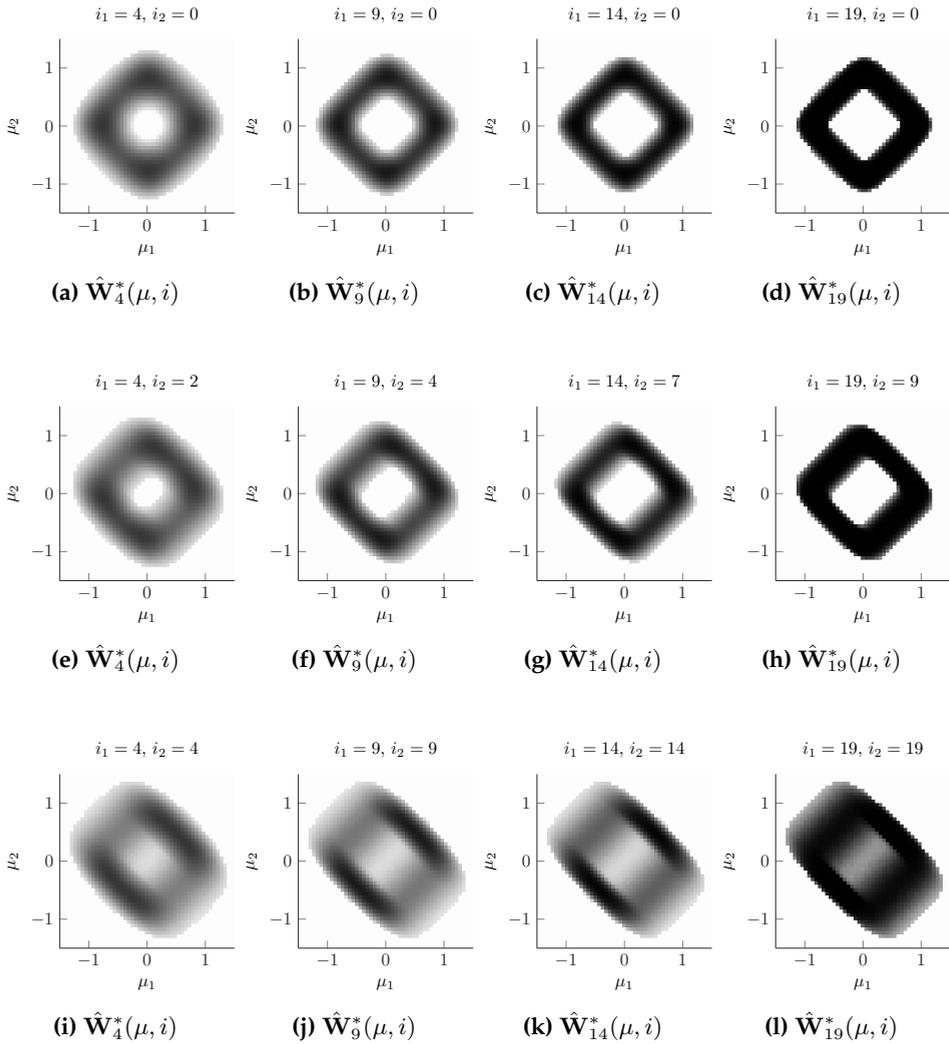


Figure 7.11: The figure represents parts of the value functions that have been computed while doing the backwards mappings (7.29). Each figure gives the value of the value function \hat{W}_k^* with $k = 19, 14, 9, 4$ for a given (i_1, i_2) over the set of possible mean values $\mu = (\mu_1, \mu_2)$. At each point $((\mu_1, \mu_2), (i_1, i_2))$ the plot depicts again the discounted probability (with white for 0 and black for 1) of reaching the required confidence level within the remaining time horizon in case the posterior distribution at that time is given as $\mathcal{N}((\mu_1, \mu_2), f_{\mathcal{I}}((i_1, i_2)))$.

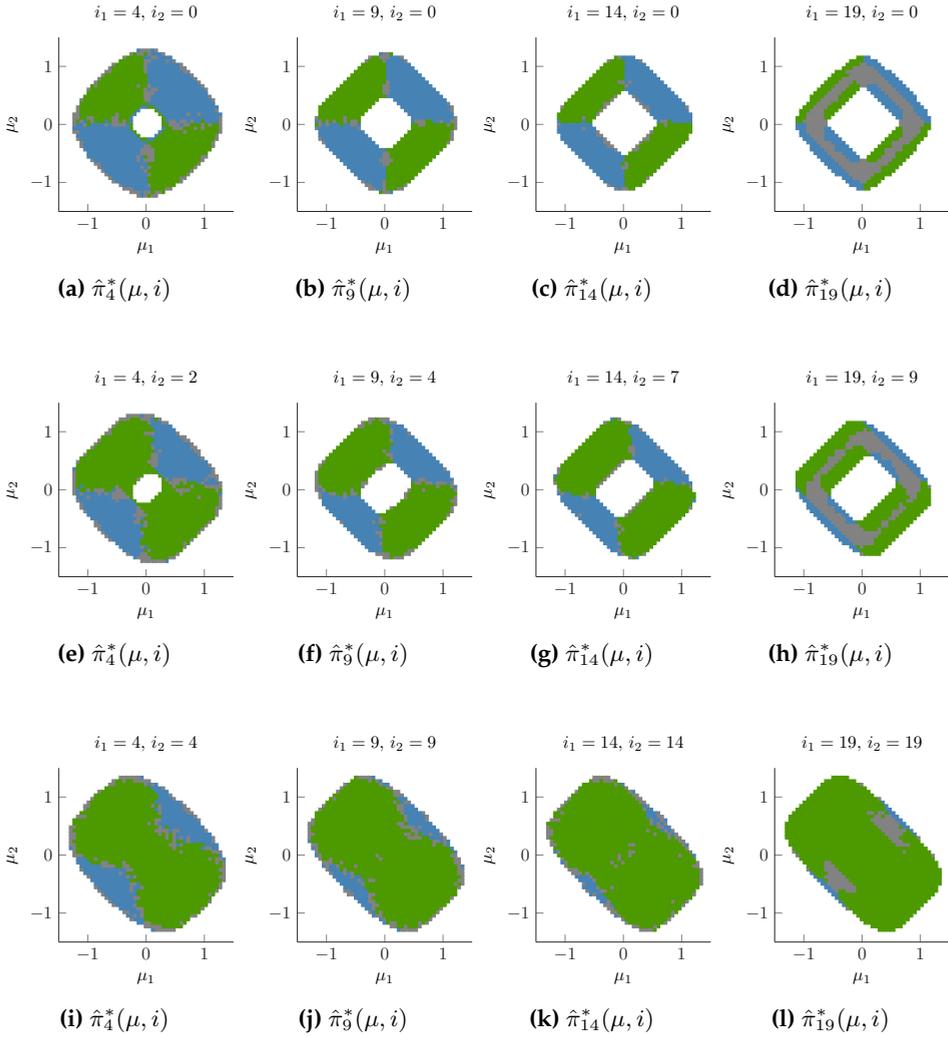


Figure 7.12: The figure represents parts of the policies that have been computed as arguments of the maximisation in the backwards mappings (7.29). Similar to Figure 7.11 gives the value of the value function \hat{W}_k^* with $k = 19, 14, 9, 4$ for a given (i_1, i_2) over the set of possible mean values $\mu = (\mu_1, \mu_2)$. The states are coloured with blue (for $(1, 1)$), green $((1, -1))$, and grey $((1, 0))$ as a function of the action (\bar{i}_1, \bar{i}_2) selected by the policy. When multiple actions gave the same value the (grey) input action $(1, 0)$ has been chosen.

gives only a couple of percentages benefits by the end of the experiment in comparison to \mathbf{D} -optimal. In comparison, the online solution gives around 5 percent of additional improvement upon the offline design.

To see the influence of the prior to these results, we now change the amount of knowledge initially in the prior, that is, we redo the computations with initial Fisher information matrix

$$\mathcal{I}(0) = \frac{1}{0.05} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

The results of the three experiment design are again summarised in Table 7.1b. It can be observed that this smaller initial variance, or larger initial \mathcal{I} means that the impact of additional information obtained by the experiment is lower. Hence within the same time limit, less experiments that are initialised in $A \setminus K$ end in a decisions. Still the closed loop design has a bigger impact in this lower variance case than in the higher variance case.

Time	5	10	15	20	Time	5	10	15	20
\mathbf{D} -Optimal (%)	43.4	58.4	66.1	71.2	\mathbf{D} -Optimal (%)	28.6	41.8	50.2	56.2
Offline (%)	46.2	60.0	68.0	72.7	Offline (%)	33.3	47.3	55.48	61.1
Online (%)	50.2	64.8	72.4	77.1	Online (%)	36.1	50.2	58.6	64.3

(a) Default initial variance.

(b) Small initial variance

Table 7.1: Comparison between two experiment design problems. Both tables give the percentage of runs, initialised in $A \setminus K$ that have reached K at several time instances of the maximal experiment length ($N_s = 20$)

In conclusion, we have observed that with respect to reaching a desired confidence level, the offline design already gives an improvement over the \mathbf{D} -optimal design, but the online design yields the best results. Remark that we have only considered the time-dependent policy and value functions for the reduced experiment MDP. It would be of interest to look at longer time horizons and fixed point solutions, to make a full comparison.

7.6 Conclusions

We have considered the design of experiment inputs for data-driven verification. For this we have developed two problem formulations, both of which have been stated in a Bayesian framework. In the first, we have given an utility-based Bayesian experiment design where the expected confidence has been maximised with an offline selected input sequence. In the second, we have considered the optimisation with respect to the probabilistic reachability of a given confidence threshold. This has been tackled by designing either time dependent or time independent policies for the online selection of the input action. We have reformulated this problem as a stochastic optimal control problem, solvable via dynamic programming.

For both problems, we have analysed the use of a stationary framework for the input design. The main focus of the presented work has been the understanding of the type of experiments that are optimal for verification related problems.

The use of an online experiment design has been shown to be beneficial. This should be verified on larger and more realistic case studies. Further it is of interest to consider advanced machine learning techniques such as those used in [Huan and Marzouk, 2016] to solve the experiment design problem. Of special interest is solving the recursions using machine learning.

