

MODEL-BASED AND
DATA-DRIVEN FORMAL
SYNTHESIS OF POWER
SYSTEMS

Ben Wooding

MODEL-BASED AND DATA-DRIVEN FORMAL SYNTHESIS OF POWER SYSTEMS

THESIS

submitted for the degree of
Doctor of Philosophy

by

Ben WOODING

School of Computing
Newcastle University



2 August 2023

This thesis has been approved by the supervisors:

Prof. Sadegh Soudjani

Prof. John Fitzgerald

Defence Panel:

Prof. Antoine Girard,

Prof. Ken Pierce,

Centre national de la recherche scientifique (CNRS)

Newcastle University, UK



**Engineering and
Physical Sciences
Research Council**

The work presented in this thesis has been supported by *United Kingdom Research and Innovation (UKRI)*, funded by the *Engineering and Physical Sciences Research Council (EPSRC)*, through an EPSRC Studentship (EP/R51309X/1).

Copyright © 2023 by Ben Wooding.

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized 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.

*To my wife Lucy,
for her unending love and support*

Acknowledgments

The fear of the Lord is the beginning of wisdom...

Proverbs 9:10

This thesis is the culmination of 4 years of research at the School of Computing at Newcastle University. I want to take this opportunity to show my appreciation to those who have helped me during my research journey.

The first thank you goes to my supervisor, Sadegh Soudjani. I am extremely grateful to him for accepting me as a PhD candidate, for our regular meetings (particularly during lockdown in the pandemic), and the relationship we developed both professionally and personally. When I started, I had no understanding of control theory or power systems, and my mathematics was covered in around 4 years of rust. Under his guidance, I have grown hugely in all of these areas. Sadegh has always been very approachable, he has a brilliant way of describing tricky concepts in easy to understand ways and he encourages high-quality results without creating any additional pressure. He is undoubtedly a brilliant scientist but his greatest qualities are his easy-going nature, his kind personality and his infectious laugh. He has been a huge role model for me to continue pursuing research past this PhD.

Secondly, I want to thank my collaborators with whom I've had the most interactions; Vahid Vahidinasab, Abolfazl Lavaei, Mahmoud Salamati, and Milad Kazemi. Their efforts and contributions to my research directions have been invaluable and I have learned so much from our conversations and discussions.

Thirdly, I want to thank the members of AMBER and HyCoDeV. I want to thank Oliver Schön and Milad Kazemi, with whom I have shared a particularly close relationship during my PhD. Among others, I've particularly enjoyed the interesting lunch debates and hallway interactions I've had with John, Zhi, Marco, Ivan, Dasha, Olivia, Kostya, Alessandro, Paulius, Zhichao, Anna, Selva, Abdel, Sam, Negar, Mahdieh, Omid, and Andrea. I also want to thank my office-mates Ani, Isi, Ziqi and Alex. I have appreciated the advice from the professors and lecturers at the school including; Leo Freitas, Nigel Thomas, John Fitzgerald, Ken Pierce, Cliff Jones, Isi Mitrani, Charles Morisset, Laura Heels, Wanqing Zhao, Maciej Koutny and Sergiy Bogolomov.

I want to recognise those friends with whom I spend my social time with outside of work. Including the regular teasing I get for trying to become "not a real doctor"! I have enjoyed the board games, movie nights, sports watching, shared meals, spikeball tournaments, frisbee matches, and wedding trips. Most of all, they remind me of the important things in life and walk alongside me in my faith. I want to give special mention to Eugene, Kathryn, Johnny, John, Laura, Andy, Nigel, Sam, Josh, and Rory.

I want to give special mention to my parents and wider family who have raised me to be the man I am today. I am so grateful for all the care I have received in the last 26 years. It is particularly sweet to have the conversations with mum as she tries to understand what on earth I do. Writing notes before ultimately saying: "so, something to do with electric vehicles?". And to my dad who has inspired me to set high standards, strive to do my best in all things, and who is the prime example of what it is like to never stop learning.

Finally, I want to thank my beautiful, patient and wonderful wife Lucy. Her endless support has been so encouraging when things got tough, and I will always have special memories from during the pandemic when I was included in a house bubble together with her and Kathryn; eating together, playing board games and piercing each other's ears. Lucy keeps me humble, loves to share a good laugh, and always points me back to Jesus.

*Ben Wooding
Newcastle, June 2023*

Summary

Model-Based and Data-Driven Formal Synthesis for Power Systems

Ben Wooding

This thesis is motivated by the increased embedding of cyber components inside physical real-world systems where properties such as safety are paramount. Smart energy systems are one such example of *cyber-physical systems* (CPS). In a world looking for net-zero emissions, uncertain renewable generation sources are being relied on over traditional turbine generators. Additionally, *plug-in electric vehicles* (EVs) are becoming attractive options of reducing an individual's carbon footprint at an increasingly affordable price point.

The systemwide delicate balance between power generation and consumption is captured by the power network's frequency. *Frequency regulation*, the control mechanism to maintain safe operating limits, is a pressing area for research, with serious instability causing potential load shedding, blackouts, or cascading failures. Supervisory Control and Data Acquisition, PID control and model-predictive control are examples of some regulation techniques. At present, these techniques do not provide guarantees for the system's behaviour, e.g. satisfying grid code standards, and they may require manual intervention in emergency scenarios.

A growing area in computer science is formal methods, with the goal of verifying that system controlled by a software, satisfy a given *formal specification*. The formal specification uses temporal logic to define *complex logical properties* the system should satisfy. Examples of these properties include safety, reachability, or reach-avoid. *Formal controllers* can be synthesised with guarantees on the satisfaction of the specification. One of the main formal controller synthesis approaches is based on the construction of simplified *abstract* models using state-space discretisation. This is an immensely powerful approach but it suffers from the *curse of dimensionality*. As the number of state variables increases, the size of the abstract model explodes exponentially, making computations intractable at higher dimensions. Due to this, case studies and results in the formal methods community tend to use simpler academic examples with low dimensions.

This thesis brings together the concepts from power systems, computer science, and control engineering. The power system community benefits from the tangible guarantees provided by formal control approaches, and the formal methods

community benefit from complex real-world case studies. In particular, this thesis provides several novel contributions:

- This thesis proposes a formal controller synthesis approach for integrating a population of EVs for centralised continuous-time frequency regulation of power systems. This approach was the first application of formal methods to the frequency regulation of smart grids. A novel symbolic controller using abstraction-based schemes for the Great Britain power system is designed and simulated under a large outage event. The symbolic controller satisfies a specification with formal guarantees that the frequency returns to a specified safe interval unlike the baseline controller taken from literature.
- Furthermore, this thesis studies formal synthesis of centralised controllers for continuous-space systems with unknown dynamics to satisfy requirements expressed as linear temporal logic formulas over finite and infinite horizons. As formal abstraction-based synthesis schemes rely on a precise mathematical model of the system to build a finite abstract model, the abstraction-based schemes are not applicable when the dynamics of the system are unknown. The approach casts the computation of the growth bound of the system as a robust convex optimisation program (RCP). Since the unknown dynamics appear in the optimisation, a scenario convex program (SCP) is formulated corresponding to the RCP using a finite number of sampled trajectories. The growth bound together with the sampled trajectories are then used to construct the abstraction and synthesise a controller. The performance of the approach is demonstrated on a reduced-order power system model.
- Model reduction involves loss of information from the original system which is not accounted for formally. Simulation functions are Lyapunov-like functions that relate the output trajectories of two systems, with the mismatch between the two systems remaining within some guaranteed error bounds. This thesis approximates concrete systems with large perturbations by reduced-order abstract models. It develops *robust simulation functions* (RSF) further to consider the perturbation in the abstract system by designing an interface function for the disturbance. Accordingly, this enables controllers designed using the reduced-order form of the concrete system and reduces the computational load required for formal synthesis. The efficacy of the approach is demonstrated by synthesising a formal controller for a 9-state area of New England 39-Bus Test System (NETS), using only a 3-state abstract system.
- Finally, this thesis presents an assume-guarantee approach to decentralised compositional control of the 27-state NETS. Based on RSFs with disturbance refinement alongside the composition of multiple subsystems, the approach tackles the scalability problem associated with the *curse of dimensionality*, particularly for synthesising controllers for high-dimensional systems. This thesis proposes two control methods to provide guarantees for NETS: one using the principle of interconnected synchronous machines and another considering the power flows in the network between neighbouring subsystems.

In summary, this thesis contributes to the scalability of formal approaches through compositionality and robust simulation functions with disturbance refinement. Both model-based and data-driven controllers are synthesised for real-world examples relating to the frequency regulation of smart grids.

This outcomes of this thesis point toward some interesting future directions. To adapt to the evolving smart grid, fully distributed formal control techniques would be of interest using multi-agent control schemes. Increased complexities of smart grids would encourage further developments in data-driven control techniques including the need for parallelised tool implementations. Robust simulation functions provide an valuable support to quantify the error from model-order reduction techniques, these can be extended to find optimal interface functions which provide guarantees on the maximal error between trajectories and upper bounds on control inputs.

Contents

Acknowledgments	vii
Summary	ix
1 Introduction	1
1.1 Motivation	1
1.2 Research Goals and Original Contributions	3
1.3 Overview of the Thesis	4
1.4 Publications by the Author	5
1.5 Notation	6
2 Smart Grid: Control and Management	7
2.1 Introduction	7
2.2 Aggregate Models for a Population of Buildings	8
2.3 Frequency Control	10
2.4 System Inertia	11
2.5 Coordinated Volt/Var Control	12
2.5.1 Transmission Network	12
2.5.2 Distribution Network	13
2.5.3 Transmission/Distributed System Operator Coordination	14
2.6 Coordinated Control of Buildings as a Multi-Vector Nano Energy Hub	14
2.7 Security Aspects of Coordinated Control of Active Buildings as a Cyber-Physical System	16
2.8 Aspects Considered in this Thesis	18
2.9 Conclusion	18

3	Smart Grid: Control Techniques	19
3.1	Introduction	19
3.2	Coordination Structures for Management and Control of the Energy Systems	20
3.2.1	Coordinated Structures from the System Perspective	20
3.2.2	Coordinated Structures from an Energy Resource Perspective	22
3.2.3	Coordinated Structures from a Security Perspective	23
3.2.4	Supervisory Control and Data Acquisition (SCADA)	23
3.3	Control Techniques for Active Buildings	24
3.3.1	PID Control	24
3.3.2	Model Predictive Control (MPC)	25
3.3.3	Multi-Agent System Control (MAS)	27
3.3.4	Artificial Intelligence and Data-Driven Control	29
3.3.5	Game Theoretic Approaches	33
3.4	Why consider formal methods?	34
3.5	Conclusion	35
4	Formal Control Techniques	37
4.1	Introduction	37
4.2	Describing Systems	37
4.3	System Behaviour	39
4.4	Exact System Relationships	40
4.5	Symbolic Models	41
4.6	System Specifications	42
4.7	Formal Control Synthesis	45
4.8	Approximate System Relationships	46
4.9	System Composition	47
4.9.1	Subsystems	47
4.9.2	Assume-Guarantee Contracts	48
4.10	Conclusion	49
5	Formal Synthesis for Frequency Regulation of Power Systems	51
5.1	Introduction	51
5.2	Frequency Control in Power Systems	53
5.2.1	Frequency Regulation	54
5.2.2	Requirements on Frequency	54
5.2.3	The GB Model	54

5.2.4	Baseline Controller	55
5.2.5	Baseline Simulation	56
5.3	Temporal Logic	57
5.3.1	Formalising the Specification for Frequency	57
5.4	Formal Controller Synthesis	58
5.4.1	Grid as a Dynamical System	59
5.4.2	Symbolic Model of the Grid	59
5.4.3	Symbolic Control for the Grid	61
5.5	Implementation Results	61
5.5.1	Simulations with a Multi-Phase Controller	61
5.5.2	Formal Guarantees	65
5.5.3	Robustness of the Controller	65
5.6	Using Energy Storage Systems for Frequency Regulation	65
5.7	Conclusion	69
6	Data-Driven Abstraction-Based Control Synthesis	71
6.1	Introduction	72
6.1.1	Contributions	72
6.1.2	Related Work.	74
6.2	Preliminaries and Problem Statement	77
6.2.1	Preliminaries	77
6.2.2	Problem Statement	79
6.3	Robust Convex Programs	79
6.4	Data-Driven Abstraction	80
6.4.1	Growth Bound for Reachable Sets	81
6.4.2	SCP for the Computation of Growth Bound	82
6.4.3	Lipschitz Constant Estimation	84
6.5	Synthesis via Abstraction Refinement	87
6.6	Experimental Evaluation	89
6.6.1	DC-DC Boost Converter	89
6.6.2	Three Area Three Machine Power System	90
6.6.3	Comparison with PAC Learning	96
6.6.4	Parameter Optimisation	97
6.7	Conclusion	99

7	Robust Simulation Functions with Disturbance Refinement	101
7.1	Introduction	102
7.1.1	Motivations and State of the Art.	102
7.1.2	Original Contributions.	103
7.1.3	Related Work.	103
7.2	Preliminaries	104
7.3	Solution Methodologies	105
7.3.1	Robust Approximate Simulation with Disturbance Refinement	105
7.3.2	Linear Systems under Large Measurable Disturbance	106
7.4	Case Study	109
7.4.1	System Specification	109
7.4.2	Simulation Relation Error	110
7.4.3	Uncontrolled system.	110
7.4.4	Abstraction without disturbance refinement.	111
7.4.5	Abstraction with disturbance refinement.	111
7.4.6	Controller Synthesis Process	112
7.4.7	Baseline controller.	112
7.4.8	Controller using robust simulation functions.	113
7.5	Conclusion	113
8	Assume-Guarantee Contracts for Compositional Control of Power Systems	115
8.1	Introduction	116
8.2	Original Contributions.	117
8.3	Preliminaries	118
8.4	Frequency Specifications	120
8.5	Simulation Functions	122
8.5.1	Robust Simulation Function with Disturbance Refinement	122
8.5.2	Class of Nonlinear Systems under Large Measurable Disturbance	122
8.6	Proof of Concept	126
8.6.1	Simulation Relation Error	127
8.6.2	Controller Synthesis	127
8.7	System and Specification Interconnection	128
8.7.1	Specification Composition	130
8.8	Assume Guarantee Contracts	131
8.9	Controllers for Subsystems	131
8.10	Compositionality with Internal Disturbances	132
8.11	Conclusion	134

9 Conclusion and Future Directions	141
9.1 Conclusions	141
9.2 Future Research Directions	142
A Appendix	145
A.1 NETS Matrices from Chapter 7	145
A.2 NETS Matrices from Chapter 8	146
Bibliography	151
List of Symbols and Notation	169
List of Abbreviations	173
Curriculum Vitae	175
List of Publications	177

Introduction

This thesis discusses model-based and data-driven formal synthesis of power systems. This chapter introduces the multidisciplinary areas of interest including computer science, control engineering and power systems engineering. Briefly mentioned are some of the areas under investigation that will be covered in greater depth in the rest of this thesis. An explanation on the organisation and the notation of this thesis concludes the chapter.

1.1 Motivation

Nowadays it is common to hear something akin to: "*We live in the digital age!*". Technology and computational devices have revolutionised the way we live our lives. Now digital systems are embedded in a whole host of products; in washing machines, in homes, in cars, etc. The devices are of particular benefit to us when they can do a task for us remotely or automatically. For example, one may set the heating to turn on inside our home while we are travelling so it reaches a nice temperature on our arrival. Similarly, we can let a controller in our house detect the temperature and input some controls to maintain it within some pleasant intervals.

Systems with the embedding of these discrete-time cyber components inside the continuous-time physical world are known as *cyber-physical systems* (CPSs). CPSs describe a wide range of system application domains including transportation, medical devices, chemical plants, etc. For applications where safe operation is of high-importance, these CPSs can be described as *safety-critical*. Power systems are an application area of CPS which are safety critical as any failures in the power network will have a significant impact such as blackouts causing loss of power to hospitals and airports.

To mitigate the chances of failure events, controllers are designed in order to enforce certain behaviours within the system. In the ideal scenario, a steady-state

behaviour should be maintained forever. When disturbance events occur, the system transitions away from its steady state and a controller will try to return the system back to its optimal operating conditions. In the literature, it is common to see *PID control* and *model-predictive control* (MPC) strategies used to complete these tasks. However, on their own these techniques do not provide confidence in the performance of the controllers, or guarantees that the desired behaviour will occur as expected either always or with high-probability.

Formal methods from computer science desire to verify that safety-critical applications behave as they are expected to and never cause failures. Formal specifications are written for the system which can then be validated by these approaches. Extending this approach, formal controllers can be developed which use a model of the system dynamics and a formal specification to design controllers which guarantee satisfying the specification. This approach is very powerful but comes at great cost with scalability being an issue. This scalability issue is more commonly known as the *state-space explosion* or the *curse of dimensionality*. As system dimensions grow linearly, the state space grows exponentially.

In this thesis, I look to construct these formal controllers for the power system particularly for primary frequency regulation. The power system frequency is a representation of the balance between power generation from turbine generators in the network and power consumption from loads in the demand-side of the network. The frequency is a good metric for detecting disturbances in the system.

Traditional power system control would adjust the speeds of the generators in order to keep the frequency within its normal range. However, with the electrification of the power grids and the desire for net-zero emissions, the power grids are moving away from turbine generation towards renewable generation. These renewables are not as reliable as traditional turbine generation, *e.g.* the generation may be affected greatly by weather conditions. Additionally, as people move away from fossil fuels in transportation, the demand-side of the power network will increase with more *electric vehicles* (EVs) being connected to charge their batteries.

To tackle these problems, a new 'smarter' power grid concept is developed known as a *smart grid*. Here, demand-side loads may be used in the control approach alongside the control approaches from the generation side. What this looks like in practice could be electrical loads such as EVs temporarily pausing their charging to accommodate the necessity not to overload the power network. Additionally, *prosumers* will arise which can both generate and consume energy from the smart grid, *e.g.* *active buildings*. This thesis brings together novelties from the power system community along with powerful techniques from the formal control community to make the research of this thesis incredibly exciting!

A deeper discussion of any required concepts will be provided in Chapters 2-4 which intend to give a background for each of these research areas.

1.2 Research Goals and Original Contributions

The broad aim of this thesis is to combine the results of formal control techniques from the communities of computer science and control engineering with the control techniques used in power systems. This thesis proposes symbolic control techniques for smart grids with both model-based and data-driven approaches. Case studies will be provided using EVs and *energy storage systems* (ESSs) in demand-side primary frequency response control techniques. In the following, I summarise the main contributions.

- **Conversion of Great Britain (GB) power network specification from natural language to temporal logic specification.** A detailed description of the requirements on the behaviour of the GB power network from the literature will be provided. These will be encoded to *linear temporal logic* (LTL) specifications, which can be verified using symbolic control methods.
- **Designing formal controllers for power systems with large disturbances.** Mathematical guarantees on the correctness of controllers designed over power systems with bounded large disturbances will be provided. This provides confidence that the controller will always return to a target region and never fall into any unsafe frequency ranges that may lead to contingency events such as blackouts.
- **Formal primary frequency response using distributed energy resources.** A demonstration of how formal control techniques may be applied to demand-side response in the emerging smart grids is shown; as well as how *demand-side energy resources* (DERs) can be used to provide a fast control response that aids power system stability, particularly for primary frequency control.
- **Abstraction-based controller design for unknown systems with finite numbers of data samples.** Presented is a data-driven method to compute symbolic models (*a.k.a* finite abstractions) for systems with unknown dynamics. Robust convex programs are used to overapproximate reachable sets and solve a scenario convex program to find a feasible solution with given confidence. Provided also is a lower bound on the number of trajectories required to achieve a certain confidence on the correctness of the model and the controller.
- **Designing symbolic controllers using reduced-order models.** Model-order reduction techniques will be applied to power systems and a relation between the concrete system and its (reduced order) abstract system proved. By refining an approximate simulation function with an interface function for the disturbance the simulation relation error between two systems can be reduced. The simulation relation error can then be used to reduce the regions defined as safe states or target states, and increase the regions defined as unsafe states; making the symbolic controller robust to the model-reduction step of the controller design. This approach is shown for both linear systems and a class of nonlinear systems.

- **Formal control of interconnected power systems.** A formal design approach for symbolic controllers of interconnected power systems using assume-guarantee contracts will be provided. In combination with the prior mentioned reduced-order model control approach - this enables complex large interconnected power system models to be simplified to reduced-order subsystems controlled independently. Successful control of individual areas produces confidence that the interconnected system will behave as expected.
- **Challenging power system case studies.** Throughout this thesis, challenging non-trivial power system case studies are presented to demonstrate the potential of the formal methods approaches. In particular, case studies of the GB power network and the *New England 39-Bus Test System* (NETS) are used for the design of formal controllers.

1.3 Overview of the Thesis

This thesis discusses the applications of symbolic control techniques to demand-side primary frequency response of smart grids. The thesis is organised as follows:

- **Chapter 2** introduces the concept of smart grids and why they require being controlled and managed. Active buildings are used as a lens with which to understand smart grids, particularly from the perspective of demand-side response devices and the emerging concept of the energy prosumer.
- **Chapter 3** introduces control theoretical techniques more generally. From the classical technique of PID control, more modern techniques such as MPC, and looking at future techniques such as multi-agent systems control. None of these techniques by themselves provide formal guarantees on the behaviour of the system, but some properties can be found such as stability and robustness.
- **Chapter 4** introduces formal control techniques from the formal methods community literature. This chapter discusses that systems can be related formally with simulation relations to give symbolic models. It defines specifications that can be guaranteed by these symbolic models. It provides relaxations on system relations with quantified error. Finally, it explains interconnected system control using assume-guarantee contract formulation.
- **Chapter 5** introduces the first intersection of formal control techniques for power system control in this thesis. The symbolic control approach is applied to a model of the GB power system. *Electric vehicles* (EVs) and *energy storage systems* (ESSs) are used for demand-side control to guarantee the frequency never falls below its containment zone which may trigger blackout events.