List of Symbols

General

$\mathbf{1}_A$	indicator function for set A : $\mathbf{1}_A(x) = 1$ if $x \in A$ and 0 otherwise.
\mathbb{E}	Expected value operator
\bar{A}	For set A over domain \mathbb{A} , $\bar{A} := \mathbb{A} \setminus A$

Systems, Models, and Identification

\mathbf{S}	The true system
\mathbf{M}	A mathematical model, generally a state space model
\mathcal{G}	Set of mathematical models to which the true mathematical model belongs
$\mathbf{M}(\cdot)$	The mapping from the parameterisation to the set of mathematical models, $\mathbf{M}(\cdot) : \Theta \rightarrow \mathcal{G}$
$\mathbf{M}(\theta^0)$	True mathematical model representing \mathbf{S} in \mathcal{G}
$y_0(t)$	Actual output of the system
$\tilde{y}(t)$	Measured output of the system
$e(t)$	Measurement noise, presumed to be white noise
σ_e^2	Variance of $e(t)$, note $e(t) \sim \mathcal{N}(0, \sigma_e^2)$
$u(t)$	Input to the system
\mathbf{u}_{N_s}	A sequence of inputs of length N_s applied to the system \mathbf{S} during the experiment
$\tilde{\mathbf{y}}_{N_s}$	A sequence of outputs of length N_s measured during the identification experiment.

Experiment design: maximising expected confidence

μ	Mean of prior or posterior distribution
R	Variance of prior or posterior distribution
ut	Utility
\bar{ut}	Expected utility
J	Expected utility of optimal decision
$\mathbb{E}J$	Optimisation objective

Experiment design: maximising probability of reaching threshold confidence

Σ_θ	The experiment MDP (XMDP) modelling both the experiment and the state transitions.
\mathbb{X}_θ	The state space of the experiment MDP
A	The safe region, the complement of the to be avoided region in the reach-avoid problem.
K	The target set, for experiment design this is the set where the confidence is above a certain threshold.
$r_{x_\theta(0)}^\pi(K, A)$	The reach-avoid probability of Σ_θ initiated at $x_\theta(0)$ for the given policy π .
$s_{x_\theta(0)}^\pi(K, A)$	The discounted reach-avoid probability of Σ_θ initiated at $x_\theta(0)$ for the given policy π .
γ	Discount factor
$\mathbf{1}_B^\gamma$	γ -discounted indicator function $\gamma \mathbf{1}_B$
\mathbf{W}_k	The value function of the (discounted) reach-avoid problem. It evaluates the (discounted) probability of starting from $t = k$ reaching the target set while staying in the safe set.
\mathbf{W}_k^*	The optimal value function of the (discounted) reach-avoid problem. It evaluates the optimal (discounted) probability of starting from $t = k$ reaching the target set while staying in the safe set.
$\hat{\mathbf{W}}_k^*$	The approximate value function, that is a function used to approximate \mathbf{W}_k^* .
\mathbf{T}	Dynamic programming operator

Specifications

ψ	LTL property
Θ_ψ	Feasible set where ψ is satisfied
\mathcal{H}_0	Null-hypothesis
\mathcal{H}_1	Alternative hypothesis

7.A Derivation of Bayesian inference formulae

For a given θ , the measurements $\tilde{y}(t)$ $t = 1, 2, \dots$ are realised as

$$x(t+1) = A^t x(0) + \sum_{k=0}^t A^k B u(t-k) \quad x(0) = [0 \quad 0 \quad \dots \quad 0]^T$$

$$y_0(t) = \theta_0^T x(t),$$

$$\tilde{y}(t) = y_0(t) + e(t), \quad e(t) \sim \mathcal{N}(0, \sigma_e^2).$$

Remember that $\theta^T A x(0) + \theta^T B u(0)$ is equal to $\hat{y}(1, \theta)$ (implicitly parameterised on x_0). Therefore the conditional probability is derived as the probability density

function of the Gaussian, zero-mean, additive measurement noise $e(t) = \hat{y}(t, \theta) - \tilde{y}(t)$ with variance σ_e^2 . Denote the vector of N_s measurements as

$$\tilde{\mathbf{y}} = [y(1) \quad y(2) \quad \dots \quad y(N_s)]^T.$$

Since the measurement noise is white it follows that the probability density distribution of $\tilde{\mathbf{y}}$, given a sequence of inputs $\mathbf{u} = \{u(t)\}_{t \leq N_s-1}$, is obtained via the product of a Gaussian distributions

$$p(\tilde{\mathbf{y}}|\mathbf{u}, \theta) = \prod_{t=1}^{N_s} p(\tilde{y}(t)|\theta) = \frac{1}{\sqrt{\sigma_e^{2N_s} (2\pi)^{N_s}}} \exp \left[-\frac{\sum_{t=1}^{N_s} (\hat{y}(t, \theta) - \tilde{y}(t))^2}{2\sigma_e^2} \right].$$

The distribution over the data $\tilde{\mathbf{y}}$ is simplified using

$$\sum_{t=1}^{N_s} (\hat{y}(t, \theta) - \tilde{y}(t))^2 = (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})^T (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})$$

and the vector function $\Phi(\mathbf{u}) = [x(1) \quad x(2) \quad \dots \quad x(N_s)]$,

$$p(\tilde{\mathbf{y}}|\mathbf{u}, \theta) = \frac{1}{\sqrt{\sigma_e^{2N_s} (2\pi)^{N_s}}} \exp \left[-\frac{(\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})^T (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})}{2\sigma_e^2} \right].$$

Compute the posterior distribution

$p(\theta | \tilde{\mathbf{y}}, \mathbf{u}) = \frac{p(\tilde{\mathbf{y}}|\mathbf{u}, \theta)p(\theta)}{\int_{\Theta} p(\tilde{\mathbf{y}}|\mathbf{u}, \theta)p(\theta) d\theta}$ with the prior, $p(\theta) = \mathcal{N}(\theta_k, R_k)$. First we simplify the multiplication with the obtained data distribution

$$p(\tilde{\mathbf{y}}|\mathbf{u}, \theta)p(\theta) = \frac{1}{\sqrt{|R_k| \sigma_e^{2N_s} (2\pi)^{N_s+n}}} \exp \left[-\frac{1}{2} (\theta - \theta_k)^T R_k^{-1} (\theta - \theta_k) \right] \\ \times \exp \left[-\frac{(\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})^T (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})}{2\sigma_e^2} \right].$$

The matrix summation in the exponents can be written as

$$\sigma_e^{-2} (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})^T (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}}) + (\theta - \theta_k)^T R_k^{-1} (\theta - \theta_k) \\ = \begin{bmatrix} -1 \\ \theta \end{bmatrix}^T \begin{bmatrix} \tilde{\mathbf{y}}^T \sigma_e^{-2} \tilde{\mathbf{y}} + \theta_k^T R_k \theta_k & \sigma_e^{-2} \tilde{\mathbf{y}}^T \Phi^T(\mathbf{u}) + \theta_k^T R_k^{-1} \\ \sigma_e^{-2} \Phi(\mathbf{u}) \tilde{\mathbf{y}} + R_k^{-1} \theta_k & \Phi(\mathbf{u}) \sigma_e^{-2} \Phi^T(\mathbf{u}) + R_k^{-1} \end{bmatrix} \begin{bmatrix} -1 \\ \theta \end{bmatrix}.$$

Therefore the a-posteriori will have a Gaussian distribution with mean θ_{k+1} and variance R_{k+1}

$$R_{k+1} = [R_k^{-1} + \sigma_e^{-2} \Phi(\mathbf{u}) \Phi(\mathbf{u})^T]^{-1}, \\ \theta_{k+1} = R_{k+1} [R_k^{-1} \theta_k + \sigma_e^{-2} \Phi(\mathbf{u}) \tilde{\mathbf{y}}],$$

since we can rewrite the matrices as

$$(\theta - \theta_{k+1})^T R_{k+1}^{-1} (\theta - \theta_{k+1}) + \tilde{\mathbf{y}}^T \sigma_e^{-2} \tilde{\mathbf{y}} + \theta_k^T R_k \theta_k - \theta_{k+1}^T R_{k+1}^{-1} \theta_{k+1}.$$

In this summation the terms dependent on θ are separated from the terms independent of θ . The independent terms are related to the unlikelihood of the data based on the prior knowledge. When computing the a-posteriori distribution they are cancelled by the normalisation. Thus for a Gaussian prior, the a-posteriori probability distribution equals $p(\theta|\tilde{\mathbf{y}}, \mathbf{u}) = \mathcal{N}(\theta_{k+1}, R_{k+1})$.

$$p(\theta|\tilde{\mathbf{y}}, \mathbf{u}) = \frac{1}{\sqrt{|R_{k+1}|(2\pi)^n}} \exp \left[-\frac{1}{2} (\theta - \theta_{k+1})^T R_{k+1}^{-1} (\theta - \theta_{k+1}) \right],$$

with

$$R_{k+1} = [R_k^{-1} + \sigma_e^{-2} \Phi(\mathbf{u})\Phi(\mathbf{u})^T]^{-1},$$

$$\theta_{k+1} = R_{k+1} [R_k^{-1} \theta_k + \sigma_e^{-2} \Phi(\mathbf{u})\tilde{\mathbf{y}}].$$

Let us now compute the distribution of the data conditioned on \mathbf{u} only, and marginalised over θ

$$\begin{aligned} p(\tilde{\mathbf{y}}|\mathbf{u}) &= \int_{\Theta} p(\tilde{\mathbf{y}}|\mathbf{u}, \theta) p(\theta) d\theta \\ p(\tilde{\mathbf{y}}|\mathbf{u}) &= \int_{\Theta} p(\tilde{\mathbf{y}}|\mathbf{u}, \theta) \frac{1}{\sqrt{|R_k|(2\pi)^n}} \exp \left[-\frac{1}{2} (\theta - \theta_k)^T R_k^{-1} (\theta - \theta_k)^T \right] d\theta \\ &= \int_{\Theta} \frac{1}{\sqrt{\sigma_e^{2N_s} |R_k| (2\pi)^{N_s+n}}} \exp \left[-\frac{1}{2} (\theta - \theta_k)^T R_k^{-1} (\theta - \theta_k)^T \right] \\ &\quad \times \exp \left[-\frac{(\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})^T (\Phi^T(\mathbf{u})\theta - \tilde{\mathbf{y}})}{2\sigma_e^2} \right] d\theta \\ &= \int_{\Theta} \underbrace{\frac{1}{\sqrt{|R_{k+1}|(2\pi)^n}} \exp \left[-\frac{1}{2} (\theta - \theta_{k+1})^T R_{k+1}^{-1} (\theta - \theta_{k+1}) \right]}_{=1} d\theta \\ &\quad \times \sqrt{\frac{|R_{k+1}|}{\sigma_e^{2N_s} |R_k| (2\pi)^{N_s}}} \\ &\quad \times \exp \left[-\frac{1}{2} (\tilde{\mathbf{y}}^T \sigma_e^{-2} \tilde{\mathbf{y}} + \theta_k^T R_k \theta_k - \theta_{k+1}^T R_{k+1}^{-1} \theta_{k+1}) \right]. \end{aligned}$$