- **Chapter 6** introduces a method similar to the symbolic control approach of Chapter 5, but using data. Known as the scenario approach, a *robust convex program* (RCP) is converted to a *scenario convex program* (SCP) which needs a finite number of samples to find a solution.
- **Chapter 7** introduces *robust simulation functions* (RSF) with disturbance refinement. Here an original system model is reduced in dimension using model-order reduction techniques to an abstract system, but with guarantees on the simulation relation error that is introduced to the system by doing this. The error can then be used in the specification of the abstract system to guarantee a controller robust enough to be used on the full dimension system. The case study presented uses a linear model of the system.
- **Chapter 8** extends Chapter 7 to a class of nonlinear systems and also to interconnected systems. Each subsystem of the interconnected system uses the RSF with disturbance refinement approach to find a local controller. Using assume-guarantee contracts these local controllers can then be shown to provide guarantees over the interconnected system.
- **Chapter 9** summarises the results of this thesis and outlines directions for future research.

1.4 Publications by the Author

Most of the materials appearing in this thesis have either been published or are under submission to international conferences, journals or books. The works in which I am the first author, I have written in their entirety. Co-authors have assisted by providing comments on how to improve the document, or assistance with code. The words themselves are my own and so these documents are presented with sections directly taken from those texts. Adjustments have been made to improve the flow of this thesis. The connections between the chapters and the publications are as follows

- **Chapter 2** is based on [210]. The cited work is an Elsevier published book chapter and includes a case study combining stochastic formal control of thermostatically controlled loads with ESSs which is beyond the scope of this thesis.
- **Chapter 3** is based on [208], which is an Elsevier published book chapter and includes examples of control methods using PID, MPC and multi-agent systems.
- **Chapter 4** gives a summary of the widely known techniques established in the formal control literature. The works [19, 25, 188] are of particular note relating to this chapter.
- **Chapter 5** is based on [209]. This work has been published as part of the *Smart Energy Systems and Technologies* (SEST), 2020, conference proceedings, and the ESS case study is from [210].

- **Chapter 6** is based on [86]. This work is under review in an international journal. This was a joint work of which I contributed equally with Milad Kazemi and Mahmoud Salamati. Significant amendments have been made to the original text to present it in a way that aligns with the rest of the thesis, particularly regarding notation, style and flow.
- **Chapter 7** is based on [207]. This work has been published as part of the *European Control Conference (ECC), 2023*, conference proceedings.
- **Chapter 8** is based on [206]. This work is under submission to an international journal.

Other works completed or under submission which do not relate to the thesis topic can be found in the **List of Publications**.

1.5 Notation

In general, notation will be explained as it appears or at the start of a chapter. The reader can skip past the notation and refer back to it when it comes up. To begin, provided is some notation that will remain consistent through the thesis.

Notation. I denote the set of natural, real, positive real, and non-negative real numbers by \mathbb{N} , \mathbb{R} , $\mathbb{R}_{>0}$, and $\mathbb{R}_{\geq 0}$, respectively. The set of natural numbers including zero is denoted by $\mathbb{N}_{\geq 0}$. I use superscript $n > 0$ with these sets to denote the Cartesian product of n copies of these sets. The power set of a set A is denoted by 2^A and includes all the subsets of A . The empty set is denoted by \emptyset . For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$, and a relational symbol $\triangleright \in \{\leq, <, =, >, \geq\}$, I write $\mathbf{x} \triangleright \mathbf{y}$ if $\mathbf{x}_i \triangleright \mathbf{y}_i$ for every $i \in \{1, 2, \dots, n\}$. A matrix $M \in \mathbb{R}^{n \times n}$ is said to be non-negative if all of its entries are non-negative. I use the operator $|\cdot|$ to denote the element-wise absolute value. Symbol \mathbb{I}^n is the identity matrix in $\mathbb{R}^{n \times n}$ and $a \ll b$ represents a much less than b . I use f for system frequency and therefore the most common functions are defined with g and h instead. A function $\varrho : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a class- κ function if ϱ is continuous, strictly increasing and $\varrho(0) = 0$. All derivatives are taken with respect to time, additionally, notation often omits time for simplicity (e.g., $\mathbf{x}(t) \rightarrow \mathbf{x}$). Given functions $g^i : X^i \rightarrow Y^i$, for any $i \in \{1, \dots, N\}$, their Cartesian product $\prod_{i=1}^N g^i : \prod_{i=1}^N X^i \rightarrow \prod_{i=1}^N Y^i$ is defined as $(\prod_{i=1}^N g^i)(x^1, \dots, x^n) = [g^1(x^1); \dots; g^N(x^N)]$. I represent systems with Σ , where superscripts are used to label subsystems (i.e., Σ^i) and subscripts to represent an abstract system that could be a reduced-order model; i.e., original system (Σ_1) and its reduced-order abstract system (Σ_2). A symbolic model (or finite-abstraction) of Σ is denoted by $\hat{\Sigma}$.

Smart Grid: Control and Management

Control and management of smart grids is a growing research area that affects the global goals of net-zero emissions, increased *renewable energy* generation, and efficient energy management. *Active buildings* are a building block of future smart grids and will be the lens used in this chapter to get an insight into high-level smart grid control approaches. The passive role of buildings as energy consumers is extended in the smart grid paradigm. These *active components* now not only consume energy but also provide energy services to the network or neighbouring areas in times of need. Using active buildings as its lens, this chapter covers aggregation, frequency and voltage regulation as well as security considerations. The discussions in this chapter are based on the work [210].

2.1 Introduction

The power network is becoming increasingly intermittent as the contribution from renewable energy generation rises. To maintain stability and functionality of the power network, storage of renewable energy and demand-side control techniques are required. *Smart grids* provide the communication infrastructure to accomplish this goal. Smart grid control originated from the idea that the demand-side of the power grid can shift or shed load to reduce the strain on the network, while also maintaining consumer satisfaction and other specialist requirements [157]. The main benefit of a *smart city*, is to help its citizens by making city-related decisions. A smart grid differs from a smart city since its communication network revolves around optimising the power network, while a smart city also considers other city-level features alongside energy provision. Both smart grids and smart cities benefit from timely and relevant information transfer.

To give a supporting example, consider the energy management system of a country. The power usage of each house in a city is measured regularly, and an *aggregate model* can be created to show the energy demands of different cities. The flow of power to those cities is also calculated, and expected needs are considered. From this data, an energy model is formed, and the power system can be controlled so the needs of each city are met. The population of the city is usually proportional to the city's energy usage. Services such as hospitals, schools, factories and universities, are also more frequent for higher population numbers. City centres and industrial zones will use more resources than residential areas. Overall, a control scheme needs to consider a complex multi-objective problem consisting of these different features when being designed. One control could be the use of higher prices in areas that use more energy, with the intent to bring their usage down [151]. These energy management concepts will be looked at in greater depth in the following sections.

Buildings are seen traditionally as passive loads in the power system. Their users consume electricity, gas and water and pay for these services via predefined tariffs. The recent evolution in smart devices and the integration of *distributed energy resources* (DERs), including storage systems, enables buildings to transition toward a responsive player in energy systems. Such buildings are called *Active Buildings*.

The future of active buildings is complex. The buildings become smarter and are coupled to complex DER components. The buildings can generate their own energy using *photovoltaic panels* (PVs) and store that energy in *energy storage systems* (ESSs) for delayed consumption. An active building can charge a plug-in electric vehicle (EV) it connects to and manage building temperature through thermostatically controlled loads (TCLs). *Building management systems* (BMS) are able to control these DERs to optimise a building's energy efficiency, even to become self-sufficient away from the power grid.

As the demand-side of the power grid contains more DERs, it also becomes smarter. The smart grid, is able to control its frequency and voltage using *demand-side response* (DSR) services using these DERs. In this way active buildings are no longer purely consumers of energy from the smart grid but also generating power that can be injected back into the power network or sharing the energy with neighbours. Hence, they become **producer-consumers**, or *prosumers* for short. The next chapter of this thesis will discuss coordination structures and specific control methods using active buildings. In this chapter, a detailed discussion of how active buildings can implement smart grid control schemes to provide services is presented, including aggregation, multi-vector control, and cyber security concerns. These services include frequency response and voltage support services.

2.2 Aggregate Models for a Population of Buildings

To manage energy production and consumption and to increase stability of the network, *aggregators* and *demand-side management* (DSM) are suggested for DSR. DSR is demand-side balancing of the generation and consumption of the power

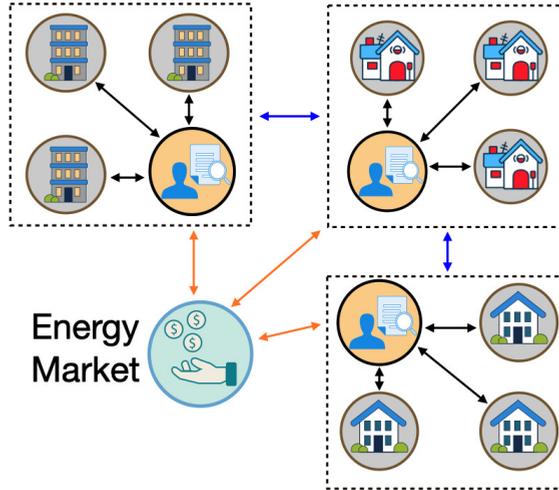


Figure 1: The aggregator is able to negotiate on behalf of the end-users it represents with the energy market and with other aggregators.

network. This is a smart grids most important function. It is vital that energy resources are carefully managed and controlled.

Currently difficulties are present with the supply-demand balance. If the consumption increases above a predicted peak value, then generation companies have to increase their energy production, or if the user consumption of energy is lower than predicted, then resources may be left unused. In both scenarios the operator will experience unnecessary additional costs, and potentially wasted resources [13].

Demand response can occur at a speed almost real-time, and can produce stabler systems with significantly reduced generation costs. In residential sectors, operators take most of the demand response benefits for themselves. Each building in a city or district level network has negligible negotiating power with the operators in the energy market. On top of this, the number of homes in a residential area contributes to scalability challenges and the operator is unlikely to manage negotiations of this scale. This combination, minimal negotiating power combined with high numbers of buildings, poses a challenge to demand response schemes [68]. A viable solution that has been developed is that of the aggregator.

The aggregator works as a middleman, between the operator and individual buildings in the residential region, see Fig. 1. The aggregator represents a group of buildings of end-users. At the energy market, the aggregator negotiates on their behalf. Since the aggregator represents a sizeable amount of demand, it is able to negotiate more effectively with the operator. This occurs predominantly with industrial buildings but is being transitioned into the residential sector as well.

The aggregator pays a fee to end-users to gain control over that building's appliances. By controlling the appliances of the buildings, the aggregator is able to

respond to peak-demand emergencies. It can turn appliances off which are using sizeable amounts of energy, such as air conditioning units. This reduces the demand-side load and returns stability to the network.

By using an aggregator, individual buildings are also able to reduce their energy costs. The aggregator is aware of the total demand profile of the buildings it represents, and pays them incentives to adjust their demand pattern. By scheduling loads to lower cost time periods, the operator is also able to save money on energy resources. The operator rewards the aggregator for its work in this service. By using this technique, both the aggregator and the end-users can have noteworthy financial benefits, despite the dominant position of the utility operator.

Demand response also has applications for EV charging schemes [170]. The aggregator can schedule when an EV charges, so as to reduce its load on the network. From a high-level perspective, the aggregator allows for collections of power units to be treated as a single large power unit. This is because the aggregator manages the details of demand response within that group. By considering an aggregator, it is possible to simplify and scale up many models which include smart buildings, EVs, and other power units.

2.3 Frequency Control

Frequency instability is the inability of the power system to maintain a global frequency value after a disturbance occurs, where the disturbance triggers rotor speed changes in the generators. The frequency may rise or fall out of its given operational range. Usually, instability is caused by an imbalance between the active power generation and the loads present in the system. Poor coordination, poor protective equipment, slow response times and lack of generation reserves are all possible reasons why instability would occur after a disturbance [28].

One of the control techniques to face the load and generation imbalance is *peak load shifting*. Over a 24 hour period, different hours of the day experience different load values. For example, during the day the system load is highest because people are awake and working. Large machines are also in operation which require large sums of power. At night, the opposite would be the case, people are asleep and machines are turned off. Load shifting attempts to move some of the demand from peak hours to off-peak hours, such as EVs which could charge overnight when their owners would be sleeping. By rescheduling the load away from peak hours, the amount of active power generation required at peak times decreases, providing multiple economic, environmental and efficiency benefits [44].

As the number of EVs connected to the electricity grid is expected to increase. A large load will be present within the system. The work [211], uses residential homes with EVs to reduce the expected power system load of higher EV penetrations. A smart charging method is proposed which utilises aggregators to schedule the EV charging times. The process is able to update, using a method similar to *model-predictive control* (MPC), depending on user travel requirements and conditions of the power system. By reducing annual peak loads, the algorithm is able

to reduce the charging costs for the EV owners and also defer system upgrades that would be required for the system with the increased load.

2.4 System Inertia

Another important consideration is the *rate of change of frequency* (RoCoF). RoCoF is inversely proportional to the available inertia in the power system. When there is low inertia, RoCoF is high which means that system is more volatile and large disturbances could destabilise the system, causing load shedding [47]. Renewable energy sources like PVs do not provide any inertia to the system [199]. This is unlike the large synchronous generators which have kinetic energy stored in their rotated mass that can be released. Therefore, with a global desire to decarbonise the energy systems and accelerate the transition to renewable based energy systems, RoCoF will increase unless managed.

One developing solution is *virtual synchronous generators* (VSGs), which provide the grid with a variable virtual inertia. The VSGs consist of short-term storage with an inverter and an efficient control mechanism. The VSG reproduces the dynamic properties of real synchronous generators and keep the advantages they provided which include adjusting active and reactive power [29].

Fixed speed wind turbines use an induction wind turbine to provide inertia response for frequency deviations, although the inertia provided is small compared to the inertia of a synchronous generator. For variable speed wind turbines, permanent magnet synchronous generators are used to provide power to the grid, but these are decoupled and so do not provide any frequency response services as well as not storing any reserve power. Techniques are developing that can release the kinetic energy stored in the rotating blades within 10 seconds to help with frequency response, this is known as *inertia emulation* and uses a *fast power reserve* [47].

For PVs, a deloading technique is used which increases the PV voltage beyond the *maximum power point* (MPP) voltage and allows the PV to retain some reserve power [215]. When the system frequency deviates, this reserve power can be released. However, all PVs will release the same value of power into the grid, and those with less reserves will reach MPP faster and stop contributing to the frequency response, this may cause a second frequency drop once the reserve power is used up.

Another technique to support renewable sources for frequency response is ESS devices, the ESS provides active power to try and prevent a second frequency drop in the system, but also acts as a backup system to provide power if there are any deficits. Coupling ESS with wind or solar can help alleviate the risks of high RoCoF [47, 69]. The work [197], shows that for low inertia systems, that inertia has little effect on the frequency probability distribution function for small disturbances. So, virtual inertia is insufficient on its own to keep the frequency near the nominal value. Instead, it was shown that the aggregate droop and deadband are the only parameters that have a major influence over the average frequency

deviations, which suggests that energy storage solutions are viable and valuable for future smart grid frequency response services.

2.5 Coordinated Volt/Var Control

With the increase in renewable penetration in the *transmission network*, and increased DERs in the *distribution network*, voltage control is facing significant challenges. On top of this the transmission network and the distribution network are increasingly coupled, meaning both sides need to be coordinated effectively for the voltage control of both networks. These challenges require a greater understanding of the problems faced and the development of advanced control techniques.

Three main control techniques exist: restoring and maintaining the system within a safe working region is known as *corrective control*, *Coordinated control* tracks a set value and mitigates disturbances, and *Preventative control* tries to fix the system before any instability occurs. Preventative control may use *fast-response dynamic reactive power* (DRP) to facilitate this. Choosing the correct control scheme depends on response times, control costs and control effectiveness [158].

2.5.1 Transmission Network

Voltage control on the transmission network has been successful for a long time, however the power network is evolving to include more renewable resources. High penetrations of wind or solar generation have a high risk of causing voltage fluctuations. The use of HVDC lines to connect onshore and offshore wind farms to the transmission network also has the potential to cause cascading failures, e.g. due to DC-blocking contingency. When the fault occurs, in 2-3 seconds the voltage will increase significantly. This may lead to the curtailment of renewable sources. For this reason, preventative control methods should be applied to wind and solar plants using DRPs for fast-response. HVDC lines have some fast-response reactive power regulation capabilities and so can also support this control scheme [91, 127].

Short-term voltage-instability is usually linked to induction motor stalling, also known as *Fault Induced Delayed Voltage Recovery* (FIDVR) issues [145]. To mitigate FIDVR, dynamic reactive power sources including *static reactive power compensators* (SVCs), *static synchronous compensators* (STATCOMs) and DERs are used. The control challenge is the optimal sizing and placement of these sources. Both aspects require the computation of the post-fault voltage trajectory which requires solving *differential algebraic equations* (DAEs). These equations are complex due to the nonlinearity of the input-state behaviour and of the solution space, which is also nonconvex. For solving the optimal placement, the selection of contingencies must be considered. Depending on the number of reactive power sources, the factors can cause a heavy computational burden. Usually solutions are found using *mixed-integer programming* and *heuristic optimisation* algorithms.

For long-term voltage instability detection, the *Thévenin equivalent* is one method used to present a model of power supply with a fixed impedance to a voltage source for analysis. When the load impedance magnitude equals the Thévenin equivalent impedance then the maximum power supply occurs. This is known as the impedance matching principle and from it the voltage stability margin can be computed. When using measurement data there are challenges from load variations and measurement noise. SCADA and *phasor measurement unit* (PMU) data are used for these computations. There are three methods which are valuable for both online and offline estimation. These are

1. using the least mean squares to fit the data to the power-voltage curve of the Thévenin equivalent,
2. using data and a real-time algorithm to estimate the Thévenin equivalent from a given bus,
3. estimation of Thévenin reactance to compute the Thévenin voltage.

Verifying and validating these measurement based approaches is an interesting open research direction [184].

2.5.2 Distribution Network

The increased penetration of DERs, has lead to stricter grid connection requirements. The interconnection of these devices has changed the grid loading pattern and influences voltage regulation device performance. In particular there are four challenges that high DER penetration has caused for voltage regulation in the distribution network.

Firstly, the *distribution system operator* (DSO) has to manage voltage rise. This is often triggered by the generation of solar energy from residential PVs [85]. The injection of too much power into the grid triggers a voltage rise in the system and to mitigate the problem the operator may impose conservative limits on PV installation. A solution is for the PV to self-regulate its voltage by reducing its active power injection or applying negative reactive power injection. Coordinating the PV inverters with demand-side management is another solution. Heat pumps, EVs, or battery storage could be coupled to the PV to increase their consumption of energy to match the PV generation injected.

Secondly, the increase in EVs in the power system could lead to overloads and large voltage drops, especially at peak times. Charging a single EV demands nearly the same energy as three houses. China, India, France and the UK have promised to phase out gas and diesel vehicles by 2040, so, with more vehicles becoming electric the load on the grid is going to increase and control methods such as load shifting and charging strategies will be necessary to avoid overloading the system [211]. Similar to the solution for voltage rise, demand-side devices can support the network to return the voltage to its nominal value. Collections of these devices can form ‘support groups’ around different network buses in the

system. A case study using this technique for control of the IEEE 24-bus reliability test system (RTS) is given in [158]. When low voltage buses are discovered in the system the support group would be tasked with returning the system voltage to its expected region.

Thirdly, the distribution network is designed for *unidirectional* power flow. *Line voltage regulators* (LVRs) are included in the network for voltage control from the load side. If a DER increases the voltage where it is situated, the LVR will try to reduce this from a load side. However, it is possible the voltage at the DER connection point would remain high. Even if the LVR could participate in *bidirectional* power flow, the DER and/or local voltage correction devices would be needed to stabilise the network. In [4], reactive power control options are coordinated to avoid the continuous operation of devices such as LVRs. This avoids both device deterioration and operation of devices at their control limit.

Finally, cloud cover variations affect PVs and can lead to voltage fluctuations and lower power quality across the distribution network. Higher deadbands and slope values would ensure system stability of the volt-var curves, but also lead to a reduced range of the voltage margin. Other solutions, such as coupling the DERs for coordinated control could be successful [118].

2.5.3 Transmission/Distributed System Operator Coordination

For voltage control it is also necessary to coordinate the control of the *transmission system operator* (TSO) and the DSO. This is to prevent reactive power exchange when the reactive power is low in the transmission network. A distribution network with coordinated DERs can flexibly adjust its reactive power consumption to provide power reserves and improve the transmission network voltage. There are two types of coordination of the TSO and DSO; rule-based methods, and distributed optimisation. Both methods involve exchanging boundary voltage and power with one another to stabilise both the transmission network and the distribution network. The distribution network can help by providing fast-response power injection into the transmission network or by reducing the demand and increasing local generation. If the distribution network control scheme ignores the transmission network conditions, it is possible that long-term instability will result. To analyse the transmission-distribution reactive power support, a co-simulation of the transmission-distribution is needed. Effective coupling of the transmission networks and distribution networks leads to economic benefits, as a transmission network does not need to invest in voltage control devices, and can leverage the DERs of the distribution networks [184].

2.6 Coordinated Control of Buildings as a Multi-Vector Nano Energy Hub

Since the buildings are a junction for the interconnection of multiple energy vectors including electricity, gas, water, renewable energy resources, and transport-

ation, they should be modelled and studied as a multi-vector nano energy hub. This is the reason that buildings will play an important role in the energy context. Considering the increasing global awareness about the holistic approach to energy issues [48], control and management of the buildings of the future need to be revised in such a way that they can be efficiently managed and controlled. This will accelerate the transition to decarbonisation. Due to the different dynamics and specifications in each energy vector, control of integrated coupled energy systems inside a building is a tough task.

District energy systems attempt to manage and control multiple vectors, this makes them complicated, they require a detailed understanding of both modelling and optimisation. Multi-vector control is the requirement to control modern energy networks that consist of coupled vectors. Previously, heat, electricity, water and gas have been controlled independently but this is no longer suitable for the future smart grid. For example, a heat pump may use electricity to create heat or a *combined heat and power unit* (CHP) creates electricity from heat and gas. Controlling just a single vector may actually provide lesser control strategies than the combined approach. In one example, a framework is developed that assesses *technical, economic and environmental* (TEE) benefits of integrating gas and electricity distributed networks with storage devices and discussed how a vector coupling storage system is able to increase whole energy systems efficiency [156].

Understanding the multi-vector components provides increased accuracy in the control. Solar and wind energy are both uncertain. Solar energy generation uncertainty comes from its proportionality to solar irradiance, but wind speed is the hardest energy source for modelling and prediction. When accurately modelled, a component such as CHP, which uses a heat byproduct from the electricity generation of natural gas, can increase its energy efficiency. In the case of CHP, this can be an increase from 30-40% to 80-90% [155].