The term in the exponent can be written as a matrix multiplication by substituting $\theta_{k+1} = R_{k+1} [R_k^{-1}\theta_k + \sigma_e^{-2}\Phi(\mathbf{u})\tilde{\mathbf{y}}]$,

$$\begin{aligned} & \sigma_e^{-2}\tilde{\mathbf{y}}^T\tilde{\mathbf{y}} + \theta_k^T R_k^{-1}\theta_k \\ & - [R_k^{-1}\theta_k + \sigma_e^{-2}\Phi(\mathbf{u})\tilde{\mathbf{y}}]^T R_{k+1} [R_k^{-1}\theta_k + \sigma_e^{-2}\Phi(\mathbf{u})\tilde{\mathbf{y}}] \\ & = \begin{bmatrix} -1 \\ \tilde{\mathbf{y}} \end{bmatrix}^T \begin{bmatrix} \theta_k^T R_k^{-1}\theta_k - \theta_k^T R_k^{-1}R_{k+1}R_k^{-1}\theta_k & \theta_k^T R_k^{-1}R_{k+1}\sigma_e^{-2}\Phi(\mathbf{u}) \\ \sigma_e^{-2}\Phi(\mathbf{u})^T R_{k+1}R_k^{-1}\theta_k & \sigma_e^{-2}I - \sigma_e^{-4}\Phi(\mathbf{u})^T R_{k+1}\Phi(\mathbf{u}) \end{bmatrix} \begin{bmatrix} -1 \\ \tilde{\mathbf{y}} \end{bmatrix}. \end{aligned}$$

Consider the matrix inversion lemma and simplify the lower right corner of the block matrix

$$\begin{aligned} & \sigma_e^{-2}I - \sigma_e^{-4}\Phi(\mathbf{u})^T R_{k+1}\Phi(\mathbf{u}) \\ & = \sigma_e^{-2}I - \sigma_e^{-4}\Phi(\mathbf{u})^T [R_k^{-1} + \sigma_e^{-2}\Phi(\mathbf{u})\Phi(\mathbf{u})^T]^{-1} \Phi(\mathbf{u}) \\ & = \langle \text{use } (A - BD^{-1}C)^{-1} = A^{-1} + A^{-1}B(D - CA^{-1}B)^{-1}CA^{-1} \rangle \\ & = \langle \text{with } A = \sigma_e^2 I, B = \Phi(\mathbf{u})^T, C = \Phi(\mathbf{u}), D = -R_k^{-1} \rangle \\ & = [\sigma_e^2 I + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})]^{-1}. \end{aligned}$$

Denote the above matrix as $R_{\tilde{\mathbf{y}}}^{-1}$, i.e. $R_{\tilde{\mathbf{y}}} := [\sigma_e^2 I + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})]$. Using the matrix inversion lemma for the left upper term for $A = R_k^{-1}$, $B = \Phi(\mathbf{u})$, $C = \Phi(\mathbf{u})^T$, $D = -\sigma_e^2 I$ gives

$$\begin{aligned} & \theta_k^T R_k^{-1}\theta_k - \theta_k^T R_k^{-1}R_{k+1}R_k^{-1}\theta_k = \theta_k^T R_k^{-1}(R_k - R_{k+1})R_k^{-1}\theta_k \\ & = \theta_k^T R_k^{-1}(R_k - [R_k^{-1} + \sigma_e^{-2}\Phi(\mathbf{u})\Phi(\mathbf{u})^T]^{-1})R_k^{-1}\theta_k \\ & = \theta_k^T R_k^{-1}(R_k - R_k - R_k\Phi(\mathbf{u})(-\sigma_e^{-2}I - \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1}\Phi(\mathbf{u})^T R_k)R_k^{-1}\theta_k \\ & = \theta_k^T [\Phi(\mathbf{u})(\sigma_e^2 I + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1}\Phi(\mathbf{u})^T] \theta_k \\ & = \theta_k^T \Phi(\mathbf{u})R_{\tilde{\mathbf{y}}}^{-1}\Phi(\mathbf{u})^T \theta_k. \end{aligned}$$

Hence the matrices inside the exponent of the density function become

$$\begin{bmatrix} -1 \\ \tilde{\mathbf{y}} \end{bmatrix}^T \times \begin{bmatrix} \theta_k^T \Phi(\mathbf{u})R_{\tilde{\mathbf{y}}}^{-1}\Phi(\mathbf{u})^T \theta_k & \theta_k^T R_k^{-1}R_{k+1}\sigma_e^{-2}\Phi(\mathbf{u}) \\ \sigma_e^{-2}\Phi(\mathbf{u})^T R_{k+1}R_k^{-1}\theta_k & R_{\tilde{\mathbf{y}}} \end{bmatrix} \times \begin{bmatrix} -1 \\ \tilde{\mathbf{y}} \end{bmatrix}.$$

Now we want to extract the mean value $\mu_{\tilde{\mathbf{y}}}$ of the data sequence $\tilde{\mathbf{y}}$ from the off-diagonal terms

$$\begin{aligned}
R_{\tilde{\mathbf{y}}}^{-1} \mu_{\tilde{\mathbf{y}}} &= \sigma_e^{-2} \Phi(\mathbf{u})^T R_{k+1} R_k^{-1} \theta_k \\
\mu_{\tilde{\mathbf{y}}} &= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} \Phi(\mathbf{u})^T R_{k+1} R_k^{-1} \theta_k \\
&= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} \Phi(\mathbf{u})^T [R_k^{-1} + \sigma_e^{-2} \Phi(\mathbf{u}) \Phi(\mathbf{u})^T]^{-1} R_k^{-1} \theta_k \\
&= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} \Phi(\mathbf{u})^T [I + \sigma_e^{-2} R_k \Phi(\mathbf{u}) \Phi(\mathbf{u})^T]^{-1} \theta_k \\
&= [\text{Note that : } (I + AB)^{-1} = I - A(I + BA)^{-1} B] \\
&= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} \Phi(\mathbf{u})^T [I - \sigma_e^{-2} R_k \Phi(\mathbf{u}) (I + \sigma_e^{-2} \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1} \Phi(\mathbf{u})^T] \theta_k \\
&= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} [\Phi(\mathbf{u})^T \theta_k - \sigma_e^{-2} \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}) (I + \sigma_e^{-2} \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1} \Phi(\mathbf{u})^T \theta_k] \\
&= \sigma_e^{-2} R_{\tilde{\mathbf{y}}} [I - \sigma_e^{-2} \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}) (I + \sigma_e^{-2} \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1}] \Phi(\mathbf{u})^T \theta_k \\
&= [\text{Push through Identity } A(I + A)^{-1} = I - (I + A)^{-1}] \\
\mu_{\tilde{\mathbf{y}}} &= R_{\tilde{\mathbf{y}}}^{-1} \left[(\sigma_e^2 I_{N_s \times N_s} + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u}))^{-1} \right] \Phi(\mathbf{u})^T \theta_k = \Phi(\mathbf{u})^T \theta_k
\end{aligned}$$

Thus

$$\begin{aligned}
p(\tilde{\mathbf{y}}|\mathbf{u}) &= \sqrt{\frac{|R_{k+1}|}{\sigma_e^{2N_s} (2\pi)^{N_s} |R_k|}} \exp \left[-\frac{1}{2} (\tilde{\mathbf{y}} - \Phi(\mathbf{u})^T \theta_k)^T R_{\tilde{\mathbf{y}}}^{-1} (\tilde{\mathbf{y}} - \Phi(\mathbf{u})^T \theta_k) \right] \\
&= * \mathcal{N}(\Phi(\mathbf{u})^T \theta_k, [\sigma_e^2 + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})])
\end{aligned}$$

The last equality =* follows from

$$\begin{aligned}
\frac{\sigma_e^{2N_s} |R_k|}{|R_{k+1}|} &= |R_{\tilde{\mathbf{y}}}| \\
\sigma_e^{2N_s} |R_k| |R_k^{-1} + \sigma_e^{-2} \Phi(\mathbf{u}) \Phi(\mathbf{u})^T| &= |\sigma_e^2 + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})| \\
\sigma_e^{2N_s} |I + \sigma_e^{-2} \Phi(\mathbf{u}) \Phi(\mathbf{u})^T R_k| &= |\sigma_e^2 + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})| \\
&= \langle \text{Sylvester's determinant theorem} \rangle \\
|\sigma_e^2 + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})| &= |\sigma_e^2 + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})|
\end{aligned}$$

Gaussian prior \rightarrow *Gaussian data distribution*

$$p(\tilde{\mathbf{y}}|\mathbf{u}) = \sqrt{\frac{1}{(2\pi)^{N_s} |R_{\tilde{\mathbf{y}}}|}} \exp \left[-\frac{1}{2} (\tilde{\mathbf{y}} - \Phi(\mathbf{u})^T \theta_k)^T R_{\tilde{\mathbf{y}}}^{-1} (\tilde{\mathbf{y}} - \Phi(\mathbf{u})^T \theta_k) \right],$$

with $R_{\tilde{\mathbf{y}}} = [\sigma_e^2 I_{N_s \times N_s} + \Phi(\mathbf{u})^T R_k \Phi(\mathbf{u})]$.

The testament of science is so continually in a flux that the heresy of yesterday is the gospel of today and the fundamentalism of tomorrow.

E. Kasner and J. R. Newman

8

Conclusions and future work

8.1 Overview & summary

Within control engineering, the importance of formal methods for the verification and control of physical systems is expected to follow the growing demands on intelligent and automated systems. The bulk of formal methods has been developed for software or hardware verification and targets the verification of finite state machines.

Instead, physical systems evolve over continuous (or more general uncountable) spaces and their dynamic evolutions are subject to uncertainties. It is for these circumstances that there exist many open challenges for the development of formal methods. We would like to know how to synthesise control modules and how to formally prove system properties, when our knowledge on the system is limited due to partial knowledge of the system dynamics, inherent stochasticity in the evolutions of the system, and noisy output measurements. As we cannot yet deal with all the aspects of this problem, we have targeted several subproblems in this thesis.

Firstly, leaving the issue of unknown system dynamics to the side, we have considered the synthesis problem. We have questioned whether control synthesis with respect to formal properties for general models evolving stochastically could be made more computationally efficient without losing formality. This has resulted in a control refinement approach based on new approximate simulation relations. Secondly, for a simpler set of models we have considered the combined issue of having both stochastic transitions and noisy output measurements.

In part two of this thesis, we retook the issue of systems with partly unknown dynamics. Specifically, their verification with respect to formal specifications was

of interest. We have formulated this in a Bayesian framework in which formal specifications are verified over the system by quantifying the credibility of their satisfaction. This has also allowed for the integration of measurement data with knowledge on the system dynamics. Additionally, we have looked at how the choice of experiments, from which the data is obtained, influences our confidence in whether or not a physical system verifies a given property. That is, we have questioned whether we can faster be sure about system properties by designing tailored experiments.

In conclusion, the results in this thesis contribute to the verification and control of physical system. Within this chapter, we recap the results obtained and discuss unsolved issues. Additionally, based on the work performed in this thesis, we single out a couple of research lines that we perceive as key for the further development of verification and control of physical systems.

8.2 Results

8.2.a Control synthesis for general Markov decision processes

Results Chapter 3: Policy refinement using approximate similarity relations

For Markov decision processes evolving over continuous spaces, verifiable control synthesis relies on reductions of the models to (finite) state space abstractions. This facilitates algorithmic policy or control constructions with quantitative guarantees. Leveraging model reductions or abstractions that enable this quantitative refinement, one obtains controllers that are correct-by-design. For this, it is key to provide formal guarantees on the abstraction step. Within Chapter 3, we have introduced a new approximate probabilistic similarity relation, which allows a useful trade-off between deviation over probability distributions on state transitions and output distances. The relations, underpinned by the use of metrics on the output space of the models, allows us to go beyond the available results that quantify exclusively either deviations in probability or deviations in output.

Discussion of the results in Chapter 3

Applications. The current treatment of the topic is still very theoretical. In the near future, application of the theory will bring more insight into the weak spots and possible improvements. One of the major issues is the unpenalised use of inputs, within the refinement no priority is given to small inputs or inputs close to the input of the abstract model. For real applications the inputs or control actions define the power consumption of the engineering system, hence it is often natural to also have requirements on them. This is especially the case when the applications are modelled with large or unbounded input sets. For these systems an interface minimising the precision of the approximate similarity relations could yield unrealistically large inputs.

Probabilistic properties. Alongside practical applications of the developed notions, a further generalisation of Theorem 3.13 in Chapter 3 to specific quantitative properties expressed via temporal logics is necessary. Theorem 3.13 can directly be used for the quantification of properties such as safety and reachability, thereby it also hints at extensions of the theory to more complicated properties. Future work should focus on defining and extending the probabilistic properties for which the approximate similarity relations can be used. It is especially of interest to question how the approximate similarity relations can be used to synthesise correct-by-design control policies with respect to temporal properties such as the PCTL ones [Baier and Katoen, 2008].

Computational tools. We have used results on control refinement for deterministic LTI systems to construct interface functions and compute relations over the state spaces in the case studies on smart buildings. For this and other model classes within the class of gMDPs, the algorithmic and automated construction of appropriate interface functions together with the optimal quantification of the ε, δ -approximate similarity relation is a topic of further research.

Within the formal verification and controller synthesis for MDP processes evolving over continuous state spaces, abstractions should be built upon a combination of model-order reduction together with finite-space discretisation. The presented results have been developed specifically to allow for model-order reductions, which are often norm based. Still more research is necessary to make a better connection between the available model-order reductions (surveyed in inter alia Gugercin and Antoulas [2004]), their computational implementations, and the developed approximate similarity relations.

Additionally, for large scale systems results of compositionality of the relations and interfaces would be beneficial. With this we target the design of (approximate) similarity relations and interfaces for large models as compositions of the similarity relations and interface functions computed on their sub-models.

Continuous time. Often one desires to consider the verification of stochastic systems evolving in continuous time. In discrete time, we have quantified the probabilistic deviation between stochastic transition kernels by the probabilistic invariance of the relation over the combined state spaces. The quantity δ bounds at every time instant the probability of leaving the relation. For continuous-time models we would have to search for an equivalent notion.

Beyond MDPs. The introduced similarity relations induce equivalence relations and pre-orders over the set of MDPs. But, the definition relies on a coupling or lifting of the stochastic kernels of the processes. Some classical cases, where exact equivalence of probabilistic distributions are known, cannot be identified within the proposed framework. Let us give a simple example to demonstrate the potential restrictiveness. Consider two Gaussian LTI systems evolving over

$(x_1, x_2) \in \mathbb{X}_a = \mathbb{R}^2$ and $z \in \mathbb{X}_b = \mathbb{R}$, respectively, with transitions

$$\mathbf{M}_a : \begin{cases} x_1(t+1) = a(t) \\ x_2(t+1) = x_1(t) \\ y_a(t) = x_2(t) \end{cases} \quad \mathbf{M}_b : \begin{cases} z_1(t+1) = b(t) \\ y_b(t) = z_1(t) \end{cases}$$

where $a(t)$ and $b(t)$ are white noise signals with a standard Gaussian distribution, i.e., $a(t) \sim \mathcal{N}(0, I)$ and $b(t) \sim \mathcal{N}(0, I)$. A simulation relation \mathcal{R} has to satisfy $\mathcal{R} \subset \{(x_1, x_2), z_1) \in \mathbb{X}_a \times \mathbb{X}_b : x_2 = z_1\}$. We observe that for this $a(t-1)$ should be related to $b(t)$. Based on this fact, we see that the stochastic transition kernels cannot be lifted. Still, the output traces $y_a(t)$ and $y_b(t)$ define the same probability measure.

A similar problem appears when we consider the case of LTI systems with output noise. Consider the standard notation for such a model, \mathbf{M}_y

$$\begin{cases} x(t+1) = Ax(t) + Bu(t) + B_w w_1(t) \\ y(t) = Cx(t) + D_w w_2(t). \end{cases}$$

Note that this can be rewritten as \mathbf{M}_{y_2}

$$\begin{cases} x(t+1) = Ax(t) + Bu(t) + B_w w_1(t) \\ x_w(t+1) = w_2(t+1) = w_3(t) \\ y(t) = [C \quad D_w] \begin{bmatrix} x(t) \\ x_2(t) \end{bmatrix}. \end{cases}$$

\mathbf{M}_{y_2} belongs to the set of gMDPs. For a given control sequence, traces of \mathbf{M}_y and \mathbf{M}_{y_2} give equal distributions over the canonical output space, but their equivalence can again not be shown with our similarity relation.

Now it is also known that \mathbf{M}_y has an innovation form, given as \mathbf{I}_y

$$\begin{cases} x(t+1) = Ax(t) + Bu(t) + Ke(t) \\ y(t) = Cx(t) + e(t). \end{cases}$$

with $e(t) \sim \mathcal{N}(0, C^T Q C + D_w D_w^T)$ for some Q (computed based on the Riccati equations). This innovation form is known to generate the same output space (for equal input sequences) as \mathbf{M}_y , additionally its state represents the available information for controller synthesis in the output-feedback case. Hence finding a way to define a similarity relation between \mathbf{I}_y and \mathbf{M}_{y_2} , would allow insight into controller refinement for output based synthesis.

We claim that to work with general partially observable MDPs, it will be necessary to extend the relations over the state spaces to relations over probability distributions of the state spaces, perhaps similar to those used by Segala [1995] for the concept of forward lifting.

8.2.b Control synthesis for partially-observable LTI models

Results Chapter 4: Extension of correct-by-design controller synthesis

For linear time invariant models, we have extended the applicability of correct-by-design controllers to partially observable ones. We have shown that any standard correct-by-design state-based controller designed for a deterministic abstraction of the system can be extended to an output-based controller with quantified bounds. For this, we have developed a hierarchical design methodology that hinges on the design of a state-observer and interface, verified with matrix inequalities. The results are applicable to LTI models disturbed by Gaussian additive noise on the state transitions and on the output measurements.

Bounds on the accuracy loss within the developed hierarchical control refinement have been expressed via the expected 2-norm. This has as a benefit that the accuracy quantification can be defined based on a feasibility check over matrix inequalities. Further, these inequalities are shown to be similar to those for the H_2 -optimal control problem. Though similar, the major difference is the structure of the controller imposed by the hierarchical refinement. The original H_2 -controller design problem, subject to these inequalities, can be transformed to a problem subject to linear matrix inequalities which in turn can be solved with interior point algorithms. The structure imposed on the controller makes this standard transform to linear matrix inequalities infeasible.

Results Chapter 5: on the separation theorem

In this chapter, we have considered the control design problem introduced in Chapter 4 and investigated computational solutions to the design problem by separating the matrix inequalities.

Guarantees for H_2 -performance with separated matrix inequalities. Targeting the algorithmic construction of output-based hierarchical control refinement, new results on the matrix inequalities for the H_2 -optimal control problem have been developed. These results allow a separation of the matrix inequalities. The matrix inequality of the controller system is replaced by an inequality for the state-estimator design and an inequality for the state-feedback design. As a result of this we can extend the proof of the separation theorem based on matrix inequalities and show that the separated matrix inequalities define sufficient conditions for the H_2 -norm. This is shown to be also of interest for multi-objective controller design.

Guaranteed hierarchical control refinement and lexicographic refinement.

Secondly, we have extended the developed theory to the hierarchical control refinement problem (of which 4 is a specific case), for which more structure is imposed on the matrix inequalities. For this we show the implications for the output-based hierarchical refinement. Beyond the standard output-based refinement, we

also develop a lexicographic refinement design. This allows us to take into account a performance objective that penalises control actions next to the objective minimising the accuracy loss on required controller specifications. In this multi-objective setting, the lexicographic approach allows for an ordering of objectives. In this case, we have prioritised safety (or functionality) above performance. In other words, we can optimise for the accuracy loss while still allowing for a small deviation from the optimum to yield a better performance. This relies on the aforementioned results obtained on the separation of estimation and control.

Discussion of chapters 4 & 5

Generalisation of the results. Current results on hierarchical control refinement utilised in the correct-by-design results rely on exact model knowledge. That is knowledge of an LTI model, representing the underlying system dynamics, and knowledge of the disturbance is assumed. In reality, the true dynamics of these systems are often not known exactly. Additionally, they are never exactly linear. Still they are often well approximated via linear models, via hybrid (that is, piece-wise affine models) or via linear parameter varying models. Thus to obtain a correct-by-design synthesis procedure for physical systems of which we do not have exact model knowledge and/or of which the dynamics cannot be represented by an LTI model, it will be necessary to extend our results to deal with modelling errors. This will require a change in the design methodology and in the computation of the guarantees via matrix inequalities. This allows to subsequently extend of the controller synthesis to hybrid models, or linear parameter varying models.

To further deal with partly unknown models it is of interest to consider the use of data-driven models. The incorporation of data-driven modelling of the system within this output-based approach also relies on the quantification of the influence of modelling errors. The knowledge of the influence of modelling errors on the functioning and performance of the controller can be used to synthesise optimal model estimation algorithms and to design optimal experiments.

Beyond the expected 2-norm. Currently, the quantification of accuracy loss is expressed based on the expected 2-norm. Together with the Markov inequality this can be used to bound deviations in the satisfaction of probabilistic properties. It is of interest to consider a quantification of the accuracy loss that directly expresses the probabilistic error, such as the one developed in Chapter 3. But the aforementioned error (or accuracy) quantification does not cover the observer-based refinements, for this the theory in Chapter 3 needs to be extended to include innovation forms as mentioned in the discussion of Section 8.2.a.