Power-to-gas (P2G) is an interesting developing technology. It is used to store excess electricity from stochastic renewable power as either hydrogen or methane gas when it cannot be utilised. These gases can then be used in other areas such as fuel for hydrogen vehicles or by injection into the gas network. P2G is still largely being tested, and there is some worry about high costs and a low conversion efficiency. But a major benefit of this technology will be the large amounts of wind energy storage it can provide.

Buildings also require modelling as they will be integral to future smart grids and district energy systems. The buildings are prosumers and take part in demand-side control.

- **White box model** - based on the physical principles of the building;
- **Black box model** - based solely on data;
- **Grey box model** - hybrid approach of white and black box modelling.

Grey box models and black box models are effective for modelling many building variables. But white box models of buildings are less effective, this is particularly for the cases of real-time optimisation with time-steps less than an hour.

In [182], an *Integrated Whole Energy System* (IWES) model is used to quantify the benefits of using a multi-vector approach with regards to active buildings. Using a whole system modelling approach shows significant economic saving opportunities. The combined flexibility can increase the proportion of electricity production from renewables and reduce the reliance on low carbon generation like nuclear power. This flexibility will also work for decarbonisation and reaching carbon emission targets. As efficiency of the active building improves the total system costs reduce. Modelling active buildings as multi-vector systems allows for system complementing behaviours to be recognised. An example of this in an active building could be *thermal energy storage* (TES) which would charge thermal energy when the carbon intensity of electric heating is low and discharge it when it is high. These can improve short-term operational costs, long-term investment costs and reducing carbon emissions. In the report, cost savings were doubled when considering a multi-vector approach.

2.7 Security Aspects of Coordinated Control of Active Buildings as a Cyber-Physical System

From a high level power system perspective, the major danger of a cyber attack is if it can permeate through the network. Localised instability, or failure, should be prevented from affecting other areas of the network. Should this mitigation fail, it might trigger cascading failures across the whole network, which would have certain economic consequences as well as possible consequences to transportation (e.g. airports), healthcare (e.g. hospitals), and/or education (e.g. universities/schools).

Sensors and actuators are critical resources for power system control. These devices may connect to an *Internet of Things* (IoT). Devices of the IoT are known to have vulnerabilities, either on the device directly or through applications that connect to them e.g. a linked smartphone app. One example attack could be the *Manipulation of Demand via IoT* (MadIoT) attack. An attack intending to deliberately cause load shedding in the system, which would have huge repercussions. In [79], a discussion is given about how the attack could be resisted. The solution describes embedding protections into the operation of the transmission grid. Another proposed solution, from [187], regards *deterministic virtual networks* (DVNs) - a lightweight encryption which would provides security, privacy, performance and energy efficiency to the IoT.

In 2017, a cyber security researcher proposed a cyber attack known as the Horus scenario which targets PV panel inverters. Consider thousands of PV panels on the rooftops of European residential buildings, an attacker might send a signal which would be picked up by these PV panels, and cause them to stop storing energy. The aggregate loss of energy across the power network would then lead to load shedding schemes across the continent.

In a response from SMA Solar Technology, the low likelihood of this attack was shown. Three factors are given that show cyber attacks of this kind require sig-

nificantly large efforts when attempting to destabilise the grid network. They are; distributed regenerative power generation, decentralised production, and the heterogeneity of the PV devices and manufacturers. In essence, it was said that diversity in the grid was the important feature for its safety. An attack would need to simultaneously affect multiple distinct types of devices. Even the large-scale use of bots would have limited success as each system would require an individually configured attack profile [172].

Ultimately, the danger of any cyber attack in the smart grid is related to its communication channels. As a smart grid relies on sending information, communication specific security requirements need to be evaluated. From known security standards, the key parameters that relate to communication are as follows:

- **Confidentiality** - only those with permission should be able to read communicated information;
- **Authentication** - the true sender should be known to the receiver of the communication;
- **Integrity** - information should arrive as it was sent, without any tampering;
- **Access Control** - access to the communication network should only be available to those with the correct clearance;
- **Non-repudiation** - a sender or receiver should not be able to deny their part in a transaction;
- **Availability** - communication channels should not fail, communication should always be possible.

An attacker would benefit should they exploit any one of these principles. Providing authentication and integrity are particularly applicable to a distributed network. Distributed network nodes are fluid and may connect and disconnect from the network at any time (e.g. EV that disconnects to drive). Sender-receiver pairs are unpredictable and depend on the current network state, as well as which nodes are active. Exploiting these could lead to man-in-the-middle attacks, impersonation, message editing or forgery [158]. One defensive authentication technique is digital signatures via hashing and decryption. The receiver compares the senders signature with a known signature of that sender. If the signatures match, the receiver can be confident in the senders authenticity and that the message had not been tampered with.

Another dangerous cyber attack is the replay attack, which would listen and copy a message as its being broadcast. The attacker then sends this message again, some time in the future, attempting to use it maliciously. If a harmful message is replayed it could cause quite problematic consequences. To address this problem, the receiver needs to know that it has received the message before. Timestamps or random number sequences embedded in the message can provide a good solution.

2.8 Aspects Considered in this Thesis

This thesis mainly deals with the aspects from Section 2.2 to Section 2.3.

Throughout this whole work, aggregate models for active buildings, energy storage systems or electric vehicles are considered. In particular, Chapter 5, Chapter 7, and Chapter 8 consider aggregate collections of electric vehicles or energy storage devices as inputs to the system for the purposes of control.

Primary frequency control is the main control approach considered in this thesis. In particular, this thesis considers primary frequency control when large disturbances are present, see Chapter 5. Data-driven approaches for frequency control are shown in the case study of Chapter 6. Chapter 7 and Chapter 8 consider frequency control for large dimension power systems, and interconnected power systems.

System inertia is implicitly considered in some of these chapters too. Lower system inertia leads to larger changes in the balance between supply and demand of the power network. Therefore, large disturbances will become more common. We consider large disturbances, as previously mentioned, and so system inertia is implicitly a part of these chapters.

Other aspects such as voltage control, cyber security, and multi-vector approaches provide interesting future directions.

2.9 Conclusion

This chapter has looked at a high-level perspective of smart grids in particular looking at their operation and control. It was discussed

- how aggregators allow residential buildings to band together to provide energy injections to the smart grid and receive payment for these grid services;
- how frequency instability and voltage instability are increasingly in need of advanced control techniques as the amount of renewable generation in the network increases, including the need for methods to compensate the reduced system inertia; and
- how coordinated control methods looking at multi-vector approaches will be more effective than control methods for single vector approaches, and diversity is important for both smart grid security and control.

In the next chapter, specific control techniques will be discussed that can be used within smart grids. Again, this will be seen through the lens of active buildings.

Smart Grid: Control Techniques

This chapter will continue to discuss smart grid control and management techniques using the lens of active buildings. Techniques specific to power systems and more generally used in control theory will be discussed for the coordination structures of *centralised*, *decentralised* and *distributed* power systems. These control methods include *PID control*, *model predictive control* (MPC), and *multi-agent systems* (MAS), among others. The work in this chapter is based on the work [208].

3.1 Introduction

Smart cities and smart grids are valuable as they allow society to make intelligent and efficient decisions, with the support of regularly updating information. Control is the central principle for a smart grid and is essential for safe and efficient operation. It is a broad term that covers various processes and system objectives.

District or city level energy management involves; monitoring network growth, balancing higher numbers of energy generating resources and storage devices, and increased decentralisation due to devices becoming increasingly interconnected, e.g. via the internet. Additionally, both grid demand and energy prices are flexible - devices can connect and disconnect at any moment and prices fluctuate up and down. All of these lead to additional control complexities [154].

Another challenge is the nonlinearity of the system being controlled. Nonlinear system optimisations require complex calculations and are time consuming. The required computations most often rely on models of the control system, but building these models with significant accuracy is difficult [128]. Available techniques to tackle these challenges include system identification [115], linearisation around a fixed point [105], and model-free data-driven control [216]. But nonlinearity is still a significant challenge to overcome for most control systems.

From the security perspective, smart grid control revolves around two-way power and information flow of the communication networks. Should an attacker gain the

correct privileges, they may be able to deliberately jeopardise grid stability with induced load surges [15]. There are several reasons this attack is unlikely, simply the more distributed and diverse the network is the more security protocols and devices that would need to be exploited. But the negative use of control schemes for destabilisation should be considered due to the potential large-scale damages if such incidents happen.

Smart infrastructures, such as smart buildings in a smart city, are increasing in prevalence. Smart infrastructure is self-monitoring, self-governing and able to communicate with other aspects of the smart grid. From this, three research topics have emerged surrounding their control. Firstly, consumers have been empowered to interact with the new resources available to them and systems with multiple decision-makers have emerged. Secondly, incentives are created to enable flexible consumption, such as cheaper prices when consuming electricity at night. This is known as *transactive control* [78]. Finally, *resilient control* is a term given to protecting the system from large system failures, and does so by leveraging the communication between the smart devices. These areas each support the higher level smart grid control, but also the end-users who receive incentives and negotiating powers in the energy market [14].

Moreover, control of the smart grid enables the increased integration of renewable energies. Renewable energy generation is uncertain and intermittent, but if properly controlled can be a big step forward to the global aim of carbon neutrality. Ultimately, all the positives and negatives presented in this section hinge on the quality of the control schemes. High quality control, from accurate system models with security defences, has no negatives of note. However, poor quality control of inaccurate models with flawed security, could be more harmful than helpful. Deciding which category a system falls into is a different challenge entirely and requires appropriate research to quantitatively characterise the *quality* of a given control scheme applied to a system.

3.2 Coordination Structures for Management and Control of the Energy Systems

This section introduces the coordinating structures available to smart grids. There are two key characteristics of such coordination frameworks, these are the *system structure* and the *energy resource type*. System structure considers the configuration of the power units included within the system and their interconnections. The energy resource type may be renewable or non-renewable energy, and this has influence over which coordination framework would optimise the system performance [195].

3.2.1 Coordinated Structures from the System Perspective

The possibility of large-scale energy transfer became a reality with the rise of the industrial revolution. The process involved extracting natural resources, trans-

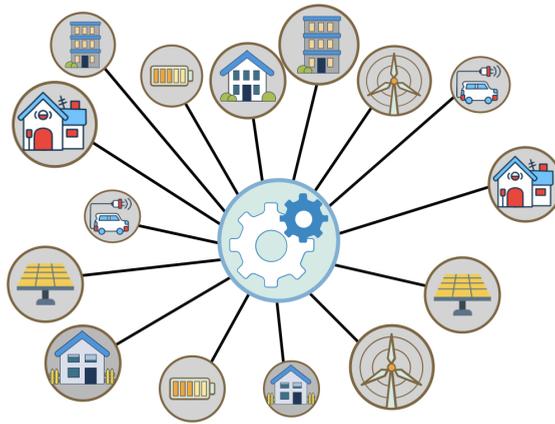


Figure 2: Centralised energy systems scheme.

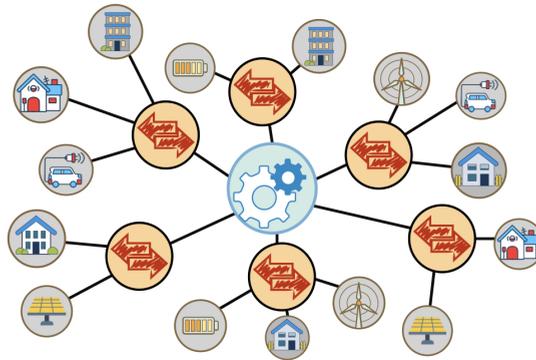


Figure 3: Decentralised energy systems scheme.

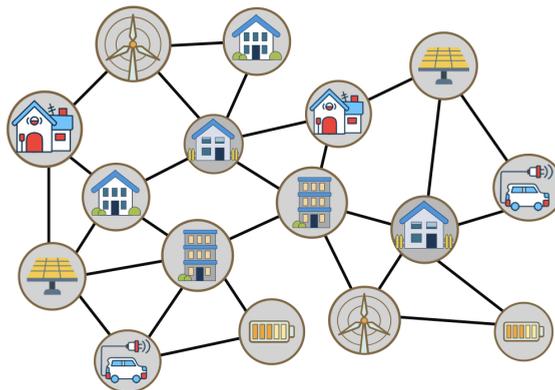


Figure 4: Distributed energy systems scheme.

porting them and then converting the resource to energy via large turbine generators. Even today, all three steps of the process are expensive and involve operations to a scale that only large corporations or bodies of government can handle. Delivering the energy to users was done through a *centralised framework*: a central large-scale generation unit sending the energy to global users far from the original generator, see Fig. 2.

As technology advances, generation units have become more readily available and at a reasonable price. This has created an opportunity for smaller companies and individuals to share in the market. A *decentralised framework* is shown in Fig. 3, where small generation units provide energy to local users and communication channels are formed between generators to share excess energy or store it.

In the decentralised and centralised approaches, the local users only consume energy from the power network. However, it is becoming more common for those users to also take part in power generation. For example, a user may decide to install a PV panel on their house to generate some electricity and reduce energy bill costs. This contributed to the development of the *distributed framework*. In this framework, local users become *prosumers*, producing and consuming energy simultaneously. A prosumer could become energy self-sufficient, but by forming a distributed network with other prosumers, any excess energy can be shared. Distributed networks can then connect to share that energy wider, see Fig. 4. As there is added complexity in the distributed framework, it is also more complex when designing control schemes compared with centralised controls [218].

The three schemes reflect the flow of energy between components in the network. They can also reflect possible communications between these components. Interpreted differently, a lack of link between two components means lack of direct energy transfer or communication between the two components. As such, all constraints should be taken into account when designing a control scheme for the network.

3.2.2 Coordinated Structures from an Energy Resource Perspective

Energy resources can be divided into two categories, renewable and non-renewable, and these influence the choice of framework to be implemented. Extraction of non-renewable fossil fuels such as oil as well as the required infrastructure for their transportation are expensive. Therefore, a large-scale centralised approach suits non-renewable generation. In contrast, a small wind turbine or photovoltaic panel does not provide much energy on its own, but the total energy can scale to large quantities of generation when these are connected within a distributed framework in large numbers. Decentralised networks, being a middle ground between centralised and distributed networks, are therefore suited to smaller generators or larger renewable generation (e.g. a solar farm).

Managing resources used for energy generation presents its own challenges. Sustainability of a particular energy resource will depend on location, legal policies and economics of a region. Wood, for example, is used in biomass, which may

become sustainable with well-planned schemes to replant uprooted trees. Uprooting trees without replanting them will lead to the complete depletion of that resource.

The use of fossil fuel energy in centralised systems is unsustainable in the long term as eventually the planet will be depleted of those resources. In contrast, renewable energy relies on sustainable energy sources (solar, wind, geothermal stones, tides, etc). As well as sustainability, countries are beginning to manage their resources with respect to carbon emissions which are significantly lower in renewable energy generation. Overall, a shift towards distributed frameworks is expected, with non-renewable resources eventually going to be depleted and net-zero emission goals forming. In the UK, legislation requires 100% carbon emission reduction relative to 1990 levels by 2050 [143].

3.2.3 Coordinated Structures from a Security Perspective

When considering security features of the frameworks, it is worth noting that should a power fault occur in a centralised or decentralised network that all or some of the network will be negatively impacted by the generation loss. A real example is from 9th August 2019. Initially triggered by a thunder strike, generation unit losses affected one million customers, health care buildings, transport and water facilities up to two days after the initial event [49]. The distributed framework is more forgiving when experiencing a fault. Each unit in the network is connected to multiple other units. A power unit can fail and rather than affecting the rest of the network it can be disconnected so the failure does not propagate through the network.

3.2.4 Supervisory Control and Data Acquisition (SCADA)

Supervisory control and data acquisition (SCADA) is a control system architecture used in industry for monitoring and control of power systems, smart grids and also power generation and transmission. It is also used in building control and for other public infrastructures such as traffic lights and water management systems. SCADA uses wired and wireless communication across four network layers. At the highest level, an operator can communicate with the process layer via the *Human-Machine Interface (HMI)*. The process layer forwards this onto logic devices; the *Remote Terminal Units (RTUs)* and *Programmable Logic Controllers (PLCs)* which can use the given information for aggregate control of field devices. At the lowest level *field devices* control and monitor the physical processes being observed by the system. These are sensors, pumps and other low-level pieces of equipment, that the observer may use to control the system as a whole. These devices provide feedback to SCADA via the HMI which helps check if the actual behaviour matches the desired behaviour [54], see Fig. 5.

As automation within industry increases and costs of operation reduces, the use of SCADA systems is expected to keep rising. However, the rise of the *Internet of Things (IoT)* has also impacted SCADA systems, and a transition from onsite

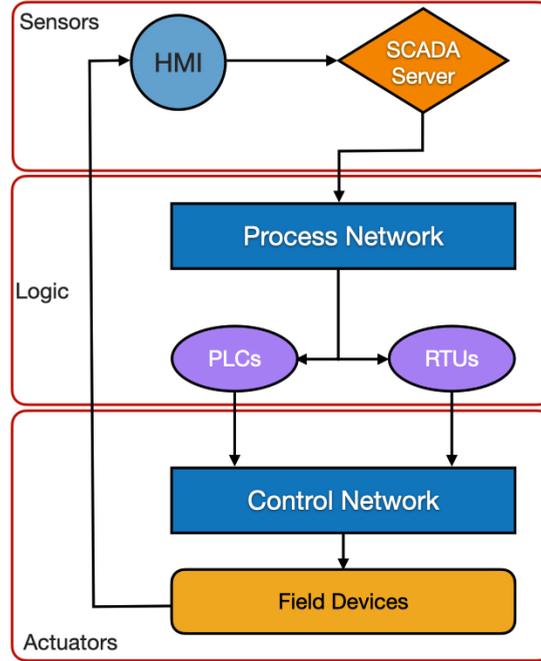


Figure 5: Control system architecture of SCADA.

and standalone systems to internet connected and remotely accessible systems is occurring. SCADA systems were never designed for network connectivity or network security. The focus was on reliability and device protection by isolating devices. SCADA is a complex system and with many inter-dependencies, so as devices go online, SCADA may become vulnerable. Another security concern, is that due to high installation costs, they remain operational for 8-15 years [141]. Relying on possible outdated or legacy systems could leave entry points to cyber-attacks.

3.3 Control Techniques for Active Buildings

The previous sections have discussed the purpose of control and systems that might need controlling. In this section, the control methods and techniques that are applicable for power systems will be introduced.

3.3.1 PID Control

The most common industrial controller choice is the *PID controller*, an acronym for proportional (*P*), integral (*I*), derivative (*D*) control [55]. This is due to the simplicity of its operation and tuning, as well as its widespread use. PID controllers use the current and previous error measurements of a system for regulation.

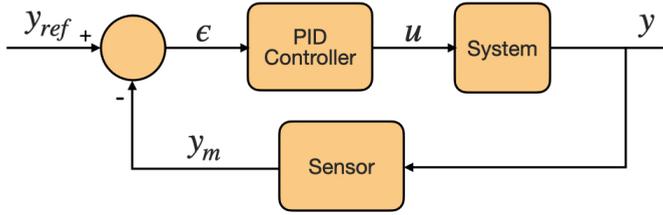


Figure 6: PID controller - the system output is measured and checked against a reference value.

The system error is the difference between a reference value of the system and the equivalent measured value. Tuning of the system can be completed by adjusting the constant values, P , I and D , of the controller, which affect each of proportional, integral and derivative areas of the control equation respectively.

In Fig. 6, an example plant, is shown with a PID controller. The system error, ϵ , is fed into the PID controller. The PID controller uses ϵ , the time step, k , and constants P , I , and D , to choose a new system input u ,

$$u(k+1) = P\epsilon(k) + I \sum_{i=1}^k \epsilon(i) + D(\epsilon(k) - \epsilon(k-1)).$$

The new input is passed to the system which produces the new output y . The values y_{ref} and y_m represent the reference output of the system and the measured output respectively.

PID control is valuable for *Single-Input Single-Output (SISO)* systems and uncoupled two input two output systems. For *Multi-Input Multi-Output (MIMO)*, other techniques such as MPC are more valuable [57, 89]. Examples of PID control in the power system literature include [53] for multi-area load frequency control, [60] for load frequency control modelling wind farms, and the survey paper [75].

3.3.2 Model Predictive Control (MPC)

Model Predictive Control (MPC), also called *Receding Horizon Control (RHC)* [30], is a control technique with some freedom involved in its implementation and is one of the fastest growing control techniques. MPC has a broad range of applications such as clinical anaesthesia, the cement industry, and robotics [34]. MPC algorithms use a model to represent a system, and attempt to minimise a cost function. Although, the implementation can be different, there are three important consistent ideas involved in MPC:

1. Use of a process model to create a prediction horizon;
2. Calculate a set of future control signals;

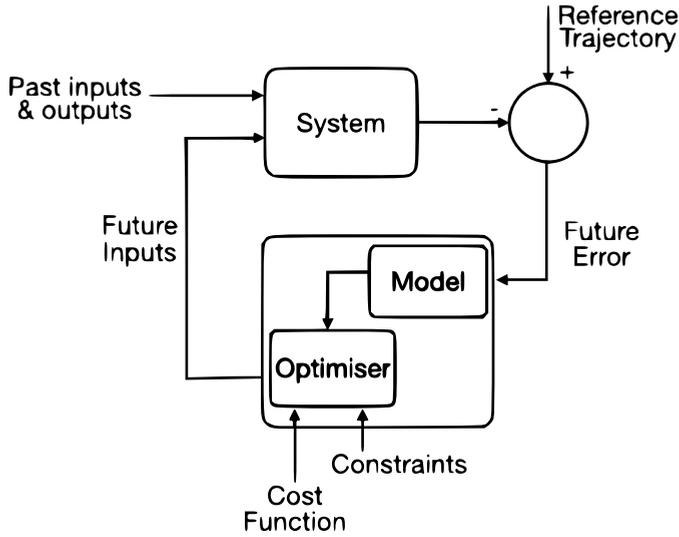


Figure 7: Model Predictive Control (MPC) - a model based control approach which predicts the next input that should be used. The model then fully updates for the next time step.

3. Use a receding strategy - the first value in the control sequence is applied to the process as it moves forward in time. At each time step a new prediction horizon and set of future control signals are calculated.