Are guarantees on optimality possible in this framework? Only parts of the design of a hierarchical control refinement for this output feedback problem can be performed fully automatically and based on given functional requirements and performance objectives. An optimisation based design leveraging the results

in Chapter 5 is either optimal for the observer gain or for the state-feedback gain. Specifically, given a (sub-)optimal gain for the observer a state-feedback gain can be computed obtained by inter alia the presented lexicographic optimisation. Of course, towards the future, we should question whether we can solve both gains in a single optimisation algorithm.

8.2.c Verification of partially unknown systems

Results Chapter 6: data-driven and model-based verification

The strength of formal techniques, such as model checking, is bound to the fundamental requirement of having access to a given model, obtained from the knowledge of the behaviour of the underlying system of interest. In Chapter 6, we have developed a new framework to deal with the verification of physical systems, allowing both partly unknown dynamics over uncountable (continuous) variables and noisy output measurements. Within this Bayesian framework the confidence in a formal property is related to the uncertainty of a model inferred from data.

Discussion of Chapter 6

The contributions of this work to the state-of-the art in data-driven verification is two-fold. First we present a Bayesian framework in which model-based reasoning is combined with data-driven reasoning, with as goal the verification of partially unknown systems. Secondly, we show that under some conditions we can obtain analytical solutions of the otherwise involved computations. Within the work presented by Polgreen, Wijesuriya, Haesaert, and Abate [2016] the data-driven and model-based approach has been used for the verification of finite-state Markov processes. Instead, in this thesis, we have applied the developed principled method to the verification of physical systems, which can be modelled by a set of linearly parameterised LTI models, and we have derived solutions for a subset of the LTL properties.

Linear parameterised LTI models. We have worked with systems modelled using linearly parameterised model sets defined through orthonormal basis functions. Practically, this has been widely used for the modelling of physical systems, such as the thermal dynamics of buildings [Virk and Loveday, 1994].

The insight developed through the analytical expressions induced by choosing this simple set models is important. Firstly, it gives insight into the development of experiments tackled in the subsequent chapter. Additionally, the computations scale similarly to those developed for reachability problems. Therefore it allows for applications on relatively large scale models, thereby providing insight on the behaviour of the models in the model set and the confidence quantification in a way that can easily be understood and interpreted. Still the use of linearly parameterised model set presents only a first step towards the goal of data-driven

verification and more research should be done to non-linear parameterised models, hybrid or piece-wise affine models, and stochastic models.

Beyond the subset of LTL properties To move beyond the current limited set of LTL properties, we want to include negations in the syntax. Combined with the ‘and’ operator, the negation defines the or operator as $a \vee b \equiv \neg(\neg a \wedge \neg b)$. We can see that this ‘or’ operator can create disjunct sets in the parameter space, which is something we can currently not cover with the developed theory.

Still we would like to point out that the current set of properties allows additionally for the verification of more complex properties via implication-based reasonings. Consider for example the case where a property includes the ‘or’ operation and the satisfaction of a simple property can be used to imply its verification.

Results Chapter 7: Bayesian experiment design

Within this Bayesian framework of data-driven and model-based verification, we have questioned what makes an experiment optimal. We have shown that this experiment design problem distinguishes itself from the standard experiment design problems for estimation, in that its main goal is not the accurate estimation of a parameterised model. Instead, the overall objective of this work is to verify or falsify a property of the underlying system of interest. As such there is only an indirect need to minimise the accuracy (and especially the variance) of the estimate. We have formulated the problem in two ways. First we have considered the maximisation of the expected confidence. The resulting classical Bayesian experiment design problem has been solved in an offline fashion, where we have designed the input sequence of the experiment before the start of the experiment. Secondly, we have considered the maximisation of the likelihood that a confidence, supporting either accepting or rejecting the property, is achieved. For this, we have looked for a sequence of (time-independent) policies that defines the inputs online based on the measurements up to that time instance. In this case, we have shown that the experiment design problem can be written as a stochastic reach-avoid problem. This has allowed us to rewrite the problem as a dynamic programming problem. Furthermore, we have looked at reductions of both the offline and the online case based on stationarity assumptions.

Discussion: Experiment design

The presented results introduce a new online perspective on experiment design. Still further analysis of the problem and the solutions will be necessary to develop more understanding. Of special interest is the comparison with experiments designed with full information of the true model.

Since the open-loop formulation scales badly in the number of free parameters, computational approaches to solve the induced stochastic reach-avoid problem will be key. Let us consider next the restrictiveness of the assumptions and some of the other aforementioned issues.

Beyond Linear parameterisations & state information. Currently, linearity of both the system dynamics and the parameterisation is exploited to obtain analytical solutions to the Bayesian inference formulas. Both assumptions are not necessary when analytical solutions can be replaced by either numerical approximations or empirical solutions.

Computational issues. Within Chapter 7, we have already elaborated on how future research can attenuate the computational difficulties associated with the stochastic optimal control problem. Together with a reduced order descriptions of the experiment MDP, the usage of more advanced machine learning techniques should bring the practical application of Bayesian experiment design within reach.

Hypothesis testing & statistical implications. The proposed Bayesian experiment design formulation is not limited to the verification of the specific type of properties introduced in this thesis. It would therefore be of interest to look into a more general setting of hypothesis testing. More precisely, it would be of specific interest to develop results supporting the use of this method for sequential hypothesis testing as introduced by Wald [1945].

8.3 Future work

A lot of work lies ahead for the verifiable control synthesis over partially unknown physical systems. In the previous section, we already discussed our results and what is still missing. Toward the original goal of verification and control of physical systems, we draw three main lines of development that we consider important

1. Control synthesis for partially observable gMDPs;
2. Computational tools for control synthesis;
3. Identification for correct-by-design control synthesis.

Distilled from the results in this thesis and from the original goal of controller synthesis and verification of physical systems, we perceive these lines as natural and promising research directions that can build further on the presented results. In addition, more fundamental research within the identification community on the topic of

4. Experiment design for hypothesis testing

would be beneficial both for failure detection and for identification. Results on this topic will also aid in the verification of physical systems.

Let us elaborate on these lines of research next.

8.3.a Control synthesis for partially observable gMDPs

At this moment the automated controller synthesis for partially observability gMDPs is a very hard topic to tackle. For linear time invariant models, we have reasoned that a generalisation of the introduced notions of similarity relations (cf. Chapter 3 and Section 8.2.a) can yield results that are beneficial for the output-based controller synthesis with guarantees on the obtained accuracy. More precisely, of interest is first the generalisation of approximate similarity relations over uncountable spaces to the cases introduced in Section 8.2.a. This would, as a first step, allow for the quantification of both the probabilistic deviation and the deviations in the output trajectories for the control structure proposed in Chapter 4. Secondly, it would also allow for solving the synthesis problem on reduced models of partially observable gMDPs while retaining guarantees on the policy refinement.

8.3.b Computational tools for control synthesis

Next to the development of theory, it is important to consider the computational load and to develop software allowing for the continuous testing on real life case studies. For the development of software and hardware there exists packages that support the design phase with formal methods tools. Recently, non-commercial toolboxes have been developed for the verification and correct-by-design controller synthesis of simple continuous space models. Exemplar are the TuLiP toolbox¹ and the Pessoa toolbox² both contain methods to synthesise correct-by-design controllers for non-stochastic systems with full state measurements and with full knowledge of the system dynamics. Alternatively, the FAUST² toolbox [Esmail Zadeh Soudjani and Abate, 2013], which has been used in this thesis (together with Pessoa), can handle state-based control of Markov decision processes. Toolboxes like the above support the development and assessment of further research on the topic of verification and control synthesis as they go to the core of the objective of formal methods. That is, they provide automatised manners to synthesis and verify controller designs.

Therefore it is of interest to support the theoretical research contributions with the continued development and extension of these toolboxes towards higher scale systems, systems that include stochasticity on state transitions and output measurement, and of which the system dynamics are only partially known.

8.3.c Identification for correct-by-design control

We intend to pursue the following questions: can we design a controller, synthesised based on formal properties, for uncertain models inferred from data?

¹<http://tulip-control.sourceforge.net>

²<https://sites.google.com/a/cyphylab.ee.ucla.edu/pessoa/>

Moreover, can we do it in such a way that the controller is reactive to the environment and such that it can adapt to new specifications with little computational effort.

Such a control structure enabling the adaptation to any satisfiable property would need formula-free guarantees on its accuracy. When designing a controller as proposed in Chapter 5, the system is abstracted to a finite state model and can attain formula-free guarantees. Such a finite state model represents the languages that the system can generate and is hence still formula free. The resolution of the control problem for this finite state model finalises the hierarchical control structure. Of interest it the efficient gathering of data that supports the identification of the continuous state model, the control interfaces and also the (formula-free) finite state model.

8.3.d Experiment design for hypothesis testing

The problem of experiment design for hypothesis testing on dynamical models has been only partly investigated. Up to now, priority has been given to experiment design increasing the quality of data-driven models with respect to the resulting performance of model-based control design. But with the increase of attention to the prevention and detection of failures in control engineering, we expect that online hypothesis testing will become more and more important. Henceforth the importance of these techniques will only increase. As mentioned before, further research in the statistical implications of online decision making (sequential hypothesis testing) and in computationally viable approaches to solve the design problems will be necessary.

Bibliography

- A. Abate. Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey. *ENTCS*, 297:3–25, 2013.
- A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic Reachability and Safety for Controlled Discrete Time Stochastic Hybrid Systems. *Automatica*, 44(11):2724–2734, 2008.
- A. Abate, R. C. Hillen, and S. A. Wahl. Piecewise affine approximation of fluxes and enzyme kinetics from in-vivo ^{13}C labeling experiments. *International Journal of Robust and Nonlinear Control*, pages 1120–1139, 2012. Special Issue on System Identification for Biological Systems.
- A. Abate, M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic model checking of labelled Markov processes via finite approximate bisimulations. In *Horizons of the Mind – P. Panangaden Festschrift*, pages 40–58. Springer Verlag, 2014a.
- A. Abate, F. Redig, and I. Tkachev. On the effect of perturbation of conditional probabilities in total variation. *Statistics & Probability Letters*, 2014b.
- R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In D. Sangiorgi and R. de Simone, editors, *Proceedings on Concurrency Theory (CONCUR'98)*, pages 163–178. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998. ISBN 978-3-540-68455-8.
- P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy Build.*, 43(7):1511–1522, 2011.
- C. Baier and J.-P. Katoen. Principles of model checking. *MIT Press*, 2008.
- M. Barenthin, X. Bombois, H. Hjalmarsson, and G. Scorletti. Identification for control of multivariable systems: Controller validation and experiment design via lmis. *Automatica*, 44(12):3070–3078, 2008.
- E. Bartocci, L. Bortolussi, and G. Sanguinetti. Learning temporal logical properties discriminating ECG models of cardiac arrhythmias. *CoRR*, abs/1312.7523, 2013.
- G. Batt, C. Belta, and R. Weiss. Model checking genetic regulatory networks with parameter uncertainty. In *HSCC*, pages 61–75. Springer, 2007.
- C. Belta, L. C. G. J. M. Habets, and V. Kumar. Control of multi-affine systems on rectangles with applications to hybrid biomolecular networks. In *Conf.on CDC*, pages 534–539, 2002.