MPC uses a model to predict the future plant outputs. This is also known as the prediction horizon, $y(k+i | k)$ for $i = 1, 2, \dots, N$, where k is the current time step, and y is a set of future outputs for a time horizon, N . This is calculated using past control signals, u , measured past outputs, y_k , up until current time step, k , as well as as predicted future control signals.

The set of future control signals is represented by $u(k+i | k)$ for $i = 0, 1, \dots, N$. From this, a control sequence which minimises the objective function at each time step, $k+i$, is also known. The objective function consists of a cost function and system constraints. It attempts to keep the process near to the reference value, $y_{ref}(k+i)$. The optimisation uses the error between the reference values and the measured values.

Calculating the optimal control signal considers the system constraints and the cost function. The overall cost function consists of the sequence of manipulated values (Z_k), a cost function for tracking error (J_y), a cost function for manipulated variable tracking (J_u), a cost function for change in manipulated variables ($J_{\delta u}$) and a cost function for constraint violations (J_ϵ),

$$J(Z_k) = J_y(Z_k) + J_u(Z_k) + J_{\delta u}(Z_k) + J_\epsilon(Z_k).$$

The real-time solver computes the future sequence of manipulated variables, but only one control signal is exported for the next time step. At time k the control

signal $u(k | k)$ is applied to the process. The remaining values of the sequence are discarded. For the next time step $k + 1$, a new prediction horizon and sequence of control signals are calculated. Therefore, at time $k + 1$ the control signal $u(k + 1 | k + 1)$ is used. This is more optimal than using the value $u(k + 1 | k)$ computed in the first sequence since it considers the newly measured output value y_{k+1} which was unknown to the sequence at time step k . This is known as the receding strategy. Therefore as k increases the accuracy of MPC should improve as it has its predicted values corrected by real measurements.

MPC is advantageous as it is relatively simple to understand for those without control system experience and has applications in a wide variety of systems including unstable systems or those with complex dynamics. MPC deals easily with multi-variable processes, and can compensate for dead time. Using feedforward control it is able to compensate for uncertainties in the system.

There are two major drawbacks. The first is that the required computation in the MPC algorithm can be very heavy. Implicit MPC is run online in the microprocessor, and there is no benefit in computing the best control signal after the time step it was needed has already passed. It is possible to compute the MPC algorithm offline, and store these values in lookup tables. For explicit MPC this is done and then the tables are imported to the microprocessor [89]. The second drawback, is the reliance on an accurate process model. MPC is computed *a priori*, with prior knowledge of the process. If the model used to compute the MPC algorithm has the wrong system dynamics, then the difference in the model and the real system will create discrepancies between the predicted control and the real control of the process. Although MPC is designed to adapt to errors in the model, the more significant the errors are the more difficult it will be for MPC to correct them. Examples of MPC for frequency control include [50], the work [90] considers battery energy storage systems and the recent work [150] considers heat pumps.

3.3.3 Multi-Agent System Control (MAS)

An energy system could be alternatively described as a group of components which interact to produce a reliable service at the lowest cost to consumers. This description involves decomposing an energy system into smaller structures which interact with one another. Multi-agent system (MAS) approaches support this framework for both modelling and control, and in [157] a definition for MAS is given:

“A multi-agent system is a system composed of a collection of autonomous and interacting entities called agents, evolving in an environment where they can automatically perceive and act to satisfy their needs and objectives.”

Agents receive information from the environment through sensors. Such as a building with a PV panel on its roof might have a sensor which detects the amount

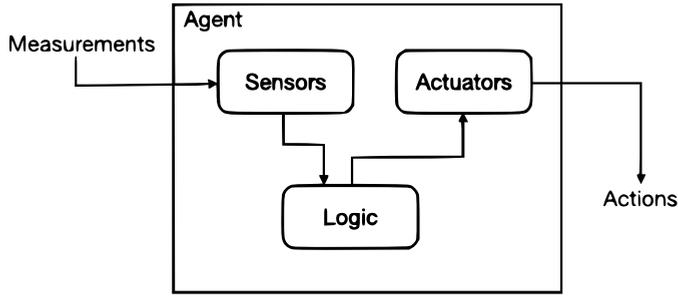


Figure 8: A model of an agent which uses measurements of the environment to impact the agent's choice of actions.

of solar irradiance currently being absorbed. This information is then passed through the agents logic which decides what action should be taken, this might be whether to sell that energy to the grid if it is in excess of the buildings own current energy requirements. This choice is passed to the actuators which complete the task so the decided action occurs within the environment, potentially changing the global state, see Fig. 8. As the agents are separated from their environment, MAS is a *distributed* approach to energy management. MAS can also be described as *proactive*; the agents follow their own objectives or objectives that cover the whole system, and *social*; the agents interact via cooperation or competition dependent on their objectives. For the MAS approach to be effective, an understanding of agent-agent and agent-environment interactions is paramount.

There are three main types of agents within the MAS framework. The first and most basic are *reactive agents* which have minimal responses to environment precepts. They provide fast responses when needed and also have useful results when modelled as a large group of simple agents, rather than a single large agent. An example of a reactive agent is an EV which connects or disconnects to the grid without using any smart charging scheme.

An *intelligent agent* has more functionality than a reactive agent since it is able to use its resources to achieve some objective. An example of an intelligent agent is a smart building with an *energy management system* (EMS), the agent uses its resources to satisfy the control objectives of its EMS.

The final agent is a *learning agent*, this agent gains information by analysing the outcome of its actions. The learning agent will have a good grasp of the environment it is in, which is used for decision making. The aforementioned smart building example, could become a learning agent by predicting the behaviour of users of the building, and using this for its resource management.

MAS control is bottom-up, due to the local knowledge of its agents and the flexible interactions they have. The agents only know what they need to know, and this reduces data transfer across the network, making MAS scalable. It is unnecessary for two agents that will never interact to be aware of one another. The agents adapt to the situations they are in, making them flexible to faults where neighbouring agents fail or when new agents are added to the neighbourhood.

Other system approaches rely on predicting network changes and the costs associated with maintenance or redesign of the network. MAS can avoid these costs. As these agents are cooperating or competing with one another autonomously the control approach required for the system is to distribute the control tasks between the agents. Decisions are made locally, with neighbouring agents grouping to form microgrids. The microgrids can then act as agents themselves to satisfy the global control objectives of the whole grid in a decentralised structure.

Difficulties of MAS revolve around the proactivity of the system. As each agent has a local objective and groups produce global objectives, it may occur that competing objectives arise. These challenges make distributed control more complex than the centralised control approach where an action would be demanded to suit the global specification. To help with this, a clear communication framework between agents is needed, as well as definitions of the roles of the agents within the network. Additionally, this benefits future hardware that will be incorporated into the networks. Examples of multi-agent systems in the literature include multi-agent reinforcement learning [161], control with electric vehicles [8], and multi-agent systems with communications constraints [114].

3.3.4 Artificial Intelligence and Data-Driven Control

Data-driven control (DDC) approaches, including *machine learning* (ML) control are growing in prominence. Having been originally developed for static data, the methods and algorithms have been shown to be equally valuable when considering dynamic systems. The broad range of techniques allows ML control to optimise both linear and nonlinear systems. The most notable techniques that will be discussed in this section are: *reinforcement learning* (RL) and *artificial neural networks* (ANN) for designing control strategies; and *genetic algorithms* (GA) and *genetic programming* (GP) for finding a parameterised controller.

A major benefit of ML, compared to other control techniques already discussed in this chapter, is the possibility for model-free control. Real-world control problems are especially difficult because they involve highly nonlinear dynamics, and an objective to maximise or minimise a certain property. System identification techniques may be impractical due to cost, complexity or other reasons (human based systems have ethical considerations). ML relies on sensor data only to optimise an objective function. It is a powerful tool when system models are unavailable. Some example real-world fields that ML techniques can help include epidemiology, robotics and fluid dynamics, but there are many more [31].

For the system that will be controlled, a cost function is minimised to satisfy the system objective. The ML controller will pass inputs to the physical system which reduce this cost value. To do this, a best strategy must be learned by the ML controller. This is completed offline, with the outputs of the system and the cost function passing to the ML controller. Using this data and differing ML algorithms a control strategy is formed. With more data the controller can gain more experience and provides better control functions.

Reinforcement Learning - Experience Based Control

RL is an algorithm that acquires experience over time to improve its control policy. *Markov decision processes (MDPs)* are the most commonly used framework for RL. MDPs incorporate uncertainty in their description of system dynamics and control laws. This promotes optimisation and exploration of the state space in equal measure.

RL is run by an agent who is in charge of choosing the control policy. Typically, the solution to an RL control strategy is binary, either the strategy is successful or it is not. To improve this, algorithms have formed a *value function*, these are known as Q-learning algorithms. This function denotes the value of success represented by the current state, and can be considered proportional to the likelihood of a winning strategy. For example, if a good control policy is in effect, the value function for the overall system state should be high. As RL learns from experience, the algorithm might initially choose low scoring strategies, but after many iterations will learn to choose higher scoring control strategies with more chance of ultimately being successful [186]. An example from the literature includes [40] for optimal primary frequency control, and also the work [185] that uses reinforcement learning to solve complex non-convex stability problems using energy storage systems.

Artificial Neural Networks - Data-Driven Control

ANN were developed from the *perceptron*, a mathematical model of a synapse in the brain. Perceptrons are quite limited, only returning a binary value, but when layered on top of one another they can learn system behaviours. ANN are easy to create and implement. As more research into ANN is done, other techniques such as *recurrent neural networks* and *deep neural networks* have been developed. In essence, ANN is used to predict an output for any given input based on all previously known inputs and outputs.

Neural networks consist of three sections, see Fig. 9. The first section is the input layer, where the system inputs are passed into the algorithm. The inputs are then passed to the hidden layer, with N-layers of neurons. The neurons in the first layer receive all the normalised inputs from the input layer and compute an output value which is then passed to each neuron in the next layer. Those neurons complete the same process, with the inputs they receive, until all layers have computed some output. The final output values are then passed to the output layer for the overall output prediction. From previous prediction errors, the ANN can adjust weights associated with each layer to improve future performance, this is known as *backpropagation* [162].

At an individual neuron the output formula is given by

$$y = \text{tansig}\left(B + \sum_{n=1}^I (i_n \times w_n)\right),$$

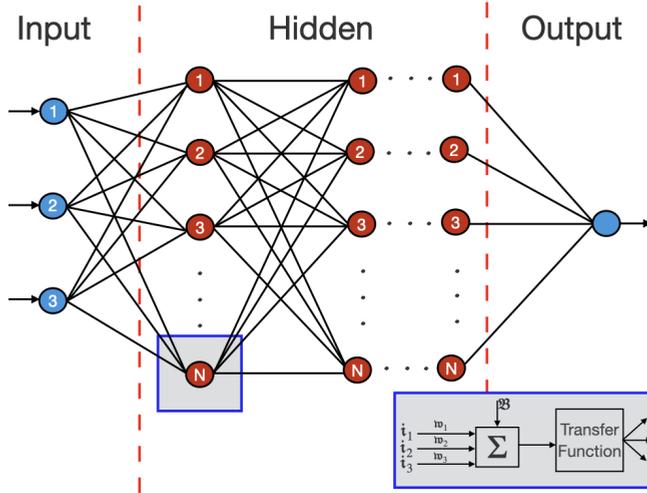


Figure 9: Artificial Neural Network with multiple hidden layers. A model representation of a neuron is also given with inputs, i , bias, B , and weights, w .

with total neuron inputs I , y is the neurons output, B is a bias value, i_n is the input and w_n is the weight associated to the respective input. The summation of inputs multiplied by their weights is calculated, then a bias value is added. This result is then passed through a transfer function such as $\text{tansig}()$, $\text{logsig}()$ or another, depending on the specific algorithm in use. In practical terms, there are many useful tools already available to help with the creation of an ANN. One example is the neural network toolbox for MATLAB [43].

To run the ANN algorithm, a large data set is required of measured input and output values. The data is divided into a training data set (majority of the data) and a validation data set (remaining data). The strength of ANN lies in its training algorithm, which decides the weights and the bias to apply to each neuron. The algorithm optimises the values of w and b for the data, using a mean square error or sum of errors in the process. Once the training has finished forming the model, the algorithm can predict the expected output values of a given input.

Occasionally, the ANN model focuses on smaller details present in the training set over the general trends of the input to output mapping. This is known as *overfitting* and to avoid this, validation data is used. The model is run using the inputs from the validation data, but the output values remain hidden. The model predicts outputs from the validation data that can be checked against the actual outputs from that data which had been hidden. If there is a significant error in these predictions then it is likely that overfitting has occurred within the model, and the training should be redone. Examples in the literature include [135] and [7].

Genetic Algorithm - Parameter Synthesis

When the structure of a control law is given but the parameters are unknown, ML control becomes parameter identification. Genetic algorithms (GAs) are meta-heuristic optimisation approaches based on the laws of natural selection from evolutionary biology, and also have applications outside of ML control. With each new generation, the most successful (or fittest) individuals climb to the top of the rankings. An individual is a member of a generation, k , with a random assignment of parameter values. This combination of parameters is called its DNA, and is written as a numeric sequence. These parameters are what the algorithm tries to optimise using a set of *genetic rules*.

The fittest individuals in a generation have minimised their cost function or error scores. Once a generation has been computed, all the individuals are ordered by cost function value. A probability is assigned to each individual, relative to this cost value. A lower cost will equate to a higher probability of being selected for the next generation, $k + 1$. Optionally, a chosen number of individuals can be immediately moved to $k + 1$, with probability 1. This is known as *elitism* and for the remaining individuals, any who are selected for generation $k + 1$ will undergo one of three genetic rules:

1. *Replication* - The individual is moved to generation $k + 1$ immediately with no modifications.
2. *Crossover* - Two individuals swap a portion of their DNA, then both move to generation $k + 1$ in their new modified forms.
3. *Mutation* - An individual has its DNA randomly modified before moving to generation $k + 1$.

The remaining individuals are then discarded. As the GA iterates through the generations, the fittest individuals with lowest cost function scores appear. GA stops when these individuals have converged or a stopping condition is met.

GA is useful as it does not require the iterations of a brute force algorithm, trying every possible set of parameters, and it scales to high dimension spaces better than other algorithms, such as Monte Carlo sampling. However, GA does not guarantee converging on an optimal solution. Additionally, choosing the size of the population and the number of generations affect the algorithms performance. An example in the literature includes [42].

Genetic Programming - Control Law Identification

GP is an extension of GA, but can be treated as its own control technique. It is used to optimise the parameters of the system and even the structure of the system. GP can also find appropriate control laws. It does this by completing a GA approach for different structures as well as different parameters for those structures.

GP uses a recursive tree structure to encode the complex functions of sensor signals. By using such a generalised framework it is possible to identify the structures of nonlinear control laws for highly nonlinear systems. GP works especially well on problems where the solutions can be checked quickly, this allows GP to test a large number of individual control laws for their suitability. A successful GP approach will find not just the optimal parameters for the model but also the optimal model to use. Example from the literature include [93] and [149].

3.3.5 Game Theoretic Approaches

Game theory is an optimisation method with mathematical foundations. Entities, or players, attempt to achieve individual objectives and take actions to complete them. Deciding on the best actions to take, and the outcome of those actions, is the essence of game theory. A player's utility is the value of their success in the game, relative to the other players. Depending on the strategy, a player's utility can increase or decrease. Game theoretic approaches usually try and manage the utility of the different players to satisfy an overall global objective. There are two main areas of game theory; noncooperative and cooperative approaches.

Under these two main umbrellas, games can also be either dynamic or static. Dynamic games consider time when playing, and thus allow for consecutive player moves. Static games on the other hand do not consider time, and so player moves are either simultaneous or alternating.

In *noncooperative game theory*, players cannot communicate with one another and so must choose actions without coordinating their choices [24]. The global solution to the game often takes the form of a *Nash equilibrium*. A Nash equilibrium is a state where no player can improve their utility. Essentially, this is a draw between all players, and changing the strategy will only result in worsened performance for the player who changes strategy. Finding such an equilibrium is not always simple and certainly not guaranteed [159].

Cooperative game theory allows communications between the players. Since the players can communicate, they have the option of whether they wish to cooperate with one another or not. This leads to two main cooperative strategies. Firstly, the *Nash bargaining strategy* where the players communicate with one another to determine a contract under which they agree to cooperate. This strategy allows for competition, but with some agreed trading. Secondly, the *coalition strategy* where the players group together as one coalition. This strategy is fully cooperative, and the players unite under the same objective. Once it is clear which type of game is being played by the players, the in-game strategies can be considered [41].

There are three essential parts to a player's turn, no matter which game they are under. The player must consider the global game state to understand the current utility and set of actions that player has. The player must estimate their prospective utility for their actions. The player updates their strategy based on those observations and chooses an action. There is some variety regarding the algorithms which are used to complete these three steps:

- *Best response dynamics* - simplest approach, chooses the action which maximises the players utility but does not guarantee convergence to an equilibrium point.
- *Fictitious play* - considers the actions of all players before choosing an action. For zero-sum games, it will always converge to a Nash equilibrium.
- *Regret matching* - a strategy which chooses the least detrimental action, as opposed to the most beneficial action.

There are also other algorithms such as reinforcement learning and stochastic learning that can be chosen [171].

An assumption from classical game theory is that the players are rational, in reality this is not necessarily the case. Small changes from the optimal strategy via non-optimal play could have disastrous knock-on consequences. To combat this there are analytical techniques to avoid such faults from occurring. Overall, the robustness of the algorithm design and model is essential to safe operations.

Game theory approaches have many power system applications. Some examples include; cooperative energy exchange, distributed control of microgrids and smart grid communication technologies. Classical approaches of game theory for demand-side management focus on the relationship between individual buildings and the operator and the optimal strategy for buildings is to use an aggregate behaviour when negotiating with the operator [164]. Examples from the literature include [58,59,168].

3.4 Why consider formal methods?

Previous results on frequency regulation rely on simulations and do not provide any formal guarantees ensuring the desired behaviour of the frequency over time. This particularly holds when comparing with classical control techniques which provide no guarantees or safety certification. For safety-critical systems such as power systems, proving strong guarantees on system behaviour mathematically would be of significant value, although challenging. I discuss the details of formal control techniques in the next chapter.

Examples of applying formal methods to smart grids in the literature include a symbolic controller design for time-varying DC microgrids [218]. The works [84, 174, 178, 180, 213] propose formal modelling and synthesis approaches for demand response of thermostatically controlled loads and microgrids. In [72], formal software engineering has potential to be applied to the smart grid domain such as using a refrigerator for active power. Formal software engineering techniques are used for self-healing smart grids in [94] and in [6] the smart grid components are formally described in Z , the formal specification language. Formal techniques for smart grid power line communication are discussed in [192].

3.5 Conclusion

This chapter looked at the current state-of-the-art control techniques used in smart grids. Discussed were

- the conventional control and monitoring method SCADA and its uses in many public infrastructures as well as smart grids;
- control approaches that use a reference tracking approach including PID control and MPC;
- different types of agents and their interactions in the MAS control framework;
- data-driven control approaches including neural networks, reinforcement learning and genetic algorithms/programming; and
- game theoretical approaches and their uses for cooperative or non-cooperative optimisation for energy exchange or pricing.

The next chapter will look at *formal control approaches* and provide a background of systems, specifications and the techniques to be applied in order to have rigorous guarantees of system behaviour and safety certification. The next chapter will form the foundation for the theory used in the later chapters.

Formal Control Techniques

This chapter provides a background to the main theoretical contributions of this thesis. Firstly, a rigorous description of a system is provided that describes its characteristics and behaviours. Then system relationships are defined, including simulation and bisimulation relations. Symbolic models, system specifications and formal controller synthesis are introduced, these form the key foundation of the contributions in all following chapters. Additional discussions are provided for approximate simulation relations, system composition and assume-guarantee contracts which are used particularly in the later chapters.

4.1 Introduction

A *transition system*, or simply *system*, can be described as a mathematical model of a dynamical phenomenon. Different models of the same phenomenon can be used in different tasks, and relationships between those systems can be described. The work [188], considers *infinite-state systems* described using *differential equations* due to *continuous-time dynamics* that are *deterministic* in nature. *Formal controller synthesis* will use *symbolic models* of these systems to provide the necessary guarantees described by the *formal specification*. Effort will be made to use the adjectives "infinite" or "finite" when describing state spaces, and "discrete" or "continuous" when discussing time throughout this thesis, but it is not uncommon to also hear state spaces described as continuous or discrete within the literature.

4.2 Describing Systems

Consider the following general definition of a system:

Definition 4.1 (System). A system Σ is a septuple (X, X_0, U, V, g, Y, h) , consisting of:

- $X \subseteq \mathbb{R}^n$ is the set of states;
- $X_0 \subseteq X$ is the set of initial states;
- $U \subseteq \mathbb{R}^p$ is the set of control inputs;
- $V \subseteq \mathbb{R}^q$ is the set of external disturbances;
- $g : X \times U \times V \rightarrow 2^X$ is a transition relation describing the evolution of the system;
- $Y \subseteq \mathbb{R}^m$ is the set of external outputs, or observations;
- $h : X \rightarrow Y$ is the external output map.

The evolution of the system can be characterised by

$$\Sigma: \begin{cases} \mathbf{x}' \in g(\mathbf{x}, \mathbf{u}, \mathbf{v}), \\ \mathbf{y} = h(\mathbf{x}). \end{cases} \quad (4.1)$$

where $\mathbf{x}, \mathbf{x}' \in X, \mathbf{y} \in Y, \mathbf{u} \in U$, and $\mathbf{v} \in V$, where \mathbf{v} is measurable and potentially large.

Generally $\mathbf{x} \in X$ are considered to be internal to the system (hidden) while $\mathbf{y} \in Y$ are external (visible). For control purposes, consider the initial set of states as either the full state space X or some subset of X dependent on any regions of the state space deemed unsafe or that would cause the system to violate the formal specification. The evolution of the system is captured in the transition map where $(\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{x}') \in g$ would describe a transition from *predecessor* state \mathbf{x} under input \mathbf{u} and external disturbance \mathbf{v} to some *successor* state \mathbf{x}' . The system definition can be simplified under several conditions:

- if $X_0 = X$ then: $\Sigma = (X, U, V, g, Y, h)$,
- if $Y = X$ and $h = \mathbb{I}$ then: $\Sigma = (X, X_0, U, V, g)$,
- if $V = \emptyset$ then: $\Sigma = (X, X_0, U, g, Y, h)$
- any combination of the above, with the most reduced system definition being: $X_0 = X, Y = X, h = \mathbb{I}, U = \emptyset$, and $V = \emptyset$, then: $\Sigma = (X, g)$.