- C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, and G. J. Pappas. Symbolic planning and control of robot motion [grand challenges of robotics]. *Robotics & Automation Magazine, IEEE*, 14(1):61–70, 2007.
- D. Bertsekas and S. E. Shreve. *Stochastic Optimal control : The discrete time case*. Athena Scientific, 1996.
- R. R. Bitmead, M. R. Gevers, I. R. Petersen, and R. Kaye. Monotonicity and stabilizability – properties of solutions of the Riccati difference equation: Propositions, lemmas, theorems, fallacious conjectures and counterexamples. *Systems & Control Letters*, 5(5):309 – 315, 1985.
- F. Blanchini and S. Miani. *Set-Theoretic Methods in Control*. Birkhäuser Basel, 1st edition, 2007.
- R. Blute, J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled markov processes. In *Logic in Computer Science, 1997. LICS'97. Proceedings., 12th Annual IEEE Symposium on*, pages 149–158. IEEE, 1997.
- V. I. Bogachev. *Measure theory*. Springer Science & Business Media, 2007.
- X. Bombois, G. Scorletti, M. Gevers, P. Van den Hof, and R. Hildebrand. Least costly identification experiment for control. *Automatica*, 42(10):1651–1662, 2006.
- X. Bombois, A. J. D. Dekker, C. R. Rojas, and H. Hjalmarsson. Optimal experiment design for hypothesis testing applied to functional magnetic resonance imaging. *IFAC Proceedings Volumes*, 44(1):9953–9958, 2011.
- V. S. Borkar. *Probability theory: an advanced course*. Springer Verlag, 2012.
- L. Bortolussi and G. Sanguinetti. Learning and designing stochastic processes from logical constraints. In *QEST*, pages 89–105. Springer, 2013.
- L. Bortolussi and G. Sanguinetti. Smoothed model checking for uncertain continuous time Markov chains. *CoRR*, abs/1402.1450, 2014.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. CUP, Cambridge, 2004. ISBN 9780511804441.
- S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, jan 1994. ISBN 978-0-89871-485-2.
- L. Brim, M. Češka, S. Dražan, and D. Šafránek. Exploring parameter space of stochastic biochemical systems using quantitative model checking. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 1–17. Springer, 2013.
- L. Busoniu, R. Babuska, B. D. Schutter, and D. Ernst. *Reinforcement Learning and Dynamic Programming Using Function Approximators*. Automation and Control Engineering. CRC Press, 2010.

- P. E. Caines and D. Q. Mayne. On the discrete time matrix Riccati equation of optimal control – a correction. *International Journal of Control*, 12(5):785–794, 1970a.
- P. E. Caines and D. Q. Mayne. On the discrete time matrix Riccati equation of optimal control. *International Journal of Control*, 12(1):785–794, 1970b.
- D. Cattaruzza, A. Abate, P. Schrammel, and D. Kroening. Unbounded-time analysis of guarded lti systems with inputs by abstract acceleration. In *International On Static Analysis*, pages 312–331. Springer, 2015.
- K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995.
- K. Chatterjee, M. Chmelik, R. Gupta, and A. Kanodia. Qualitative analysis of POMDPs with temporal logic specifications for robotics applications. *CoRR*, abs/1409.3360, 2014.
- Y. Chen and T. D. Nielsen. Active learning of Markov decision processes for system verification. In *Conference on Machine Learning and Applications*, pages 289–294, 2012.
- E. M. Clarke. The birth of model checking. In *25 Years of Model Checking*, pages 1–26. Springer, 2008.
- E. M. Clarke and J. M. Wing. Formal methods: State of the art and future directions. *ACM Computing Surveys (CSUR)*, 28(4):626–643, 1996.
- B. L. Cooley and J. H. Lee. Control-relevant experiment design for multivariable systems described by expansions in orthonormal bases. *Automatica*, 2001.
- H. Cramér. *Methods of mathematical statistics*. Princeton: Princeton University Press, 1946.
- P. R. D’Argenio, B. Jeannot, H. E. Jensen, and K. G. Larsen. Reduction and refinement strategies for probabilistic analysis. In *Proceedings of the Second Joint International Workshop on Process Algebra and Probabilistic Methods, Performance Modeling and Verification, PAPM-PROBMIV ’02*, pages 57–76, London, UK, 2002. Springer-Verlag. ISBN 3-540-43913-7.
- M. Davis and M. Zervos. A new proof of the discrete-time LQG optimal control theorems. *Automatic Control, IEEE Transactions on*, 40(8):1450–1453, Aug 1995.
- V. Desaraju, H. C. Ro, M. Yang, E. Tay, S. Roth, and D. Del Vecchio. Partial order techniques for vehicle collision avoidance: Application to an autonomous roundabout test-bed. In *IEEE International Conference on Robotics and Automation*, pages 82–87. IEEE, 2009.
- J. Desharnais, A. Edalat, and P. Panangaden. A logical characterization of bisimulation for labeled markov processes. In *Logic in Computer Science, 1998. Proceedings. Thirteenth Annual IEEE Symposium on*, pages 478–487. IEEE, 1998.

- J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labelled Markov processes. *Inf. Comput.*, 179(2):163–193, 2002.
- J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating labelled Markov processes. *Inf. Comput.*, 184(1):160–200, jul 2003.
- J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theoretical Computer Science*, 318(3):323–354, 2004.
- J. Desharnais, F. Laviolette, and M. Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. *Conf. on Quantitative Evaluation of Systems*, pages 264–273, Sept. 2008.
- E. W. Dijkstra. Notes on structured programming, 1970.
- X. Ding, M. Lazar, and C. Belta. LTL receding horizon control for finite deterministic systems. *Automatica*, 50(2):399–408, feb 2014.
- A. D’Innocenzo, A. Abate, and J.-P. Katoen. Robust PCTL model checking. In *Proceedings of the 15th ACM international conference on Hybrid Systems: computation and control*, pages 275–285, 2012.
- A. Edalat. Semi-pullbacks and bisimulation in categories of markov processes. *Mathematical Structures in Computer Science*, 9(05):523–543, 1999.
- S. Esmail Zadeh Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Proc. of Quantitative Evaluation of Systems*, pages 59–68, Aachen, Germany, 2011.
- S. Esmail Zadeh Soudjani and A. Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.
- S. Esmail Zadeh Soudjani, C. Gevaerts, and A. Abate. FAUST ²: Formal Abstractions of Uncountable-STATE STOchastic Processes. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Lecture Notes in Computer Science, pages 272–286. Springer Berlin Heidelberg, 2015.
- S. S. Farahani, V. Raman, and R. M. Murray. Robust Model Predictive Control for Signal Temporal Logic Synthesis. *IFAC-PapersOnLine*, 48(27):323–328, 2015.
- G. F. Franklin, J. D. Powell, and M. L. Workman. *Digital control of dynamic systems*. Addison-Wesley Menlo Park, second edition, 1990.
- G. F. Franklin, D. J. Powell, and A. Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 4th edition, 2001. ISBN 0130323934.
- G. Frehse, S. K. Jha, and B. H. Krogh. A counterexample-guided approach to parameter synthesis for linear hybrid automata. In *HSCC*, pages 187–200. Springer Berlin Heidelberg, 2008.

- G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler. Spaceex: Scalable verification of hybrid systems. In S. Q. Ganesh Gopalakrishnan, editor, *Proc. 23rd International Conference on Computer Aided Verification (CAV)*, LNCS. Springer, 2011.
- J. Fu and U. Topcu. Probably approximately correct MDP learning and control with temporal logic constraints. *CoRR*, abs/1404.7073, 2014.
- L. Gerencsér and H. Hjalmarsson. Adaptive input design in system identification. In *Proceedings of the 44th IEEE Conference on Decision and Control*, pages 4988–4993, Dec 2005.
- L. Gerencsér, H. Hjalmarsson, and J. Martensson. Identification of ARX systems with non-stationary inputs – asymptotic analysis with application to adaptive input design. *Automatica*, 45(3):623 – 633, 2009.
- M. Gevers. Identification for control: From the early achievements to the revival of experiment design. *European Journal of Control*, 11(4):335 – 352, 2005.
- M. Gevers, X. Bombois, R. Hildebrand, and G. Solari. Optimal experiment design for open and closed-loop system identification. *Communications in Information and Systems*, 11(3):197–224, 2011.
- R. Ghaemi and D. D. Vecchio. Control for Safety Specifications of Systems With Imperfect Information on a Partial Order. *IEEE Trans. Automat. Contr.*, 59(4): 982–995, 2014.
- A. Girard and G. J. Pappas. Approximate bisimulation relations for constrained linear systems. *Automatica*, 43(8):1307–1317, aug 2007.
- A. Girard and G. J. Pappas. Hierarchical control system design using approximate simulation. *Automatica*, 45(2):566–571, 2009.
- S. Giro and M. N. Rabe. Verification of partial-information probabilistic systems using counterexample-guided refinements. In *Proceedings on Automated Technology for Verification and Analysis*, LNCS, pages 333–348. Springer, 2012.
- G. Goodwin, J. C. Murdoch, and R. L. Payne. Optimal test signal design for linear S.I.S.O. system identification. *International Journal of Control*, 17(1):45–55, 1973.
- A. Gosavi. Reinforcement learning: A tutorial survey and recent advances. *INFORMS Journal on Computing*, 21(2):178–192, 2009.
- G. Gössler. Component-based modeling and reachability analysis of genetic networks. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 8(3): 672–682, 2011.
- M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- S. Gugercin and A. C. Antoulas. A survey of model reduction by balanced truncation and some new results. *International Journal of Control*, 77(8):748–766, 2004.