Definition 4.2 (Determinism). *A system Σ is deterministic if for any state \mathbf{x} , any input \mathbf{u} , and any disturbance \mathbf{v} , there is **at most** one successor state:*

$$(\mathbf{x} \xrightarrow{(\mathbf{u}, \mathbf{v})} \mathbf{x}') \wedge (\mathbf{x} \xrightarrow{(\mathbf{u}, \mathbf{v})} \mathbf{x}'') \Rightarrow (\mathbf{x}' = \mathbf{x}''), \text{ for all } \mathbf{x} \in X, \mathbf{u} \in U, \mathbf{v} \in V.$$

For any input \mathbf{u} , any disturbance \mathbf{v} and any state \mathbf{x} , if it is possible to have two or more distinct successor states e.g., a distribution of successor states, then the system is *non-deterministic*. Different control approaches are needed depending on the system's determinism. To describe these principles graphically, the most common approach is to use circles for system states with a transition between states represented by an arrow. A notation is added on top of the arrows which

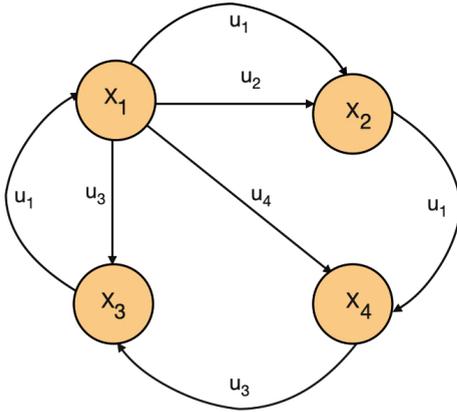


Figure 10: Determinism - for each input u_i applied to a state x_i , a state transition is represented by an arrow. There is at **most** one successor state for each input-state pair.

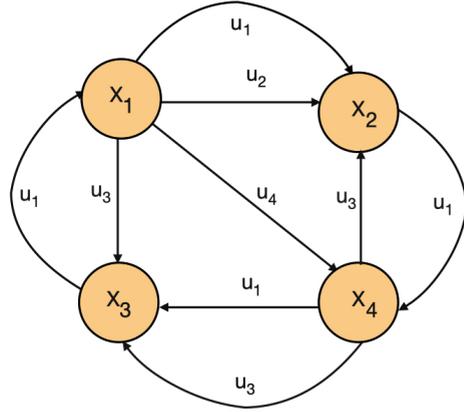


Figure 11: Non-determinism - for each input u_i applied to a state x_i , a state transition is represented by an arrow. There may be two or more successor states for an input-state pair.

shows the input required for the state to undergo the transition. In Fig. 10 a deterministic system is provided and Fig. 11 shows a non-deterministic system. The disturbance v is omitted for simplicity, but usually uncertainty in the form of noise or unknown disturbances cause the non-determinism.

For a state x and any input $u \in U$, denote the set of successor states by $Post_u(x)$ and the set $U(x)$ as the set of inputs $u \in U$ where $Post_u(x) \neq \emptyset$. The systems considered in this thesis are assumed to be *non-blocking* meaning that for every state x the set $U(x) \neq \emptyset$.

4.3 System Behaviour

As stated before, the system states and transitions between the system states, known as *paths* or *trajectories*, are considered to be *infinite internal behaviours* of the system, indicated by ρ_x .

$$\rho_x = x_0 \xrightarrow{(u_0, v_0)} x_1 \xrightarrow{(u_1, v_1)} x_2 \xrightarrow{(u_2, v_2)} x_3 \dots$$

Therefore the system's behaviour may not be immediately obvious when designing a controller, except when the set of outputs matches the set of states, $Y = X$. Instead, system behaviour can be described using the *infinite external behaviour*, written ρ_y , where depending on the system design, the system may be aware of the input selected for the transition and perhaps an upper bound on the anticipated disturbance;

$$\rho_y = y_0 \xrightarrow{(u_0, \bar{v})} y_1 \xrightarrow{(u_1, \bar{v})} y_2 \xrightarrow{(u_2, \bar{v})} y_3 \dots$$

with $h(\mathbf{x}_i) = \mathbf{y}_i, \forall i \in \mathbb{N}_{\geq 0}$. The set of all infinite external behaviours *initialised* by state \mathbf{x} can then be captured by the set $\mathfrak{B}_{\mathbf{x}}(\Sigma)$.

Definition 4.3 (Infinite External Behaviour). *The infinite external behaviour generated by system Σ , denoted $\mathfrak{B}(\Sigma)$, is defined by:*

$$\mathfrak{B}(\Sigma) = \bigcup_{\mathbf{x} \in X_0} \mathfrak{B}_{\mathbf{x}}(\Sigma).$$

Finite external behaviour can be defined similarly. For non-blocking systems, as considered in this work it is straightforward to prove that $\mathfrak{B}(\Sigma) \neq \emptyset$.

4.4 Exact System Relationships

For infinite-state systems, designing controllers to enforce a behaviour is difficult. Mostly this is due to the complexity of the system and the computational time and effort required to do any calculations over an infinite-state space. An *abstraction* $\hat{\Sigma}$ of a transition system Σ preserves some of its details required for analysis and control, but ignores aspects which do not influence the results [25]. This is useful because an abstraction-based controller for the desired behaviour of Σ can be designed using $\hat{\Sigma}$. However, this is only possible if an appropriate relationship can be shown to exist between the two systems.

First, consider some general relationships that can be shown between two systems, the simplest being with regard to the system behaviours.

Definition 4.4 (Behavioural Inclusion). *Given two systems Σ and $\hat{\Sigma}$ with $Y = \hat{Y}$, Σ is behaviourally included in $\hat{\Sigma}$, written $\Sigma \preceq_{\mathfrak{B}} \hat{\Sigma}$, if $\mathfrak{B}(\Sigma) \subseteq \mathfrak{B}(\hat{\Sigma})$.*

Using the definition of behavioural inclusion it can be seen that if $\Sigma \preceq_{\mathfrak{B}} \hat{\Sigma}$, then all of the behaviours of Σ are included within the behaviours of $\hat{\Sigma}$. If $\hat{\Sigma}$ satisfies the specification desired, then Σ is guaranteed to satisfy it too. However, if $\hat{\Sigma}$ does not satisfy the specification then nothing is learned about Σ . This is known as a *sound abstraction*.

There also exists a stronger relationship between systems known as a *complete abstraction*.

Definition 4.5 (Behavioural Equivalence). *Given two systems Σ and $\hat{\Sigma}$ with $Y = \hat{Y}$, then Σ is behaviourally equivalent to $\hat{\Sigma}$, written $\Sigma \cong_{\mathfrak{B}} \hat{\Sigma}$, if $\Sigma \preceq_{\mathfrak{B}} \hat{\Sigma}$ and $\hat{\Sigma} \preceq_{\mathfrak{B}} \Sigma$.*

Behavioural equivalence is stronger than behavioural inclusion because the two systems have the exact same set of behaviours, $\mathfrak{B}(\Sigma) = \mathfrak{B}(\hat{\Sigma})$. This means that if $\hat{\Sigma}$ satisfies the specification desired then Σ is guaranteed to satisfy it too. Additionally, if $\hat{\Sigma}$ violates the specification desired then Σ is proven to also violate the specification, for sound abstractions the latter guarantee would be unknown. For non-blocking systems, the infinite external behaviour equivalence implies finite external behaviour equivalence and vice-versa.

Definition 4.6 (Simulation Relation). Consider two systems Σ and $\hat{\Sigma}$ with $Y = \hat{Y}$. A relation $\mathfrak{R} \subseteq X \times \hat{X}$ is a simulation relation from Σ to $\hat{\Sigma}$ under the following conditions:

1. $\forall \mathbf{x}_0 \in X_0, \exists \hat{\mathbf{x}}_0 \in \hat{X}_0$ such that $(\mathbf{x}_0, \hat{\mathbf{x}}_0) \in \mathfrak{R}$;
2. $\forall (\mathbf{x}, \hat{\mathbf{x}}) \in \mathfrak{R}$ then $h(\mathbf{x}) = \hat{h}(\hat{\mathbf{x}})$;
3. $\forall (\mathbf{x}, \hat{\mathbf{x}}) \in \mathfrak{R}$ then $\mathbf{x} \xrightarrow{(u,v)} \mathbf{x}'$ in Σ implies the existence of $\hat{\mathbf{x}} \xrightarrow{(\hat{u}, \hat{v})} \hat{\mathbf{x}}'$ in $\hat{\Sigma}$ satisfying $(\mathbf{x}', \hat{\mathbf{x}}') \in \mathfrak{R}$.

Σ is simulated by $\hat{\Sigma}$, written $\Sigma \preceq_{\mathfrak{S}} \hat{\Sigma}$, if there exists a simulation relation \mathfrak{R} from Σ to $\hat{\Sigma}$.

For infinite-state systems, using behaviours only to define relationships between systems is challenging. New definitions are needed for stronger relationships such as the *simulation relation*, which relates the initial states, system observations and state transitions between two systems. If one can show $\Sigma \preceq_{\mathfrak{S}} \hat{\Sigma}$, then by implication $\Sigma \preceq_{\mathfrak{B}} \hat{\Sigma}$, but the reverse is not guaranteed, unless it holds that different successors of a state always have different outputs, known as *output determinism*. Similarly to before, the notion of a simulation relation can be extended to a *bisimulation relation*, written $\Sigma \cong_{\mathfrak{S}} \hat{\Sigma}$.

Definition 4.7 (Bisimulation Relation). Given two systems Σ and $\hat{\Sigma}$ with $Y = \hat{Y}$, Σ is bisimilar to $\hat{\Sigma}$, written $\Sigma \cong_{\mathfrak{S}} \hat{\Sigma}$ if both $\Sigma \preceq_{\mathfrak{S}} \hat{\Sigma}$ and $\hat{\Sigma} \preceq_{\mathfrak{S}} \Sigma$. Alternatively, Σ is bisimilar to $\hat{\Sigma}$ if there exists a relation \mathfrak{R} such that:

- \mathfrak{R} is a simulation relation from Σ to $\hat{\Sigma}$;
- \mathfrak{R}^{-1} is a simulation relation from $\hat{\Sigma}$ to Σ .

As before, it is fairly straightforward to show that $\Sigma \cong_{\mathfrak{S}} \hat{\Sigma} \implies \Sigma \cong_{\mathfrak{B}} \hat{\Sigma}$.

Further relations exist such as the *alternating simulation relation*: a stronger relation than simulation relations that relates state-input trajectories of two systems [188]. However, these only differ from simulation relations in non-deterministic systems when $Post_u(\mathbf{x}) \geq 1$, and are therefore beyond the scope of this thesis.

4.5 Symbolic Models

One tactic used for formal controller synthesis is to find a relationship between the infinite system Σ and a second finite system $\hat{\Sigma}$. $\hat{\Sigma}$ is known as the *symbolic model* or the *finite abstraction* of Σ . One of the most powerful relationships that can be acquired is the relationship between an infinite-state system and the finite-state representation of that system known as its *symbolic model* (a.k.a. *finite abstraction*, or *quotient system*). Many formal approaches have been developed to handle only finite-state systems, these approaches become more challenging when the cardinality of X increases due to the '*curse of dimensionality*', also known as the '*state-space*

explosion'. As the number of states increases, the size of the state space increases *exponentially*, which makes the problem intractable for high-dimension systems. Therefore, symbolic models are used most commonly to address this problem by reducing the size of a finite system, or map an infinite system to a finite system. Although, symbolic models reduce the size of the state space, this is different from model-order reduction that reduces the number of dimensions (covered in later sections).

Definition 4.8 (Symbolic Model). *Let $\Sigma = (X, X_0, U, V, g, Y, h)$ be a system and let Ω be an equivalence relation on X such that $(\mathbf{x}, \mathbf{x}') \in \Omega$ implies $H(\mathbf{x}) = H(\mathbf{x}')$. X/Ω denotes the set of all equivalence classes on set X . The symbolic model of Σ by Ω denoted $\hat{\Sigma}$ is the system $(\hat{X}, \hat{X}_0, \hat{U}, \hat{V}, \hat{g}, \hat{Y}, \hat{h})$ where:*

- $\hat{X} = X/\Omega$;
- $\hat{X}_0 = \{\hat{\mathbf{x}} \in \hat{X} \mid \hat{\mathbf{x}} \cap X_0 \neq \emptyset\}$;
- $\hat{U} = U$;
- $\hat{V} = V$;
- $\hat{\mathbf{x}} \xrightarrow{\mathbf{u}, \mathbf{v}} \hat{\mathbf{x}}'$ if there exists $\mathbf{x} \xrightarrow{\mathbf{u}, \mathbf{v}} \mathbf{x}'$ in Σ with $\mathbf{x} \in \hat{\mathbf{x}}$ and $\mathbf{x}' \in \hat{\mathbf{x}}'$;
- $\hat{Y} = Y$;
- $\hat{h}(\hat{\mathbf{x}}) = H(\mathbf{x})$ for some $\mathbf{x} \in \hat{\mathbf{x}}$.

The symbolic model is designed using techniques from *group theory*: a *quotient group* (also known as *factor group*) is formed where similar elements are aggregated together using an equivalence relation that preserves some of the structure of the system while also '*factoring out*' or '*abstracting away*' unnecessary parts of the structure [188]. In this thesis the most common use of symbolic models is to *grid* the state space, dividing the state space into a grid of squares with height η_x . The set of infinite states within each grid cell are then represented by one single state in the cell. This state is known as the *representative point* and is often the centre of the cell. Further details of the formulation and uses of symbolic models are discussed in Chapter 5 and Chapter 6.

4.6 System Specifications

By finding relationships between the behaviours of systems, properties of the systems can also be related such as the set of *reachable states* of the systems.

Definition 4.9 (Reachable States). *A state $\mathbf{x}_n \in X$ is reachable in system Σ if there exists an initialised finite internal behaviour where the trajectory ends at state \mathbf{x}_n :*

$$\mathbf{x}_0 \xrightarrow{(\mathbf{u}_0, \mathbf{v}_0)} \mathbf{x}_1 \xrightarrow{(\mathbf{u}_1, \mathbf{v}_1)} \dots \xrightarrow{(\mathbf{u}_{n-2}, \mathbf{v}_{n-2})} \mathbf{x}_{n-1} \xrightarrow{(\mathbf{u}_{n-1}, \mathbf{v}_{n-1})} \mathbf{x}_n.$$

An output $\mathbf{y} \in Y$ is reachable in Σ if there exists a reachable state $\mathbf{x} \in X$ where $H(\mathbf{x}) = \mathbf{y}$. The reachable set of Σ , written $\text{Reach}(\Sigma)$, is the set of all its reachable outputs.

Using the above definition for reachability it can be inferred:

$$\Sigma \preceq_{\mathfrak{B}} \hat{\Sigma} \implies \text{Reach}(\Sigma) \subseteq \text{Reach}(\hat{\Sigma}),$$

and:

$$\Sigma \cong_{\mathfrak{B}} \hat{\Sigma} \implies \text{Reach}(\Sigma) = \text{Reach}(\hat{\Sigma}).$$

The reachable set is desirable in defining two key elements of formal specifications; *reachability* and *safety*. Reachability specifications ask the question ‘*is it possible?*’, while safety specifications ask the question ‘*is it always true?*’. For reachability, some target set \mathcal{T} is defined, and the system is tested to see if it is possible for the system to reach its target. For safety, some region to avoid \mathcal{A} is defined, and the system is tested to see if it will always avoid that region. Notice then that safety can be shown as $\text{Reach}(\Sigma) \cap \mathcal{A} = \emptyset$.

If a system fails, the economic repercussions are high. Whether that is purely in lost revenue or repair costs and other aspects affected by a system failure. By designing a rigorous specification that the system holds, it is possible to provide guarantees on system behaviour.

Model checking verifies the system operation against a given *specification* [19]. The two main specifications for systems are safety and reachability. Specifications can also be combined, such as a *reach and stay specification* to reach a safe region and then remain within it, or *reach and avoid specification* to reach a safe region while avoiding unsafe regions along its path.

Formula of complex temporal logic are used to define system specifications, the two main ones are: *Linear Temporal Logic (LTL)*, which is used for specifying linear time properties, and *Computational Time Logic (CTL)*, used for branching time properties. In this thesis, LTL specifications only will be considered, as CTL specifications go beyond the scope of this work, but further details on both can be found in [19].

Here the notation Y^ω is introduced as an infinite sequence of elements in Y and similarly U^ω , V^ω , and X^ω . These sequences have been called paths or trajectories so far. Now introduced are the set of *labels* or *atomic prepositions* AP . The paths can then be mapped using a *labelling function* $L : Y \rightarrow 2^{AP}$ to a new sequence known as *words*. These notions are largely equivalent to observations y and the set of observations Y .

Definition 4.10 (LTL Syntax). *A (propositional) linear temporal logic (LTL) formula ψ over a given set of atomic prepositions AP is defined recursively as:*

$$\psi := \top \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \bigcirc\psi \mid \psi_1 \mathbf{U} \psi_2,$$

where the proposition $p \subset AP$ is of the set of atomic prepositions and ψ , ψ_1 and ψ_2 are LTL formulas.

LTL formulas can be constructed from the set of atomic prepositions, *boolean operators*, and *temporal operators*. Standard boolean operator notation is used including; *true* \top , *negation* \neg , and *conjunction* \wedge . Graphical temporal operator notation

is also used including *next* \bigcirc , and *until* \bigcup , where until is right-associative; e.g. $\psi_1 \bigcup \psi_2 \bigcup \psi_3 = \psi_1 \bigcup (\psi_2 \bigcup \psi_3)$. Additional operators can be defined from existing operators, such as; *disjunction* \vee , *implication* \implies , *eventually* \diamond , and *always* \square :

$$\begin{aligned}\psi_1 \vee \psi_2 &:= \neg(\neg\psi_1 \wedge \neg\psi_2), \\ \psi_1 \implies \psi_2 &:= \neg\psi_1 \vee \psi_2, \\ \psi_1 \iff \psi_2 &:= (\psi_1 \implies \psi_2) \wedge (\psi_2 \implies \psi_1), \\ \diamond\psi &:= \top \bigcup \psi, \\ \square\psi &:= \neg\diamond\neg\psi.\end{aligned}$$

Further powerful and often used expressions in LTL include *eventually always* $\diamond\square\psi$, and *always eventually* $\square\diamond\psi$.

A word $\rho_u = \rho_u(0), \rho_u(1), \rho_u(2), \dots \in U^\omega$ is called an input word. A word of Σ produced by the input word ρ_u and originating at $x_0 \in X$ is a word $\rho_x = \rho_x(0), \rho_x(1), \rho_x(2), \dots$ where $\rho_x(k) \in X$ and $\rho_x(k+1) \in g(\rho_x(k), \rho_u(k), \rho_v(k))$ for all $k \geq 1$. The set of transitions initialised at x is denoted $\Sigma(x)$ and the output word $\rho_y = \rho_y(0), \rho_y(1), \rho_y(2), \dots \in Y^\omega$ where $\rho_y(k) = h(\rho_x(k))$. A word ρ_y satisfying an LTL formula ψ may be written as $\rho_y \models \psi$ if $\rho_y(0) \models \psi$. The set of all words that satisfy the formula ψ are then called the *language* and denoted \mathcal{L}_ψ .

Definition 4.11 (LTL Semantics). *The satisfaction of LTL formula ψ over a set of atomic propositions AP at position $k \in \mathbb{N}_{\geq 0}$ by word $\rho_y = \rho_y(0), \rho_y(1), \rho_y(2), \dots \in L(Y^\omega)$, with labelling function L , written $\rho_y(k) \models \psi$, is defined recursively as:*

- $\rho_y(k) \models \top$;
- $\rho_y(k) \models p$ for some $p \in AP$ if $\rho_y(k) = p$;
- $\rho_y(k) \models \neg\psi$ if $\rho_y(k) \not\models \psi$;
- $\rho_y(k) \models \psi_1 \wedge \psi_2$ if $\rho_y(k) \models \psi_1$ and $\rho_y(k) \models \psi_2$;
- $\rho_y(k) \models \bigcirc\psi$ if $\rho_y(k+1) \models \psi$;
- $\rho_y(k) \models \psi_1 \bigcup \psi_2$ if
 $\exists i \in \mathbb{N} : \rho_y(k+i) \models \psi_2$, and $\forall j \in \mathbb{N} : 0 \leq j < i, \rho_y(k+j) \models \psi_1$.

Again further properties can be expressed by combining these semantics into more powerful specifications. A few examples of LTL specifications in realistic applications are as follows. A robot visits two regions p_1 and p_2 infinitely often is denoted by $\psi = (\square\diamond p_1) \wedge (\square\diamond p_2)$, commonly "always eventually p_1 and always eventually p_2 ". An autonomous vehicle overtaking a parked vehicle while maintaining a minimum safe distance can be written as $\psi = p_1 \bigcup p_2$, commonly " p_1 until p_2 ", where p_1 indicates the safe region and p_2 the target location of the vehicle after overtaking. The frequency of the power network must always remain inside the range $p_1 = [49.5, 50.2]$ Hertz at all times and return back to $p_2 = [49.8, 50]$ if it goes down to $p_3 = [49.5, 49.8]$. This specification can be written as $\psi = \square p_1 \wedge \square(p_3 \implies \diamond p_2)$, commonly "always p_1 and always if p_3 then eventually p_2 ".

4.7 Formal Control Synthesis

Formal control synthesis leverages off the previous discussions on specifications and the equivalence properties. For a system Σ , by finding an equivalent simpler system $\hat{\Sigma}$, computations for control can be done in less time. The shift to a simpler system is done via an abstraction. The system Σ is partitioned, and each partition is treated as a single *symbolic state* in a *symbolic model*. The symbolic model $\hat{\Sigma}$, is related mathematically to the system Σ in one of the ways discussed in previous sections.