- B. M. Gyori, D. Paulin, and S. K. Palaniappan. Probabilistic verification of partially observable dynamical systems. *CoRR*, abs/1411.0976, 2014.
- S. Haesaert, R. Babuska, and A. Abate. Sampling-based Approximations with Quantitative Performance for the Probabilistic Reach-Avoid Problem over General Markov Processes. *CoRR*, abs/1409.0553, sep 2014.
- S. Haesaert, A. Abate, and P. M. J. Van den Hof. Correct-by-design output feedback of LTI systems. In *Proceedings of the Conference on Decision and Control*, pages 6159–6164, Dec 2015a.
- S. Haesaert, P. M. J. Van den Hof, and A. Abate. Data-driven property verification of grey-box systems by Bayesian experiment design. In *American Control Conference*, pages 1800–1805, 2015b.
- S. Haesaert, P. M. J. Van den Hof, and A. Abate. Data-driven and model-based verification: A bayesian identification approach. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 6830–6835. IEEE, 2015c.
- S. Haesaert, P. M. J. Van den Hof, and A. Abate. Experiment design for formal verification via stochastic optimal control. In *European Control Conference*, 2016.
- J. Helton, M. de Oliveira, M. Stankus, and R. Miller. Ncalgebra, 2015 release edition. <http://www.math.ucsd.edu/ncalg>, 2015.
- D. Henriques, J. G. Martins, P. Zuliani, A. Platzer, and E. M. Clarke. Statistical model checking for Markov decision processes. In *QEST*, pages 84–93, 2012.
- T. Henzinger and H. Wong-Toi. Using hytech to synthesize control parameters for a steam boiler. In *Formal Methods for Industrial Applications*, pages 265–282. Springer Berlin Heidelberg, 1996. doi: 10.1007/BFb0027241.
- O. Hernández-Lerma and J. B. Lasserre. *Discrete-time Markov control processes*, volume 30 of *Applications of Mathematics (New York)*. Springer Verlag, 1996.
- P. S. C. Heuberger, P. M. J. Van den Hof, and O. H. Bosgra. A generalized orthonormal basis for linear dynamical systems. *Automatic Control, IEEE Transactions on*, 40(3):451–465, 1995.
- P. S. C. Heuberger, P. M. J. Van den Hof, and B. Wahlberg. *Modelling and identification with rational orthogonal basis functions*. Springer London, 2005.
- R. Hildebrand and M. Gevers. Minimizing the worst-case ν -gap by optimal input design. In *Proceedings of 13th IFAC Symposium on System Identification*, pages 665–670, 2003.
- H. Hjalmarsson. Iterative feedback tuning an overview. *International journal of adaptive control and signal processing*, 16(5):373–395, jun 2002.
- H. Hjalmarsson. From experiment design to closed-loop control. *Automatica*, pages 393–438, 2005.

- H. Hjalmarsson. System identification of complex and structured systems. *European journal of control*, 15(3):275–310, 2009.
- H. Hjalmarsson, M. Gevers, and F. De Bruyne. For model-based control design, closed-loop identification gives better performance. *Automatica*, 32(12):1659–1673, 1996.
- O. Holub and K. Macek. HVAC simulation model for advanced diagnostics. In *Symp. Intelligent Signal Processing*, pages 93–96. IEEE, Sept. 2013.
- X. Huan and Y. M. Marzouk. Sequential Bayesian optimal experimental design via approximate dynamic programming. *CoRR*, abs/1404.7073:34, apr 2016.
- V. Ionescu and M. Weiss. Continuous and discrete-time Riccati theory: A Popov-function approach. *Linear Algebra and its Applications*, 193:173–209, Nov. 1993.
- H. Jansson. *Experiment design with applications in identification for control*. PhD thesis, KTH, Stockholm, 2004.
- B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proceedings Sixth Annual IEEE Symposium on Logic in Computer Science*, pages 266–277, July 1991. doi: 10.1109/LICS.1991.151651.
- D. P. Joseph and T. J. Tou. On linear control theory. *American Institute of Electrical Engineers, Part II: Applications and Industry, Transactions of the*, 80(4):193–196, Sept 1961.
- A. A. Julius and G. J. Pappas. Approximations of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 2009.
- R. M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19(7):371–384, 1976.
- Y. Kouskoulas, D. Renshaw, A. Platzer, and P. Kazanzides. Certifying the safe design of a virtual fixture control algorithm for a surgical robot. In *Proceedings of the 16th international conference on Hybrid systems: computation and control*, pages 263–272. ACM, 2013.
- J. Kreiker, A. Tarlecki, M. Y. Vardi, and R. Wilhelm. Modeling, analysis, and verification—the formal methods manifesto 2010 (dagstuhl perspectives workshop 10482). *Dagstuhl Manifestos*, 1(1), 2011.
- M. Lahijanian, S. B. Andersson, and C. Belta. A probabilistic approach for control of a stochastic system from LTL specifications. In *Proceedings of the Conference on Decision and Control*, pages 2236–2241, 2009.
- M. Lahijanian, J. Wasniewski, S. B. Andersson, and C. Belta. Motion planning and control from temporal logic specifications with probabilistic satisfaction guarantees. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3227–3232, 2010.

- K. G. Larsen and A. Skou. Bisimulation through probabilistic testing (preliminary report). In *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '89*, pages 344–352, New York, NY, USA, 1989. ACM. ISBN 0-89791-294-2. doi: 10.1145/75277.75307. URL <http://doi.acm.org/10.1145/75277.75307>.
- C. A. Larsson. *Application-oriented experiment design for industrial model predictive control*. PhD thesis, KTH, 2014.
- A. Legay and S. Sedwards. Lightweight Monte Carlo algorithm for Markov decision processes. *CoRR*, abs/1310.3609, 2013.
- A. Legay, B. Delahaye, and S. Bensalem. Statistical model checking: An overview. In H. Barringer, Y. Falcone, B. Finkbeiner, K. Havelund, I. Lee, G. Pace, G. Roşu, O. Sokolsky, and N. Tillmann, editors, *Runtime Verification*, volume 6418 of *LNCS*, pages 122–135. Springer Berlin Heidelberg, 2010.
- K. Lesser and A. Abate. Controller Synthesis for Probabilistic Safety Specifications using Observers. *IFAC-PapersOnLine*, 48(27):329–334, 2015.
- K. Lesser and M. Oishi. Reachability for partially observable discrete time stochastic hybrid systems. *Automatica*, 50(8):1989–1998, 2014.
- D. V. Lindley. The philosophy of statistics. *Journal of the Royal Statistical Society: Series D (The Statistician)*, pages 293–337, 2000.
- T. Lindvall. *Lectures on the coupling method*. Courier Corporation, 2002.
- L. Ljung. *System identification - Theory for the user*. PTR Prentice Hall, 2nd edition, 1999.
- H. Mao and M. Jaeger. Learning and model-checking networks of I/O automata. In *Proc. of Asian Conference on Machine Learning*, 2012.
- M. Mazo Jr, A. Davitian, and P. Tabuada. PESSOA: towards the automatic synthesis of correct-by-design control software. In *Work-in-progress HSCC*, 2010.
- A. Mesbah, X. Bombois, M. Forgone, J. H. A. Ludlage, P. E. Modèn, H. Hjalmarsson, and P. M. J. Van den Hof. A unified experiment design framework for detection and identification in closed-loop performance diagnosis. In *51st IEEE Conference on Decision and Control (CDC)*, pages 2152–2157, Dec 2012a.
- A. Mesbah, X. Bombois, J. H. A. Ludlage, and P. M. J. Van den Hof. Experiment design for closed-loop performance diagnosis. *IFAC Proceedings Volumes*, 45(16): 1341–1346, 2012b.
- A. Mesbah, X. Bombois, M. Forgone, H. Hjalmarsson, and P. M. J. Van den Hof. Least costly closed-loop performance diagnosis and plant re-identification. *International Journal of Control*, 88(11):2264–2276, 2015.
- S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Communications and Control Engineering Series. Springer-Verlag London Ltd., 1993.

- O. Mickelin, N. Ozay, and R. M. Murray. Synthesis of correct-by-construction control protocols for hybrid systems using partial state information. In *2014 American Control Conference*, pages 2305–2311. IEEE, 2014.
- R. Munos and C. Szepesvári. Finite Time Bounds for Fitted Value Iteration. *Journal in Machine Learning Research*, 9:815–857, 2008.
- V. Peterka. Bayesian Approach to System Identification. *Trends Prog. Syst. Identif.*, 1981.
- A. Pnueli. The temporal logic of programs. In *Symposium on Foundations of Computer Science*, pages 46–57. IEEE, 1977.
- E. Polgreen, V. B. Wijesuriya, S. Haesaert, and A. Abate. Data-efficient bayesian verification of parametric markov chains. In *Quantitative Evaluation of Systems*, volume 9826 of *Lecture Notes in Computer Science*, pages 35–51. Springer, 2016.
- L. Pronzato. Adaptive optimization and D-optimum experimental design. *Annals of Statistics*, pages 1743–1761, 2000.
- L. Pronzato. Optimal experimental design and some related control problems. *Automatica*, 44(2):303 – 325, 2008.
- L. Pronzato and É. Walter. Robust experiment design via stochastic approximation. *Mathematical Biosciences*, 75(1):103–120, 1985.
- L. Pronzato and É. Walter. Robust experiment design via maximin optimization. *Mathematical Biosciences*, 89(2):161–176, 1988.
- A. Puri and P. Varaiya. Decidability of hybrid systems with rectangular differential inclusions. In D. L. Dill, editor, *International Conference on Computer Aided Verification*, Lecture Notes in Computer Science, pages 95–104. Springer Berlin Heidelberg, 1994.
- B. C. Reginato, R. Z. Freire, G. H. D. C. Oliveira, N. Mendes, and O. Abadie, Marc. Predicting the temperature profile of indoor buildings by using orthonormal basis functions. In *Conf. on Building Performance Simulation Association*, United Kingdom, 2009.
- C. R. Rojas, J. S. Welsh, G. C. Goodwin, and A. Feuer. Robust optimal experiment design for system identification. *Automatica*, 43(6):993–1008, 2007.
- N. Saeedloei and G. Gupta. A logic-based modeling and verification of CPS. *ACM SIGBED Review*, 8(2):31–34, 2011.
- C. Scherer and S. Weiland. Linear matrix inequalities in control. *Lecture Notes, Dutch Institute for Systems and Control, Delft, The Netherlands*, 3, 2000.
- C. Scherer, P. Gahinet, and M. Chilali. Multi-objective output-feedback control via LMI optimization. *IEEE Transactions on Automatic Control*, 42:896–911, 1997.

- S. Schupp, E. Ábrahám, X. Chen, I. B. Makhoulf, G. Frehse, S. Sankaranarayanan, and S. Kowalewski. Current challenges in the verification of hybrid systems. In *International Workshop on Design, Modeling, and Evaluation of Cyber Physical Systems*, pages 8–24. Springer, 2015.
- R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, Massachusetts Institute of Technology, 1995.
- R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 1995.
- K. Sen, M. Viswanathan, and G. Agha. Learning continuous time Markov chains from sample executions. In *QEST*, pages 146–155, 2004a. doi: 10.1109/QEST.2004.1348029.
- K. Sen, M. Viswanathan, and G. Agha. Statistical model checking of black-box probabilistic systems. In R. Alur and D. Peled, editors, *CAV*, volume 3114 of *LNCIS*, pages 202–215. Springer, 2004b.
- K. S. Shanmugan and A. M. Breipohl. *Random Signals: Detection, Estimation and Data Analysis*. Wiley, 1988.
- H. J. Skala. The existence of probability measures with given marginals. *Ann. Probab.*, 21(1):136–142, 01 1993.
- J. D. Stigter, D. Vries, and K. J. Keesman. On adaptive optimal input design: a bioreactor case study. *AIChE J.*, 52(9):3290–3296, 2006.
- A. A. Stoorvogel and A. Saberi. The discrete algebraic Riccati equation and linear matrix inequality. *Linear Algebra Appl.*, 274(1-3):317–365, apr 1998.
- V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36(2):423–439, 04 1965.
- C. Striebel. Sufficient statistics in the optimum control of stochastic systems. *Journal of Mathematical Analysis and Applications*, 12(3):576 – 592, 1965.
- S. Summers and J. Lygeros. Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46:1951–1961, 2010.
- P. Tabuada. *Verification and control of hybrid systems*. Springer US, 2009.
- P. Tabuada and G. J. Pappas. Model checking LTL over controllable linear systems is decidable. *Hybrid Systems: Computation and Control*, 2623:498–513, 2003.
- P. Tabuada and G. J. Pappas. Linear time logic control of discrete-time linear systems. *IEEE Transactions on Automatic Control*, 51(12):1862–1877, 2006.
- C. W. Therrien. *Discrete random signals and statistical signal processing*. Prentice Hall, 1992.
- I. Tkachev and A. Abate. On infinite-horizon probabilistic properties and stochastic bisimulation functions. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 526–531. IEEE, 2011.

- Transportation Research Board, the Committee on Electronic Vehicle Controls and Unintended Acceleration, and National Research Council (US). *The Safety Promise and Challenge of Automotive Electronics: Insights from Unintended Acceleration*, volume 308. Transportation Research Board, 2012.
- P. M. J. Van den Hof and R. J. P. Schrama. Identification and control—closed-loop issues. *Automatica*, 31(12):1751–1770, 1995.
- P. M. J. Van den Hof, P. S. C. Heuberger, and J. Bokor. System identification with generalized orthonormal basis functions. *Automatica*, pages 1821–1834, 1995.
- M. Y. Vardi. From philosophical to industrial logics. In *Logic and Its Applications*, pages 89–115, Berlin, Heidelberg, 2009. Springer-Verlag.
- G. S. Virk and D. L. Loveday. Model-based control for HVAC applications. In *Conf. on Control Applications*, pages 1861–1866. IEEE, 1994.
- A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.
- H. Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971.
- W. M. Wonham. On the separation theorem of stochastic control. *SIAM Journal on Control*, 6(2):312–326, 1968.
- J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald. Formal methods. *ACM Computing Surveys (CSUR)*, 41(4):1–36, Oct. 2009.
- B. Yordanov, G. Batt, and C. Belta. Model Checking Discrete-Time Piecewise Affine Systems: Application to Gene Networks. In *Europa Control Conference*, 2007.
- M. Zamani, P. Mohajerin Esfahani, R. Majumdar, A. Abate, and J. Lygeros. Symbolic control of stochastic systems via approximately bisimilar finite abstractions. *IEEE Transactions on Automatic Control*, 59(12):3135–3150, Dec 2014.
- M. Zarrop. A chebyshev system approach to optimal input design. *IEEE Transactions on Automatic Control*, 24(5):687–698, Oct 1979.
- C. Zhang and J. Pang. On probabilistic alternating simulations. In C. Calude and V. Sassone, editors, *Theoretical Computer Science*, volume 323 of *IFIP Advances in Information and Communication Technology*, pages 71–85. Springer Berlin Heidelberg, 2010. ISBN 978-3-642-15239-9.
- L. Zhang, H. Hermanns, and D. N. Jansen. Logic and model checking for hidden Markov models. In *Proc. on Formal Techniques for Networked and Distributed Systems*, pages 98–112. Springer, 2005.
- M. Zhivich and R. K. Cunningham. The real cost of software errors. *IEEE Security Privacy*, 7(2):87–90, March 2009.

P. Zuliani, A. Platzer, and E. M. Clarke. Bayesian statistical model checking with application to Stateflow/Simulink verification. *Formal Methods in System Design*, pages 338–367, 2013.

Summary

Verification and control of physical systems: Data-driven and model-based approaches

Sofie Haesaert

The use of design software for the formal verification and synthesis of critical hardware and software has led to a reduction of design mistakes within computer sciences. This raises the question whether we can also develop semi-automated methods to support control engineers in the verification of physical systems and the design of controllers. In this thesis, I look at how such methods should be tailored to engineering design problems, that is, to the problems present when dealing with physical systems.

Within control engineering dealing with uncertainty, expressed in a probabilistic framework, is key. It is present in stochastic state-transitions of many physical phenomena, in the disturbance on sensor measurements, and often also as a lack of knowledge or uncertainty on the model dynamics. Throughout the thesis, methods are explored to deal with these kinds of uncertainty, thereby paving the way to new controller synthesis methods.

To aid in the synthesis of controllers for stochastic systems operating over uncountable state spaces, I give new results for hierarchical controller refinement, which underpinned by the use of metrics, allows in particular for a useful trade-off between deviations over probability distributions on states, and distances between model outputs.

For the specific case of linear time invariant systems disturbed by Gaussian noise on state transitions and output measurements, I extend the literature on correct-by-design controller synthesis to output-based correct-by-design controller synthesis.

Furthermore, I develop a new measurement-driven and model-based formal verification approach, applicable to dynamical systems with partly unknown dynamics. Since the information content of the measurements in the data-driven verification approach is crucial, in this thesis experiment design is formulated via stochastic optimal control to verify the system as data-efficiently as possible.

Acknowledgments

During my final year of my masters' degree, I wrote a proposal for my own Ph.D. project under the guidance of my current supervisors Paul and Alessandro. This thesis presents my scientific journey trying to push the ideas behind identification for control into the field of formal verification and automated controller design. My experience as a PhD student has been greatly shaped by the people around me, and I will do my best to acknowledge their much appreciated support, presence and guidance.

First of all, I would like to thank Alessandro. I value the constant advice and guidance you have given me in the past five years. After being my supervisor during my masters' thesis you agreed to spend another four years as a co-promotor for my PhD degree. Even though your move to Oxford created a physical distance between us, it never created a big hurdle. I greatly admire your qualities as a supervisor, and appreciate the freedom you have given me to explore and learn, while still giving support when needed. You have taught me many important things, like how to write and how to supervise Master students. As a mentor you have fostered my entrance into academia, who at key times was able to put everything in perspective for me. Thank you.

As my promotor at TU Eindhoven, Paul has had a crucial supervisory role in the past four years. Paul, I would like to thank you for allowing me to grow as a researcher. I appreciate the trust you have put in me by allowing me to drive my own research. You taught me the importance of problems statements and of explaining my work step-by-step. Together stuck in a multidisciplinary field, the struggle to communicate clearly on ideas across different fields and backgrounds has been an enriching experience.

I would like to thank the members of my PhD committee prof.dr.ir. J. Schoukens, prof.dr.ir. A. Rantzer, prof.dr.ir. J.F. Groote, prof.dr. K.G. Larsen and prof.dr. S. Weiland. A special word of thanks goes to professors Robert Babuska and Frank Lewis under whose supervision I did my first research during my masters' degree. This experience made the decision to do a PhD so much easier. The most enriching part of doing research is attending, participating and starting challenging discussions. For this reason I have enjoyed collaborating with great minds like Siep, Carsten, Sadegh, Kendra, and Elizabeth.

After many visits to Oxford, this wonderful city with its eccentric university has stolen a part of my heart. This is also due to the warm welcome offered by Alessandro and his family. I fondly remember many interesting conversation during

the 10 am meetings at the espresso machine with Alessandro and his growing research group. With "duck, die ..." my Oxfordian office mates changed my knowledge of the English language. I value the evenings spent with Sadegh, Nathalie, Kendra and Dario and many others eating a nice meal, drinking excellent wine, or sipping at a cocktail while looking over the scenery of Oxford.

I have been a proud member of the Control Systems group of TU Eindhoven. The many presentations and interesting discussions have been very valuable. Together with my colleagues I enjoyed coffee breaks, lunch breaks and social get-togethers in the past four years. Many of my close friendships developed during my stay in this group.

I am grateful for having a group of wonderful friends in my home town, who have never asked me to do the normal thing. The same goes for my Dutch friends; Chris, Pepijn and Robbert. Pepijn, you chose the same group as me in Eindhoven to do your PhD and having a good friend like you at work has been really nice. I always enjoy our dinners with that bottle of wine (or two).

I would like to thank my family and friends, both new and old, for their support and friendship. Mum and dad, thank you for giving me a warm home and for supporting me throughout my studies. With your belief and encouragement, I have been able to develop to the person I am today.

Sofie
Eindhoven, 2017

David Alina Marjan Yanin Quan Leyla Valerie
 Dhruv Edwin Jolien Constantijn Sena Jan Henrik
 Daming **Paul** Koen Marcella Bahadir Ruxandra **Pepijn**
 Margo Christof **Thank you** Alejandro
 Ruben Thijs Riet Ioannis Robbert Harm Ingrid Tom
 Michel Chris Arne Diana Alessandro
 Siep **PAPA** Tuan 陈飞 Jeroen Hernan Mircea
 Barbara Roland

Curriculum Vitae

Sofie Haesaert was born in Leuven, Belgium on the 2nd of May, 1989. She finished high-school in the Paridaens institute in Leuven in 2007. Afterwards she started studying for her Bachelor Mechanical Engineering at the TU Delft. In 2010, she received her B.Sc. degree cum laude. She graduated cum laude the Systems and Control M.Sc. degree in 2012 at the Delft Center for Systems and Control, TU Delft. In 2013, she started working towards her Ph.D. degree at the Control Systems group of the Electrical Engineering Department at TU Eindhoven. Her Ph.D. project, on the topic of identification for control of dynamic networks in Biology, was funded by NWO in the scope of the NWO graduate Program organised by DISC.

Her investigations into the topic of identification for control together with dr. A. Abate, who moved to the Computer Science department at Oxford University around that time, put forward the question what control would entail for dynamic networks. During her visits at Oxford University, the novel area of automatic verification and generation of control software for engineering systems caught her interest. Her thesis is on the development of data-driven methods and model-based methods for physical systems.

Her research interests cover data-driven methods such as identification and reinforcement learning and numerical tools for controller design.