The desired specification is written using LTL and is given along with $\hat{\Sigma}$ to a formal software tool. The tools attempt to synthesise a formal controller on $\hat{\Sigma}$ which guarantees the satisfaction of the specification. If a formal controller can be synthesised then the controller can be applied to Σ . From the definition of the symbolic model, the controller for $\hat{\Sigma}$ provides the same system guarantees when applied to Σ due to their equivalence relation. Examples of tools used in controller synthesis include SCOTS [163] for non-probabilistic models and MAS-COT [119,121], FAUST² [175], SReachTools [196], AMYTISS [102], StocHy [35], SySCoRe [193], and PRISM [97] for systems under stochastic uncertainty. A friendly competition to compare some formal tools runs yearly [1,2]. Examples of the employed techniques include uniform space discretisation in [102,163], model-order reductions in [73,193], adaptive and sequential gridding [35,175,177], higher-order approximations [176], and symbolic abstractions with fairness conditions [21,22,120,121].

In summary, formal controller design requires three main steps:

1. **Define the formal specification** - using LTL, the desired system specification is written in complex temporal logic. Choosing this specification can be tricky as a good understanding of the system is required in order to define wanted behaviour and unwanted behaviour;
2. **Create the symbolic model** - here the original system dynamics are converted to a symbolic representation by representing partitions of the state space as symbolic states;
3. **Synthesise symbolic controller** - the symbolic model and the specification are combined to work out legal transitions around the system. Transitions which do not lead to the satisfaction of the specification are removed from the symbolic model. Once this process has finished, any remaining state transitions form the symbolic controller. From the mathematical properties of the system such as bisimulations, the symbolic controller can be applied to the original system with formal guarantees that the specification will be satisfied.

At this point *exact* system relationships have been discussed where system output trajectories are the same, either due to behavioural inclusions and simulation relations, or through behavioural equivalence and bisimulations. They provide the

foundations for Chapter 5 and Chapter 6. In the following chapters the notation Σ will be explicitly used for the original infinite-state system and $\hat{\Sigma}$ for its finite-state symbolic model.

4.8 Approximate System Relationships

Previous sections have discussed *exact* simulation relations where $\hat{\Sigma}$ simulates Σ by describing the transitions of $\hat{\Sigma}$ in order to match the transitions of Σ and produce the same output trajectories. This is powerful when reducing an infinite-state system to a finite-state system through symbolic models but struggles in other context such as *model-order reduction*. Model-order reduction takes a large dimensional system and finds a lower-dimension approximation. Denote Σ_1 as the high-dimension original or *concrete* system and Σ_2 as the low-dimension reduced-order *abstract* system.

As system models increase in size, so do the number of system states n . Even with a vast state space, dominant structures composed of a small percentage of states that characterise the system can be found using model reduction [31]. These *reduced-order models* (ROMs) provide efficient, low-dimension systems which maintain key system input-output features. A controller designed on a ROM can be reapplied to the full system and used during real-time simulations. This may be otherwise impossible due to the computational load of high-dimension full state systems; *a.k.a* the curse of dimensionality.

One model-order reduction approach is *balanced model reduction* which reorders the states based on their *Hankel singular values*; more simply from states which are most sensitive to change (*high-energy states*) to least sensitive to change (*low-energy states*). The matrices can then be truncated to keep only the most relevant states [131].

To extend the notions of *exact* simulation relations to ROMs, *approximate* simulation relations are introduced [65]. These relax the requirement on equality of the output trajectories, instead requiring the output trajectories simply to remain close.

Definition 4.12 (Metric System). *A system Σ is said to be a metric system if the set of outputs Y is equipped with a metric $m : Y \times Y \rightarrow \mathbb{R}_{\geq 0}$.*

Definition 4.13 (Approximate Simulation Relation). *Consider two metric systems Σ_1 and Σ_2 with $Y_1 = Y_2$, and $\epsilon \geq 0$, a relation $\mathfrak{R}_\epsilon \subseteq X_1 \times X_2$ is called an approximate simulation relation of Σ_1 by Σ_2 , of precision ϵ , if for all $(\mathbf{x}_1, \mathbf{x}_2) \in \mathfrak{R}_\epsilon$:*

1. $m(h_1(\mathbf{x}_1), h_2(\mathbf{x}_2)) \leq \epsilon$
2. $\mathbf{x} \xrightarrow{(\mathbf{u}_1, \mathbf{v}_1)} \mathbf{x}'_1$ in Σ_1 implies $\mathbf{x}_2 \xrightarrow{(\mathbf{u}_2, \mathbf{v}_2)} \mathbf{x}'_2$ in Σ_2 satisfying $(\mathbf{x}'_1, \mathbf{x}'_2) \in \mathfrak{R}_\epsilon$.

Σ_2 approximately simulates Σ_1 with precision ϵ (denoted $\Sigma_1 \preceq_\epsilon \Sigma_2$) if there exists \mathfrak{R}_ϵ , an approximate simulation relation of Σ_1 by Σ_2 , of precision ϵ , such that for all $\mathbf{x}_1 \in X_{10}$

there exists $\mathbf{x}_2 \in X_{20}$ such that $(\mathbf{x}_1, \mathbf{x}_2) \in \mathfrak{R}_\epsilon$, where X_{10} and X_{20} are the set of initial states of Σ_1 and Σ_2 respectively.

When $\epsilon = 0$, the exact simulation relation $\Sigma_1 \preceq_{\mathfrak{S}} \Sigma_2$ is recovered. These notions also extend to *approximate bisimulations*.

Definition 4.14 (Approximate Bisimulation). *Given two systems Σ_1 and Σ_2 with $Y_1 = Y_2$, and $\epsilon \geq 0$, then Σ_1 is approximately bisimilar to Σ_2 , written $\Sigma_1 \cong_{\mathfrak{S}}^\epsilon \Sigma_2$ if $\Sigma_1 \preceq_{\mathfrak{S}}^\epsilon \Sigma_2$ and $\Sigma_2 \preceq_{\mathfrak{S}}^\epsilon \Sigma_1$. Alternatively, Σ_1 is approximately bisimilar to Σ_2 if there exists a relation \mathfrak{R}_ϵ such that:*

- \mathfrak{R}_ϵ is an approximate simulation relation from Σ_1 to Σ_2 ;
- $\mathfrak{R}_\epsilon^{-1}$ is an approximate simulation relation from Σ_2 to Σ_1 .

Again, when $\epsilon = 0$, the exact bisimulation relation $\Sigma_1 \cong_{\mathfrak{S}} \Sigma_2$ is recovered. The concept of the approximate relationship forms the foundations for Chapter 7 and Chapter 8. In particular, upper bounds of the metrics are computed based on the fundamental notion of simulation and bisimulation functions defined by Lyapunov like inequalities. A simulation function of Σ_1 by Σ_2 is a positive function defined on $X_1 \times X_2$, bounding the distance between the observations associated to the couple $(\mathbf{x}_1, \mathbf{x}_2)$ and non-increasing under the dynamics of the systems.

4.9 System Composition

Similar to model-reduction techniques, another approach to reduce the difficulties associated with the curse of dimensionality is *compositionality*. It may be possible to *decompose* large-dimension system down into several lower-dimension *subsystems*. This enables symbolic controllers to be designed for the subsystems separately, which together provide a control technique that covers the large-dimension system. Composition is discussed in detail in Chapter 9 as it is applied to the large-dimension New England 39-bus Test System (NETS).

4.9.1 Subsystems

Definition 4.15 (Subsystems). *Consider a network of N subsystems, where each subsystem i can be modelled by $\Sigma^i = (X^i, U^i, V^i, W^i, g^i, Y_1^i, Y_2^i, h_1^i, h_2^i)$, and $i \in \{1, \dots, N\}$, where:*

- $X^i \subseteq \mathbb{R}^{n^i}$ are state sets of subsystems;
- $U^i \subseteq \mathbb{R}^{p^i}$ are control input sets of subsystems;
- $V^i \subseteq \mathbb{R}^{q^i}$ are external disturbance sets of subsystems;
- $W^i \subseteq \mathbb{R}^{r^i}$ are internal disturbance sets of subsystems;

- $g^i : X^i \times U^i \times V^i \times W^i \rightarrow X^i$ are transition maps describing the evolution of subsystems;
- $Y_1^i \subseteq \mathbb{R}^{m^i}$ are external output sets of subsystems;
- $Y_2^i \subseteq \mathbb{R}^{m^i}$ are internal output sets of subsystems;
- $h_1^i : X^i \rightarrow Y_1^i$ are external output maps of subsystems;
- $h_2^i : X^i \rightarrow Y_2^i$ are internal output maps of subsystems.

The evolution of subsystems can be characterised by

$$\Sigma^i : \begin{cases} \dot{\mathbf{x}}^i = g^i(\mathbf{x}^i, \mathbf{u}^i, \mathbf{v}^i, \mathbf{w}^i), \\ \mathbf{y}_1^i = h_1^i(\mathbf{x}^i), \\ \mathbf{y}_2^i = h_2^i(\mathbf{x}^i), \end{cases} \quad i \in \{1, \dots, N\}. \quad (4.2)$$

where $\mathbf{x}^i \in X^i$, $\mathbf{y}_1^i \in Y_1^i$, $\mathbf{y}_2^i \in Y_2^i$, $\mathbf{u}^i \in U^i$, $\mathbf{w}^i \in W^i$, and $\mathbf{v}^i \in V^i$ are measurable and potentially large.

In the following, the definition of interconnected systems is presented where subsystems Σ^i are connected with each other via internal disturbances \mathbf{w}^i .

Definition 4.16 (Interconnected Systems). *Consider a network of N subsystems Σ^i , as defined in Definition 4.15, with a coupling matrix \mathcal{M} among them. The interconnection of Σ^i for any $i \in \{1, \dots, N\}$, is the interconnected control system $\Sigma = (X, U, V, g, Y, h)$, denoted by $\mathcal{I}(\Sigma^1, \dots, \Sigma^N)$, such that $X := \prod_{i=1}^N X^i$, $U := \prod_{i=1}^N U^i$, $V := \prod_{i=1}^N V^i$, $g := \prod_{i=1}^N g^i$, $Y := \prod_{i=1}^N Y_{z_1}^i$, and $h := \prod_{i=1}^N h_{z_1}^i$, with internal disturbances constrained by*

$$[\mathbf{w}^1; \dots; \mathbf{w}^N] = \mathcal{M}[\mathbf{y}_2^1; \dots; \mathbf{y}_2^N].$$

The evolution of the interconnected system is therefore characterised by

$$\Sigma : \begin{cases} \dot{\mathbf{x}} = g(\mathbf{x}, \mathbf{u}, \mathbf{v}), \\ \mathbf{y} = h(\mathbf{x}) \end{cases} \quad z \in \{1, 2\}.$$

In Chapter 9, system composition is combined with approximate simulation to consider $\Sigma_1^i = (X_1^i, U_1^i, V_1^i, W_1^i, g_1^i, Y_{1_1}^i, Y_{1_2}^i, h_{1_1}^i, h_{1_2}^i)$ as the original (concrete) subsystem and $\Sigma_2^i = (X_2^i, U_2^i, V_2^i, W_2^i, g_2^i, Y_{2_1}^i, Y_{2_2}^i, h_{2_1}^i, h_{2_2}^i)$ as its (possibly) lower-dimensional abstraction (with $n_2^i \leq n_1^i$).

4.9.2 Assume-Guarantee Contracts

Just as approximate simulation relations relax the definitions of simulation relations, it is also possible to relax the definitions of temporal logic specifications so they can be used in the compositional setting. This is done through the use of *assume-guarantee contracts* [26].

The properties expected from a system are called its *guarantees*. Each guarantee \mathcal{G} relies on a set \mathcal{A} of properties called *assumptions*, expressing boundary conditions for the guarantee \mathcal{G} to hold. Guarantees can be combined using conjunction, where one or more guarantees provide a contract \mathcal{C} . Assumptions if false remove all guarantees that relied on that assumption - other guarantees may still hold. Mathematically, $\mathcal{A} \implies \mathcal{G}$.

For an interconnected system with multiple subsystems; assumptions and guarantees may have some independence from one another. To subsystems they may appear to be distinct subcontracts, but for the interconnected system they combine to provide strong guarantees. So, a contract for an interconnected system with three subsystems can be defined using subcontracts

$$\mathcal{C}_1 \oplus \mathcal{C}_2 \oplus \mathcal{C}_3 \preceq \mathcal{C},$$

where \oplus is the contract composition and \preceq is a *refinement relation*. For the general contracts $\mathcal{C}_i, \mathcal{C}_j$ under assumptions $\mathcal{A}_i, \mathcal{A}_j$ and providing guarantees $\mathcal{G}_i, \mathcal{G}_j$, respectively, \mathcal{C}_i refines \mathcal{C}_j or $\mathcal{C}_i \preceq \mathcal{C}_j$ if and only if $\mathcal{A}_j \subseteq \mathcal{A}_i$ and $\mathcal{G}_i \subseteq \mathcal{G}_j$.

Satisfaction of the contract is acquired when individual subsystems hold this refinement. Controllers designed on these subsystems then provide a decentralised approach to acquiring guarantees. Using this concept of assume-guarantee contracts the subcontracts \mathcal{C}_i can be replaced by specifications ψ_i where the composition of those specifications satisfies the specification of system ψ :

$$\psi_1 \oplus \psi_2 \oplus \psi_3 \oplus \dots \preceq \psi.$$

For a composed system, the specification can give guarantees based upon the underlying assumption that subsystems guarantee their individual specifications. The contracts then compose to form a composed system contract with guarantees.

4.10 Conclusion

This chapter has defined the key terms and theory which provides a foundation for the control technique to be used in the rest of the thesis. It was seen that

- systems are defined as a tuple of states, inputs, disturbances and outputs along with the mappings to describe state transitions;
- sound abstractions and complete abstractions which relate two systems to each other allow an abstraction to be used as a representation of the desired system, reducing computational complexity;
- symbolic models provide a finite-state representation of infinite-state systems, greatly reducing the state space under analysis;
- specifications can be rigorously defined for desired system behaviour. Using formal synthesis tools, these specifications can be guaranteed to hold through the formal controller synthesis;

- exact simulation relationships can be relaxed to approximate simulation relationships, useful for quantifying the error in model-order reduction; and
- large systems can be decomposed into smaller subsystems, which when combined with assume-guarantee contracts can provide techniques to find guarantees for large system without the computational overhead.

The next chapter will implement formal control synthesis on a small model of a smart grid. The controller will be designed to use a fleet of electric vehicles as its input, as well as an example using active buildings as the input.

Formal Synthesis for Frequency Regulation of Power Systems

In this chapter, a formal controller synthesis approach is proposed for integrating a population of *plug-in electric vehicles* (EVs) and *energy storage systems* (ESSs) to the frequency regulation of power systems. This chapter is based on the works [209] and [210]. A novel symbolic controller is designed and simulated for the Great Britain power system. The proposed controller enhances the frequency response behaviour of the system when encountering a large outage event. The symbolic controller guarantees the convergence of the after-event frequency to the specified safe interval and ensures the frequency never drops below the contingent zone. The majority of this chapter will discuss the control approach using EVs, at the end of the chapter a similar case study for ESSs is provided.

5.1 Introduction

On 9th August 2019, there was a power outage event started in Cambridgeshire, Great Britain (GB), due to a lightning strike that hit overhead transmission lines, affecting about 1 GW demand (i.e. around 2.5% of total electric demand of the UK). This caused countrywide losses comprising a 740 MW power station, a 1,200 MW wind farm (on the day outputting 800 MW) and various embedded generation unit losses leading to a total generation loss as part of the initial event of over 2000 MW. The frequency of the system fell to 48.8 Hz, below the statutory limit of 49.5 Hz, at which the automatic protection system known as *Low Frequency Demand Disconnection* (LFDD) Scheme are triggered to protect the other 95% demand. Due to the LFDD, over one million customers were affected by the disruption [49]. A normal frequency range was restored within 5 minutes but essential services such as transport, health and water were still affected up to two days later. This chapter is inspired by such an event.

Frequency response is the reaction to a change in grid frequency. Most frequency responses occur from the supply-side where the turning on and off of turbines balance the generation-consumption relationship within the power system. Future smart grid technologies look to use demand-side resources to regulate grid frequency which save costs, energy and time when power disruption events occur [33,189]. This chapter studies *primary frequency response* of the GB grid based on the model of [134] and shows how a formal controller for *plug-in electric vehicles* (EVs) and *energy storage systems* (ESSs) could aid frequency recovery during system contingencies.

EVs have been proposed as a means of frequency regulation due to the fast response the EVs can provide to a power disruption event. EVs are essential to the future of smart grids due to gas and diesel vehicles slowly being phased out [80]. Frequency regulation is the most beneficial ancillary service that EVs can provide due to minimal impacts on battery degradation. One study even argues that EV battery life can be extended if EVs take part in demand-side frequency response, compared with regular EV use [191]. Other EV ancillary services are discussed in [87]. EVs respond to frequency events depending on the type of plug-in charger that is being used. *Unidirectional* chargers receive power from the grid and when signalled they stop charging to reduce grid the demand. *Bidirectional* chargers have the option to discharge energy stored in a EV back into the grid which leads to wider frequency response services. Bidirectional charging is likely to only be viable for level 2 type chargers, while unidirectional charging would be valid for all other levels of charging speed [214]. Charging strategies are discussed in [23]. This chapter uses a simple model based on [133] and [81] to simulate the aggregate behaviour of a collection of EVs.

Formal methods can be used to achieve frequency response services in the smart grid. Formal methods give guarantees for safe operation in many safety critical systems. Similarly, *formal verification* is a technique used to verify if systems meet a desired specification. In Chapter 4, the details of encoding a desired specification in LTL were discussed. Such specifications are able to accurately capture the behaviour of a system over time [11].

Formal synthesis consists of designing controllers such that the system satisfies a desired specifications, e.g. the states remain in the safe region or reach a target region. Due to the continuous (infinite-state) nature of the state space, *abstraction techniques* are a key component of formal synthesis of systems. A system can be abstracted by partitioning the state space and representing partitions by single points in the abstract state space. The mathematical properties of the abstract system can be used to ensure satisfaction of properties in the original system [11,61]. From abstraction, LTL properties can be preserved with language equivalent relations [11,126]. Safety, avoiding "bad" states, and reachability, converging to a winning region, are common requirements of formal specifications. Safety properties can be verified and enforced with control barrier functions [12,113], and in [62], zonotopes are used for reachability. Available tools for formal verification and synthesis include, but are not limited to SCOTS [163], CORA [92], Pessoa [160] and SpaceEx [56] for non-probabilistic systems.

A simplified model of the GB power system is given in [134]. Aggregate models

of EVs are described in [81, 133] in the form of differential equations with nonlinear components. In this chapter, aggregate models of a collection of EVs are adapted to generate a baseline controller of the system which can use for comparison. Requirements on the frequency are expressed (always stay in a safe interval, and do not go outside of a smaller interval for more than a specific time period) as temporal logic formulae [19]. The available software tool SCOTS [163] is used to synthesise a controller for the network that guarantees satisfaction of the temporal formula. SCOTS is a software tool for automatic controller synthesis through discrete abstractions. Linear and nonlinear differential equations are over-approximated with finite-state symbolic models and controllers are obtained in the form of finite-state machines. The symbolic model gives an abstraction that over-approximates the behaviours of the original system. If the closed-loop symbolic model satisfies the specification, then the original system also satisfies the specification due to the symbolic approximation including all the behaviours of the original system.

In brief, the novel aspects of this chapter are summarized as follows:

- A formal controller synthesis approach for integrating a population of EVs in the power system;
- Application of formal methods in frequency regulation of the network;
- Design and simulation of a novel symbolic controller for the GB power system;
- The proposed controller, enhances the frequency response behaviour of the system when encountered with a large outage event;
- The symbolic controller guarantees the settlement of the after-event's frequency in the specified safe interval;
- An additional example showing the same approach using ESSs.

This chapter is organised as follows. Section 5.2 provides the current requirements on frequency of the grid in case of a power loss. Additionally, this section contains the GB model and the baseline controller adapted from the literature for integration of EVs. Section 5.3 shows how the requirements on the frequency as a temporal logic formula can be encoded. Section 5.4 provide the formal synthesis approach for finding a controller with guarantees on satisfaction of the requirements. In Section 5.5 simulations are presented of the formal controller comparing it to the baseline controller. In Sections 5.6, additional simulations for an ESS example are given. Finally, in Section 5.7 the chapter concludes with all the findings.

5.2 Frequency Control in Power Systems

The system described is a relevant representation of the frequency control in the GB system and can be used to develop a control logic.

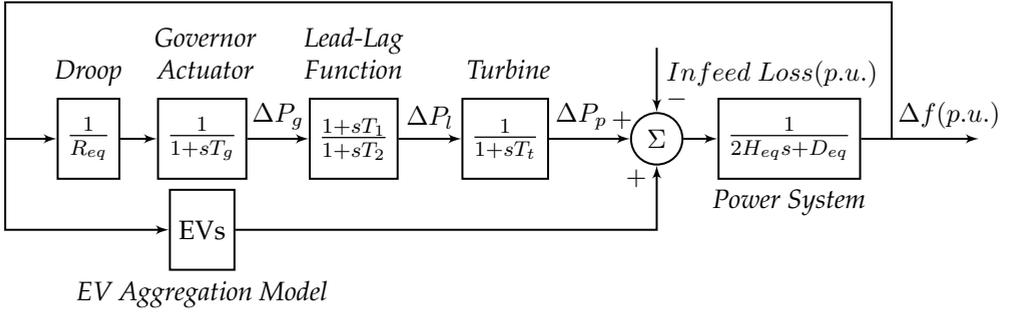


Figure 12: A simplified GB power system model including EVs for primary frequency response services, based on [134].

5.2.1 Frequency Regulation

Frequency is defined as the number of alternating current cycles per second (in Hertz) at which a system is running [138]. The *Electric System Operator* (ESO) increases or decreases system frequency using reserve and response services [137]. Positive service response increases generation or reduces demand while negative service response does the opposite. Positive response services provide power within seconds and are automatically triggered by local frequency readings while reserve services provide power after an instruction is received. If demand is greater than generation the frequency goes down, if demand is less than generation the frequency increases. The declared frequency of the GB grid is 50 Hz [136].

5.2.2 Requirements on Frequency

The focus of this chapter will be on events of *infrequent infeed losses* of 2000 MW, similar to the sequence of events mentioned in [49] that caused a 2000 MW total loss within a short period of time. When such large losses occur protocols such as LFDD are triggered to return stability to the system [138]. The current accepted maximum *normal infeed loss* for the GB grid is 1320 MW, while the maximum *infrequent infeed loss* is 1800 MW [142]. A containment zone is given for -0.8 Hz, this value is the maximum frequency deviation allowed for a loss greater than the normal infeed loss. For a normal infeed loss, the maximum deviation should stay within the statutory limits of 50 ± 0.5 Hz [173]. For plants taking part in frequency regulation, a droop characteristic of 3-5% is expected [138]. Frequency conditions are required to have a steady state within statutory limits for normal infeed losses and in the case of infrequent infeed loss, a violation should occur for no more than 60 seconds [142].

5.2.3 The GB Model

Figure 12 shows the GB grid model used in this chapter that consists of responsive synchronous plants and an aggregate group of EVs. The synchronous plants

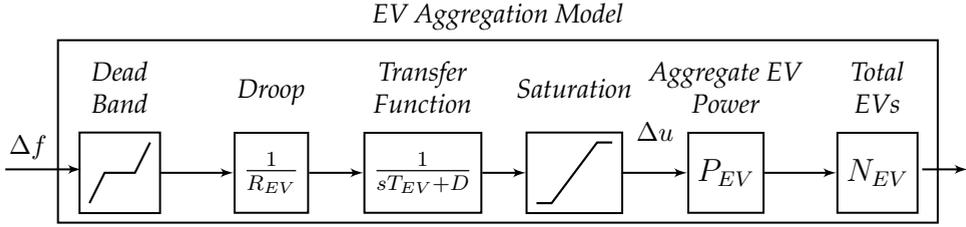


Figure 13: Baseline controller of EV frequency response services when a large power loss occurs, adapted from [81, 133].

model is discussed in depth in [134]. This chapter is desired to be a proof of concept not a replication of any specific event therefore values are set using Table 1. An extensive discussion on modelling of conventional power plants can be found in [95]. Included is an aggregate model of EVs that are in charging mode and if included in frequency response services, they will contribute to primary frequency control. There are three main frequency response conditions being considered for the EVs, depending on the charging strategy:

1. EVs do not participate in frequency response services and continue to charge;
2. EVs participate in primary frequency response when plugged in using unidirectional chargers;
3. EVs participate in primary frequency response when plugged in using bidirectional chargers.

5.2.4 Baseline Controller

The adapted aggregate model for these EVs is shown in Figure 13. The participation factor is the proportion of vehicles available to contribute to frequency control. In the baseline controller this is calculated by the components between (and including) the dead-band and saturation. The saturation in the system is used to determine the participation giving a value between 0 and 1. The participation value is multiplied by the power per unit (P_{EV}) and the number of vehicles (N_{EV}) to form the total power provided for frequency regulation. For bidirectional charging, energy can be discharged back into the grid if necessary so twice as much power is available per vehicle. It does not include moving vehicles as these would not contribute to charging demand in the system [81].

For this chapter, participation will be considered the input of the model. The output of the system is the frequency with respect to the time. As the frequency deviates from the nominal value the controller will increase the participation of the EVs in the system to provide response services and to return the frequency back to a steady state as near to the nominal value as possible. The baseline controller is adapted from [81, 133] and compared with the formal synthesis approach of this chapter that finds a symbolic controller with respect to the requirements on the frequency.

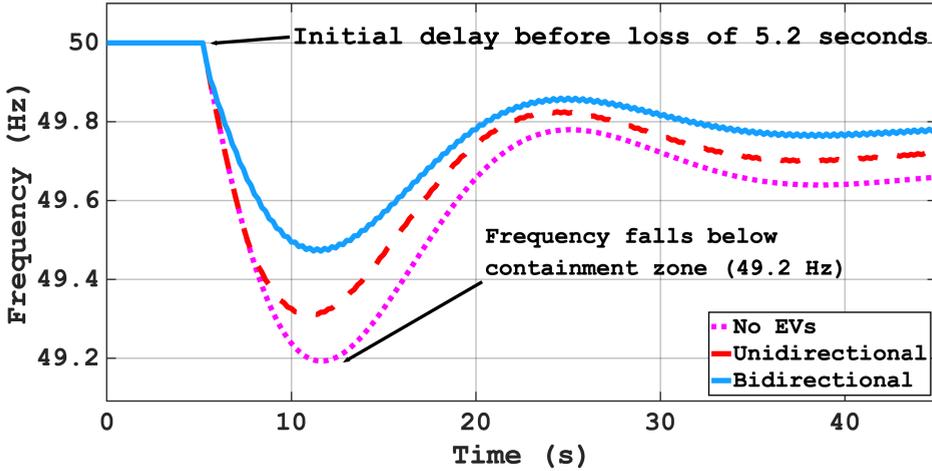


Figure 14: Frequency control under 2000 MW infrequent infeed loss using different EV charging strategies.

5.2.5 Baseline Simulation

When plotting the values from sections 5.2.3 and 5.2.4, Figure 14 is given. It can be seen that for losses of 2000 MW with no EV input to frequency regulation the containment limit of 49.2 Hz is breached, the system is in a delicate position and a large effort is required to return the frequency to stable conditions. Introducing EVs to primary frequency regulation when considering unidirectional charging improves the response of the system to large power losses. The frequency falls below the statutory limits. As this is an infrequent infeed loss, this is acceptable should the system return to the statutory limits within 60 seconds. In the case of both charging strategies this is true. Introducing bidirectional charging improves the recovery even further.

However, simulations will struggle to fully model a real GB system. Inertia changes due to the infeed loss are not considered within the system and a decrease in inertia leads to a larger rate of change of frequency. This means the maximum frequency loss could be greater than simulated and so suggested techniques may not be valid in practice. These results are therefore used as a basis for comparison.

Therefore the contribution of this chapter is not the simulation of theoretical results but the design of a controller with a given formal specification for how the system should behave. In this regard, the system will be able to show mathematically that a specification holds, using techniques such as over-approximation to provide formal proofs for the system. This chapter is a proof of concept with implications for extensions to more complex and real-time systems.

Table 1: values used for simulation adapted from [81, 133, 134].

Parameters	Unidirectional Value	Bidirectional Value
$1/R_{eq}$	-5	-5
T_g	2.5	2.5
T_t	0.5	0.5
T_1	2	2
T_2	12	12
D_{eq}	1.0	1.0
H_{eq}	4	4
T_{ev}	0.035	0.035
R_{ev}	0.5	0.5
P_{EV}	0.028	0.056
N_{EV}	25,000	25,000
deadband	50 ± 0.15	50 ± 0.15

5.3 Temporal Logic

We present again briefly some of the details of temporal logic that were discussed in Chapter 4. Temporal logic is a formalism for specifying desired properties of systems that evolve over time. *Linear temporal logic* (LTL) is a logic that provides a high-level language for describing such desired behaviour. This logic is primarily employed for the study of temporal behaviour of finite-state systems [19]. In this chapter, LTL is considered for specifying the desired behaviour of the frequency of the grid. LTL formulas ψ from Def. 4.10:

$$\psi := \top \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \bigcirc\psi \mid \psi_1 \text{ U } \psi_2,$$

where $p \subset AP$ is an observation in the system and ψ, ψ_1 and ψ_2 are LTL formulas. In addition to the aforementioned operators, LTL can also use *disjunction* \vee , *eventually* \diamond , and *always* \square operators as $\psi_1 \vee \psi_2 = \neg(\neg\psi_1 \wedge \neg\psi_2)$, $\diamond\psi := (\top \text{ U } \psi)$ and $\square\psi := \neg(\diamond\neg\psi)$ respectively.

5.3.1 Formalising the Specification for Frequency

As described in Section 5.2.1, the acceptable behaviours of the frequency as a function of time when considering power loss in the GB grid are given in natural language. These specific behaviours can be express formally in LTL as follows.

First, the frequency should never drop below the containment zone ($\mathcal{Z} = 49.2$ Hz) as, at this frequency, larger scale frequency response is required to return the system to normal limits and can include *load shedding* which is hugely disruptive. This can be written as the safety specification

$$\psi_1 := \square(f \geq \mathcal{Z}).$$

Second, the frequency should remain within the statutory limits $\mathcal{S} = [49.5, 50.5]$

Hz, for any *normal power losses*, i.e., $loss \leq \mathcal{L}$ with a predefined \mathcal{L} . This can be represented as the LTL formula

$$\psi_2 := (loss \leq \mathcal{L}) \implies \Box(f \in \mathcal{S}).$$

Third, for *infrequent infeed losses* specified by the inequality $loss \geq \mathcal{L}$ with a predefined \mathcal{L} , the frequency must return within 60 seconds to the statutory limits whenever it leaves that limit. This can be written as the LTL formula

$$\psi_3 := (loss \geq \mathcal{L}) \implies \Diamond_{60}(f \in \mathcal{S}) \quad (5.1)$$

Note that \Diamond_{60} means the condition holds eventually within the next 60 seconds. Finally, the desired behaviour of the frequency can be written as

$$\psi = \psi_1 \wedge \psi_2 \wedge \psi_3.$$

Specification for designing the controller. As the focus is an infrequent infeed loss with specification (5.1), reachability is focused on to show that LTL has the capability of defining a much richer class of behaviours. In particular, a two-stage controller is considered for the frequency regulation. The first controller is responsible for bringing the frequency inside an interval \mathcal{T}_1 and the second controller is responsible for bringing the frequency inside a smaller interval $\mathcal{T}_2 \subset \mathcal{T}_1$.

$$\begin{aligned} \psi := & \Box(f \geq \mathcal{Z}) \wedge [\neg(f \in \mathcal{T}_1) \implies \Diamond(f \in \mathcal{T}_1)] \wedge \\ & [(f \in \mathcal{T}_1 \wedge f \notin \mathcal{T}_2) \implies \Diamond(f \in \mathcal{T}_2)]. \end{aligned} \quad (5.2)$$

This specification reduces the pressure on the first controller by bringing the frequency inside the smaller interval \mathcal{T}_2 in multiple phases. Note that since only primary frequency response is considered, it is not necessary for the frequency to return to 50 Hz as other response schemes would respond in real-time scenarios to aid the full recovery. Therefore, the specification does not consider any requirement in ψ on the steady state being at 50 Hz.

Remark 5.1. *I note here, that for the following LTL case study we consider eventually without a restricted time horizon. It would be possible to verify this time horizon after the fact, or to synthesise the controller with this horizon built in, (similar to approaches that consider Markov decision processes for stochastic systems). But these ideas are beyond the scope of this chapter.*

5.4 Formal Controller Synthesis

This section discusses how to formally design a controller for integrating EVs in the grid such that the frequency satisfies the desired behaviour. Such a formal controller design requires that the time evolution of the system is written down as a dynamical system with differential equations affected by inputs and disturbances.

5.4.1 Grid as a Dynamical System

The simplified grid model of Fig. 12 can be represented as a dynamical system by converting the transfer functions into differential equations. The dynamics of such a system can be written as

$$\begin{aligned}
 \dot{f}(t) &= \frac{1}{2H_{eq}} P_p(t) + \frac{P_{EV} N_{EV}}{2H_{eq}} \mathbf{u}(t) - \frac{1}{2H_{eq}} \mathbf{v}(t) - \frac{D_{eq}}{2H_{eq}} f(t) \\
 \dot{P}_g(t) &= \frac{1}{T_g R_{eq}} f(t) - \frac{1}{T_g} P_g(t) \\
 \dot{P}_l(t) &= \frac{T_1}{T_2 T_g R_{eq}} f(t) + \frac{T_g - T_1}{T_2 T_g} P_g(t) - \frac{1}{T_2} P_l(t) \\
 \dot{P}_p(t) &= \frac{1}{T_t} P_l(t) - \frac{1}{T_t} P_p(t).
 \end{aligned} \tag{5.3}$$

Using these equations, a state space model can be constructed of the form

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{D}\mathbf{v}(t), \tag{5.4}$$

where $\mathbf{x} = [\Delta f, \Delta P_g, \Delta P_l, \Delta P_p]^T \in X \subseteq \mathbb{R}^4$ is the state vector (shifted around their nominal values), $\mathbf{u} \in [0, 1]$ is the participation ratio as the input, and $\mathbf{v} \in \mathbb{R}$ is the power loss. The state matrices are

$$\mathbf{A} = \begin{bmatrix} \frac{-D_{eq}}{2H_{eq}} & 0 & 0 & \frac{1}{2H_{eq}} \\ \frac{1}{T_g R_{eq}} & \frac{-1}{T_g} & 0 & 0 \\ \frac{T_1}{T_2 T_g R_{eq}} & \frac{T_g - T_1}{T_g T_2} & \frac{-1}{T_2} & 0 \\ 0 & 0 & \frac{1}{T_2} & \frac{-1}{T_t} \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} \frac{p_{ev} N_{ev}}{2H_{eq}} \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \frac{-1}{2H_{eq}} \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

The power loss is treated as a disturbance $\mathbf{v}(t)$ which is bounded by the maximum power loss.

5.4.2 Symbolic Model of the Grid

Definition 5.1. A symbolic model of dynamical system (5.3) for a sampling time τ is a transition system of the form $\hat{\Sigma} := (\hat{X}, \hat{U}, \hat{g})$, where \hat{X} is a finite partition of the state space of (5.3), \hat{U} is a finite subset of input set of (5.3), and $\hat{g} : \hat{X} \times \hat{U} \rightarrow 2^{\hat{X}}$ is a transition relation with $2^{\hat{X}}$ being the power set of \hat{X} , see also Def. 4.8.

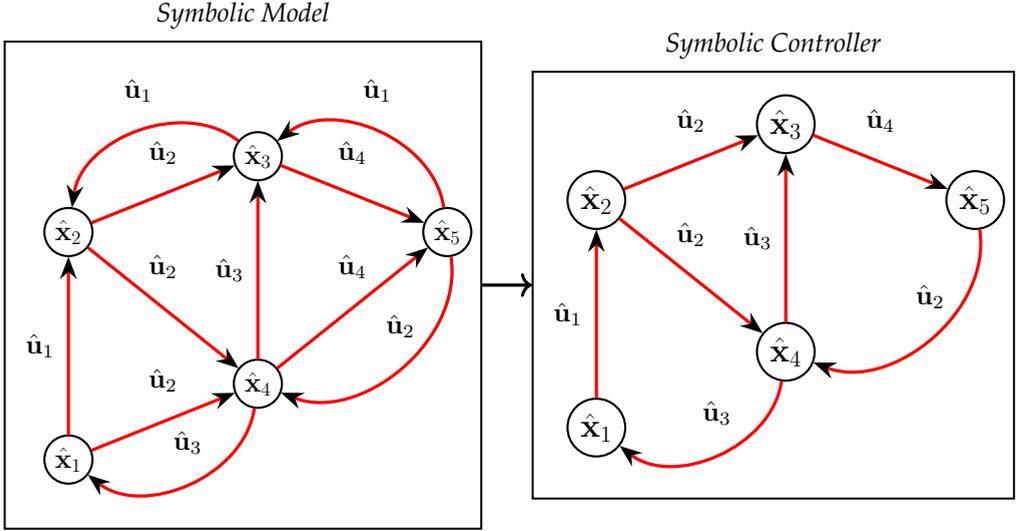


Figure 15: **Representation of the Symbolic Model $\hat{\Sigma}$ and the Symbolic Controller $\hat{\mathbf{C}}$.** The symbolic model shows the state space \hat{X} with $\{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_5\} \in \hat{X}$ representing partitions of the continuous state space and the input space \hat{U} with values $\{\hat{u}_1, \dots, \hat{u}_4\} \in \hat{U}$. The transition relation \hat{g} is shown graphically. The symbolic controller stores the appropriate inputs \hat{u} for each state \hat{x} . Thus the symbolic controller is treated as a lookup table, providing the input when the system is in a given state to guarantee specification of a satisfaction.

The transition relation $\hat{g}(\hat{x}, \hat{u})$ is defined as follows: compute all state trajectories of (5.3) starting from partition set \hat{x} under the input \hat{u} and for all possible values of the disturbance; then $\hat{x}' \in \hat{g}(\hat{x}, \hat{u})$ if \hat{x}' intersects with this set of trajectories after a fixed sampling time τ . Fig. 15 (left) shows a graphical representation of a symbolic model.

Theorem 5.1. *The particular construction of the symbolic model $\hat{\Sigma}$ implies that $\hat{\Sigma}$ overapproximates the trajectories of the original model. Thus if a controller is found on the symbolic model that satisfies a given specification, the original model will also satisfy the same specification for all disturbance trajectories.*

Available tools for computation of symbolic models and design of symbolic controllers include SCOTS. These tools usually rely on overapproximating the trajectories of the original model using growth bounds that depend on continuity properties of the differential equations (5.3). They also use fixed point computations for finding symbolic controllers. In this chapter, SCOTS is used as it enables designing symbolic controllers that have the ability to synthesise controllers for reach, reach-avoid and safety specifications. The computations in SCOTS are implemented in C++ language with a MATLAB interface to view the symbolic state space along with simulating the closed loop system.

5.4.3 Symbolic Control for the Grid

A symbolic controller $\hat{\mathbf{C}}$ for the symbolic model $\hat{\Sigma} := (\hat{X}, \hat{U}, \hat{g})$ defined in Def. 5.1 is in the form of $\hat{\mathbf{u}} = \hat{\mathbf{C}}(\hat{\mathbf{x}})$ that assigns any partition set $\hat{\mathbf{x}} \in \hat{X}$ to an input $\hat{\mathbf{u}} \in \hat{U}$ in order to satisfy the given specification on $\hat{\Sigma}$. Such a controller is used to construct a controller \mathbf{C} for the original system (5.4) as follows. Then $\mathbf{u}(t) = \mathbf{C}(\mathbf{x}(t))$ with $\mathbf{x}(t) \in \hat{\mathbf{x}}$ and $\mathbf{u}(t) = \hat{\mathbf{C}}(\hat{\mathbf{x}})$. In other words, the partition set of $\mathbf{x}(t)$ is identified and the input related to that partition set in the symbolic controller is selected as the input for the original system. Fig. 15 (right) shows a graphical representation of a symbolic controller.

In the construction of the symbolic model $\hat{\Sigma}$, the working region of state variables is selected as $\Delta f \in [-1, 0.1]$, $\Delta P_p \in [0, 3]$, $\Delta P_g \in [0, 2]$, and $\Delta P_l \in [0, 2]$. The values are generally chosen based on the time constants of the blocks in Fig. 12 and the range of inputs of these blocks. Adjustments are made to reduce computation time in simulation. Note that these are the states shifted around their nominal values. The working region is partitioned along each dimension with discretisation $\eta_x = 0.05$. For the input $\mathbf{u} \in [0, 1]$, discrete steps of 5% of the total input range is considered. From these partition sets as symbolic states and inputs, a symbolic model and a growth bound are calculated. The growth bound is calculated by taking the Jacobian of the right-hand side of (5.4) in the form of a Metzler matrix. This is the abstraction of the original system and the transition relation of this new system is computed for the fixed point computations. The fixed point computation of the reach specification using the target range is then calculated giving us a formally synthesised controller. The results of these controllers will be discussed in Section 5.5.

5.5 Implementation Results

In this section, symbolic controller synthesis is applied to the model of the GB power grid and compared with the baseline controller of Fig. 13. SCOTS was used for the design of the symbolic controllers and the simulations were implemented in MATLAB on a machine equipped with Intel Core i5-7267U 3.1GHz CPU and 8GB RAM. Computing each of the two controller's reach function takes approximately 28 seconds for unidirectional EVs and 31 seconds for bidirectional EVs.

5.5.1 Simulations with a Multi-Phase Controller

A symbolic controller has been designed for satisfying the specification ψ in (5.2). The results of the frequency response are presented in Fig. 16 and Fig. 18 for respectively bidirectionally and unidirectionally charged EVs. The containment zone ($f \leq \mathcal{Z}$) that should not be visited is shown in these figures with a box having red edges. The target regions $f \in \mathcal{T}_2$ and $f \in \mathcal{T}_1$ are shown with boxes having respectively green and black edges. The symbolic controller uses the following phases control strategy once the power loss occurs:

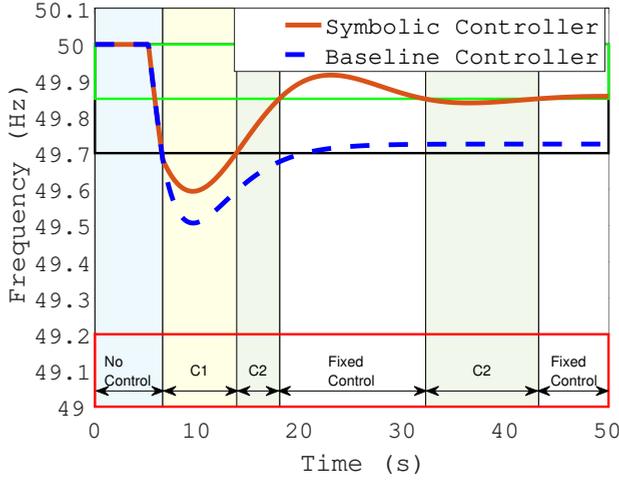


Figure 16: Symbolic control for frequency regulation with bidirectionally charged EVs. The frequency under the synthesised symbolic controller satisfies the specification ψ in (5.2) with $\mathcal{Z} = [49.2, 50]$, $\mathcal{T}_1 = [49.70, 50]$ and $\mathcal{T}_2 = [49.85, 50]$ Hz, but the baseline controller of Fig. 13 is unable shape the frequency with respect to ψ .

1. If the state is still within the larger target region \mathcal{T}_1 , no participation of EVs is required ($\mathbf{u} = 0$). This phase is highlighted in light blue in the figures with name “No Control”.
2. Whenever the frequency leaves the larger target region \mathcal{T}_1 , a low-level symbolic controller is activated to bring the frequency inside \mathcal{T}_1 . This phase is highlighted in light yellow in the figures with name $\mathfrak{C}1$.
3. When the frequency is inside \mathcal{T}_1 but outside of the smaller target region \mathcal{T}_2 , a second low-level symbolic controller is activated to bring the frequency inside \mathcal{T}_2 . This phase is highlighted in light green in the figures with name $\mathfrak{C}2$.
4. Finally, if the frequency goes inside the smaller target region \mathcal{T}_2 , the last value of participation is used. This phase is in white colour in the figures with name “Fixed Control”.

Symbolic controllers $\mathfrak{C}1$ and $\mathfrak{C}2$ have been designed by solving two reachability problems with target regions \mathcal{T}_1 and \mathcal{T}_2 using SCOTS. The selected parameters are $\mathcal{T}_1 = [49.70, 50]$, $\mathcal{T}_2 = [49.85, 50]$ Hz for bidirectionally charged EVs and $\mathcal{T}_1 = [49.55, 50]$, $\mathcal{T}_2 = [49.75, 50]$ Hz for unidirectionally charged EVs. The results of the required participation are presented in Fig. 17 and Fig. 19.

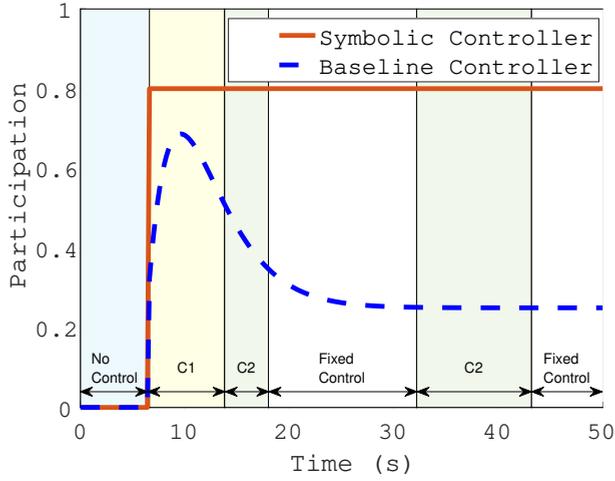


Figure 17: Percentage of participation of bidirectionally charged EVs as a function of time obtained from the synthesis approach to satisfy ψ in (5.2) and from the baseline controller of Fig. 13.

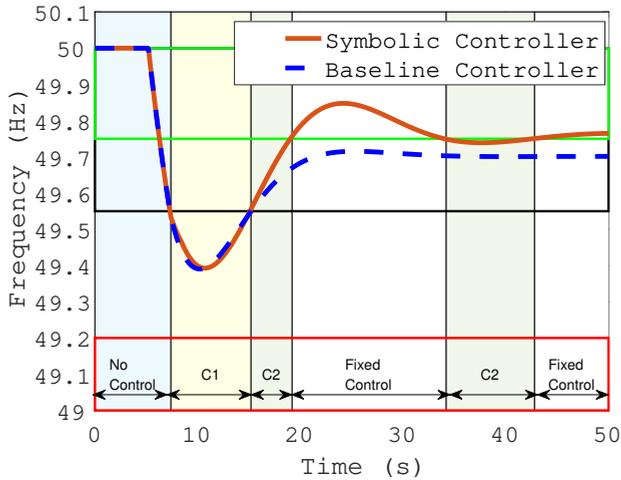


Figure 18: Symbolic control for frequency regulation with unidirectionally charged EVs. The frequency under the synthesised symbolic controller satisfies the specification ψ (5.2) with $\mathcal{Z} = 49.2$, $\mathcal{T}_1 = [49.55, 50]$ and $\mathcal{T}_2 = [49.75, 50]$ Hz, but the baseline controller of Fig. 13 is unable shape the frequency with respect to ψ .

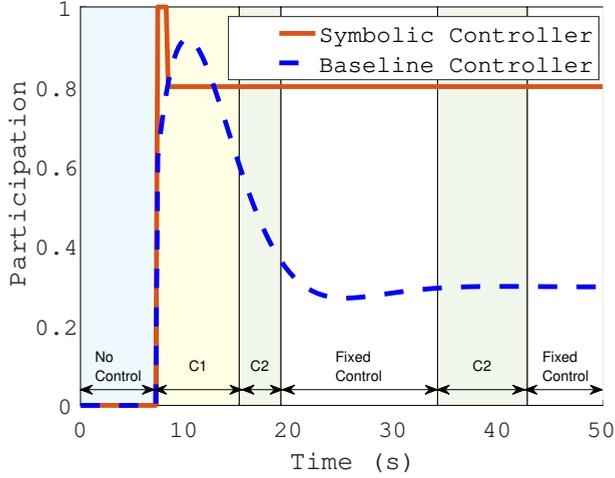


Figure 19: Percentage of participation of unidirectionally charged EVs as a function of time obtained from the synthesis approach to satisfy ψ in (5.2) and from the baseline controller of Fig. 13.

Table 2: Steady state frequency of Baseline Controller for different deadband thresholds (Hz)

Deadband	Unidirectional SS	Bidirectional SS
50 ± 0.00	49.73	49.77
50 ± 0.05	49.72	49.76
50 ± 0.10	49.71	49.74
50 ± 0.15	49.70	49.72
50 ± 0.20	49.69	49.71
50 ± 0.25	49.68	49.69
50 ± 0.30	49.67	49.68
50 ± 0.35	49.67	49.67

5.5.2 Formal Guarantees

In order to compare the performance of the approach with the baseline controller in Fig. 13, the GB model was simulated with the baseline controller having different values for the deadband threshold. The steady-state values of the frequency are reported in Table 2 for both unidirectional and bidirectional EVs. As can be seen, the highest steady-state frequency is achieved when deadband thresholds are both 50 Hz, i.e., no deadband component in the baseline controller which requires instantaneous response from the EVs. Even in such a case, the baseline controller is unable to satisfy the specification ψ in (5.2) as the steady-state is outside of the smaller target region \mathcal{T}_2 . In contrast, the multi-phase Controller based on the two symbolic controllers $\mathfrak{C}1$ and $\mathfrak{C}2$ satisfies the required specification. This comes at the cost of higher participation in comparison with the baseline controller as reported in Fig. 17 and Fig. 19.

5.5.3 Robustness of the Controller

To measure the robustness of the controller against uncertainty in the participation of the EVs, up to 10% uncertainty was allowed for the EV participation of the symbolic controller. Fig. 20 and Fig. 22 show that both cases of unidirectionally and bidirectionally charged EVs continue to facilitate satisfaction of the specification ψ , despite the uncertainty on the participation, although for bidirectional charging the specification is satisfied after a relatively longer time period (≈ 48 seconds). Fig. 21 and Fig. 23 show the variation in participation is substantial and that the fixed value assigned inside the winning region, can also fluctuate. Uncertainty has a larger effect on bidirectional charging than unidirectional charging as each bidirectional vehicle contributes double the power of its unidirectional equivalent. With increased uncertainty, the time taken to converge to the winning region also increases.

Overall, the design approach encourages more refined specifications for the frequency of the grid. It allows designing controllers automatically to satisfy those specifications with correctness guarantees and is more robust. Other approaches are unable to provide controllers automatically with correctness guarantees and require manual tuning of parameters while relying on simulations.

5.6 Using Energy Storage Systems for Frequency Regulation

In Chapter 3 and Chapter 2 active buildings were used for smart grid control. ESSs, can connect to the grid through homes or buildings and are used in aggregation for primary frequency response. In this section, the formal control approach demonstrates that it will satisfy the specification 5.2 and provide guarantees over the system. The simulation is based on the same GB power system dynamics

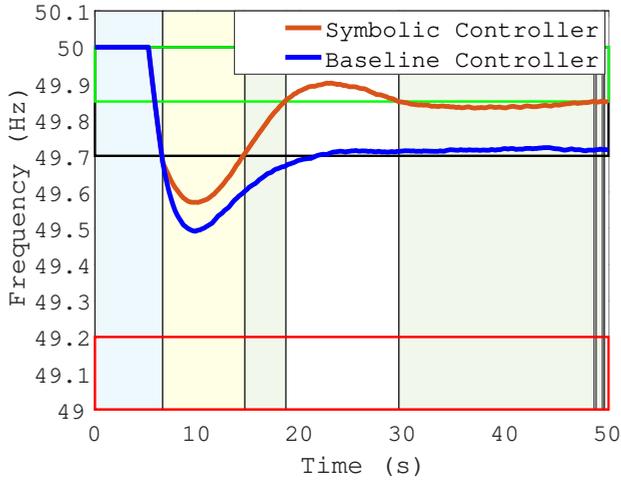


Figure 20: Symbolic control for frequency regulation with bidirectionally charged EVs with up to 10% uniformly distributed random uncertainty in participation. The frequency under the synthesised symbolic controller still satisfies the specification ψ in (5.2) but the baseline controller fails to do so.

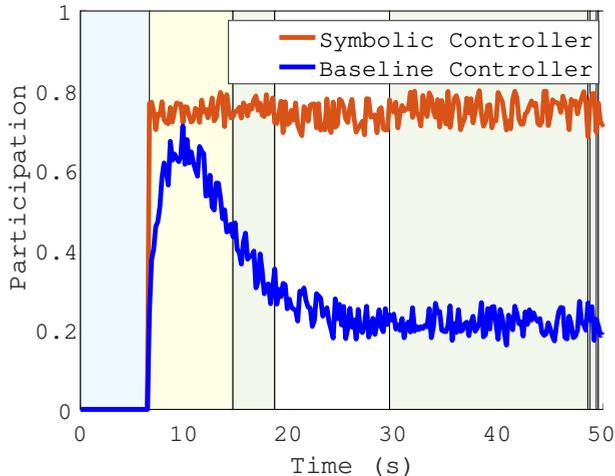


Figure 21: Percentage of participation of bidirectionally charged EVs that has up to 10% uniformly distributed random uncertainty in participation, as a function of time obtained from the synthesis approach to satisfy ψ in (5.2) and from the baseline controller of Fig. 13.

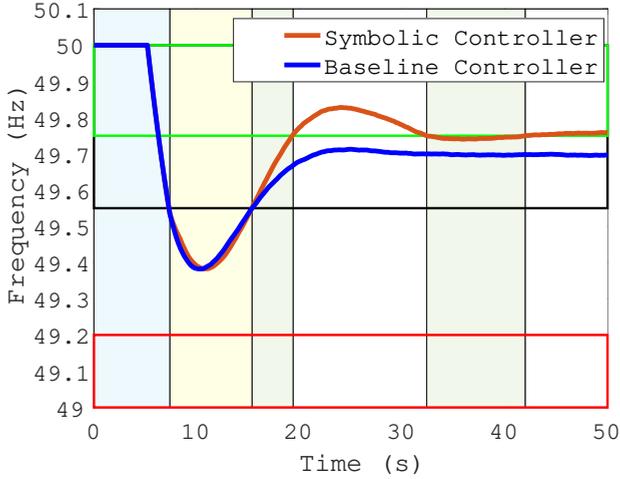


Figure 22: Symbolic control for frequency regulation with unidirectionally charged EVs that has up to 10% uniformly distributed random uncertainty in participation. The frequency under the synthesised symbolic controller still satisfies the specification but the baseline controller fails to do so.

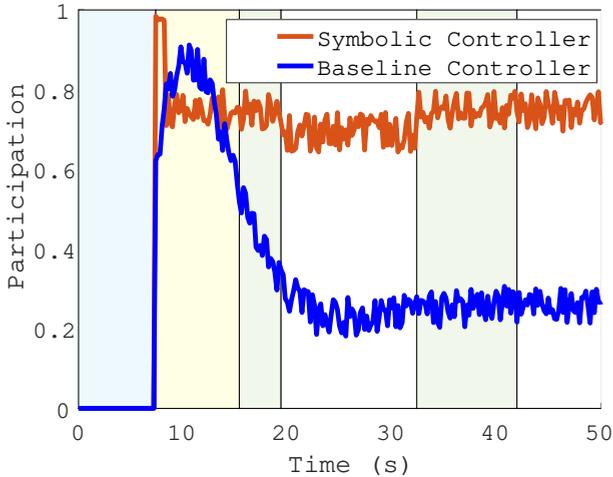


Figure 23: Percentage of participation of unidirectionally charged EVs that has up to 10% uniformly distributed random uncertainty in participation, as a function of time obtained from the synthesis approach to satisfy ψ in (5.2) and from the baseline controller of Fig. 13.

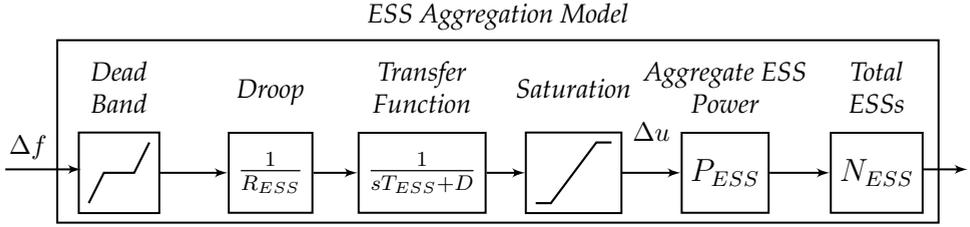


Figure 24: Baseline controller of ESS frequency response services when a large power loss occurs, adapted from [81, 133].

with a different baseline controller considering the ESSs. The expected infrequent infeed loss is 2000 MW. For the system input, the percentage of ESSs which participate in demand-side response is used.

The formal synthesis was completed inside the SCOTS software tool. Passing the formal specification, system dynamics, maximum and minimum values of the states and size of the partitions into the SCOTS tool will enable it to attempt the formal controller synthesis. As the technique involves removing states from the state space that do not lead to satisfaction, i.e. where $U(x) = \emptyset$, all states could be removed. In this case the synthesis will have failed and the SCOTS tool returns no controller to the user. As SCOTS builds the relationship $\Sigma \cong \hat{\Sigma}$, should the SCOTS tool return a controller to the user then guarantees are provided that the controller will satisfy the specification.

After successfully synthesising the controller for the system, simulations to compare the performance to a baseline controller are devised. The baseline controller accepts the system frequency, Δf , as an input and consists of a frequency dead-band, a transfer function and a saturation block to determine the percentage of ESSs required to participate in frequency regulation, Δu . This participation value is multiplied by the aggregate ESS power, P_{ESS} , and number of ESSs, N_{ESS} , to calculate a total demand response from ESSs for frequency regulation at that time instance; as shown in Fig. 24.

The control strategy implemented involved several phases which are each represented in the further figures by a shaded colour in the background. In the first phase, there is no initial ESS participation as the frequency is inside the safe interval \mathcal{T}_1 , phase shown in light blue. When the frequency leaves the safe interval, the first controller \mathcal{C}_1 attempts to return the frequency to \mathcal{T}_1 again, phase shown in yellow. Once achieved, controller \mathcal{C}_2 takes over and tries to transition the frequency to smaller region \mathcal{T}_2 , phase shown in white. Upon entering \mathcal{T}_2 , the ESS participation value is fixed, unless the frequency oscillates outside the region or normal system operation is restored, phase shown in green. In the same manner the baseline controller provides no participation when the frequency is inside the larger safe zone \mathcal{T}_1 , and this is managed by the deadband.

In Fig. 25 and the following figures, the region \mathcal{T}_1 is marked in black, \mathcal{T}_2 in green, and \mathcal{Z} in red. This model updated every 0.2 seconds. The frequency under the synthesised symbolic controller satisfies the formal specification ψ with $\mathcal{T}_1 =$

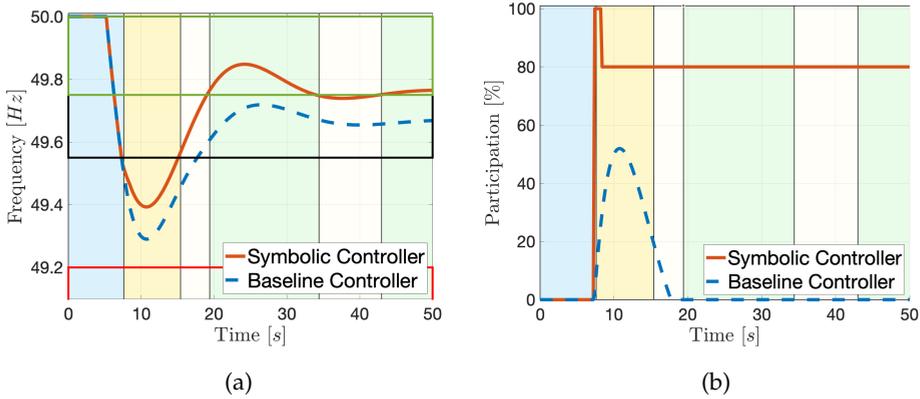


Figure 25: Symbolic control of the frequency using ESSs **(a)** and the participation of ESSs over time **(b)**.

$[49.55, 50]$ and $\mathcal{T}_2 = [49.75, 50]$ Hz, but the baseline controller is unable to shape the frequency with respect to ψ .

To test the robustness of the controller, uncertainty is added to ESS response. Adding a uniformly distributed random uncertainty of up to 10% participation was found to still satisfy the specification as shown in Fig. 26, and outperform the equivalent baseline controller. In both controller scenarios the minimum frequency value shifts slightly lower, this is because the uncertainty in the participation reduces the amount of energy used for response. This leads to a larger RoCoF and so, for the same time horizon, lower frequency values are reached.

This case study approach can also be applied to more complex and nonlinear systems, but with increased computations. Fig. 26 showed the results of the system with an uncertain participation value but extensions can also be made to include uncertainty of the plant values or increasing the participation uncertainty. To continue to provide guarantees in the cases of uncertainty, one approach uses formal methods for stochastic systems. The guarantees are to a certain confidence level, this can be a % confidence that the guarantees are provided or a full guarantee within a particular range.

5.7 Conclusion

In conclusion, this chapter proposes a new approach for integration of EVs and ESSs in frequency response services with the following features:

- A proof of concept for the design and use of symbolic controllers in primary frequency response services;
- Using temporal logic to encode the requirements on the frequency that are usually expressed in natural language;

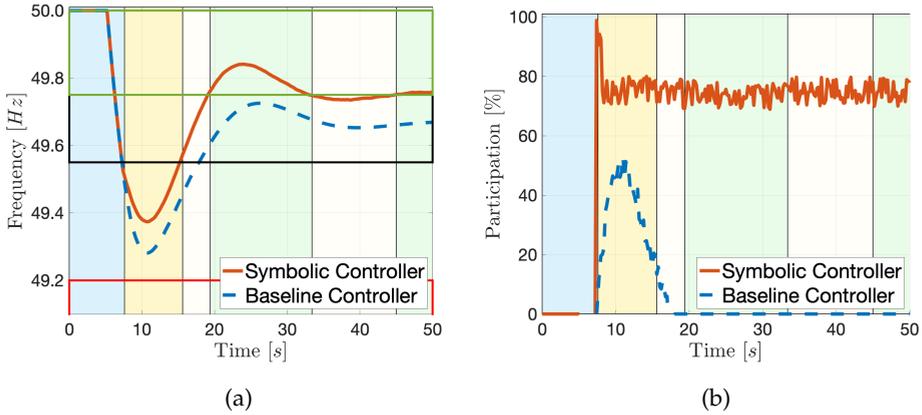


Figure 26: Symbolic control of the frequency using uncertain ESSs **(a)** and the participation of uncertain ESSs over time **(b)**.

- Formal guarantees on satisfaction of such requirements under the synthesised symbolic controller;
- Enhanced response to large frequency loss events with symbolic control due to a more robust controller design;
- The controller is robustness against uncertainty in the EVs participation;
- Simulation results that show correctness of the controller design against a more refined specification on the GB grid.

The next chapter will extend the symbolic control method from a model-based control method seen in this chapter to a data-driven control method.

Data-Driven Abstraction-Based Control Synthesis

This chapter investigates formal synthesis of controllers for continuous-space systems with unknown dynamics to satisfy requirements expressed as linear temporal logic formulas. Abstraction-based schemes relying on precise mathematical models are not applicable when the dynamics of the system are unknown. This chapter is based on the work [86], which computes a growth bound of the system using a finite number of trajectories. The computed growth bound together with the sampled trajectories are then used to construct the abstraction and synthesise a controller. The approach casts the computation of a growth bound as a robust convex optimisation program (RCP). Since the unknown dynamics appear in the optimisation, a scenario convex program (SCP) is formulated corresponding to the RCP using a finite number of sampled trajectories. The data-driven approach is demonstrated on several case studies, including a DC-DC boost converter and a 3 area 3 machine power system example.

Notation. The operator $\|\cdot\|$ denotes the infinity norm. The notation $\Omega_\varepsilon(c) := \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - c\| \leq \varepsilon\}$ denotes the ball with respect to infinity norm centred at $c \in \mathbb{R}^n$ with radius $\varepsilon \in \mathbb{R}_{>0}^n$. A probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$ is defined, where Ω is the sample space, \mathcal{F}_Ω is a sigma-algebra on Ω comprising its subsets as events, and \mathbb{P}_Ω is a probability measure that assigns probabilities to events. Primarily in this chapter, the terms finite abstractions and continuous-state systems are used instead of the equivalent terms symbolic models, and infinite-state systems. Systems are described as having action spaces, equivalent to input spaces, the space from which an action/input is selected for control.

6.1 Introduction

One of the major objectives in the design of safety-critical systems is to ensure their safe operation while satisfying high-level requirements. Through this thesis the safety-critical systems of interest are power systems, and in particular smart grids. Although safety-critical systems also include autonomous vehicles, traffic control, and battery-powered medical devices. Automatic design of controllers for such systems that can fulfil the given requirements have received significant attention recently. These systems can be represented as control systems with continuous state spaces. Within these continuous spaces, it is challenging to leverage automated control synthesis methods that provide satisfaction guarantees for high-level specifications, such as those expressed in Linear Temporal Logic [19, 25, 64, 188].

A common approach to tackle the continuous nature of the state space is to use *abstraction-based controller design* (ABCD) schemes [25, 122, 167, 188]. The first step in the ABCD scheme is to compute a finite abstraction by discretising the state and action spaces. Finite abstractions are connected to the original system via an appropriate behavioural relation such as feedback refinement relations or alternating bisimulation relations [153, 188]. Under such behavioural relations, trajectories of the abstraction are related to the ones of the original system. Therefore, a controller designed for the simpler finite abstract system can be refined to a controller for the original system. The controller designed by the ABCD scheme is described as being formal due to the guarantees on satisfaction of the specification by the original system in a closed loop with the designed controller.

ABCD schemes generally rely on a precise mathematical model of the system. This stems from the fact that establishing a behavioural relation between the original system and its finite abstraction uses reachability analysis over the dynamics of the original system that require knowledge of the dynamical equations. Although such equations can in principle be derived for instance by using physics laws, the real-world control systems are a mixture of differential equations, block diagrams, and lookup tables. Therefore, extracting a clean analytical model for systems of practical interest could be infeasible. A promising approach to tackle this issue is to develop data-driven control synthesis schemes with appropriate formal (probabilistic) guarantees.

6.1.1 Contributions

The main contribution of this chapter is to provide a data-driven approach for formal synthesis of controllers to satisfy temporal specifications. A brief overview of the approach can be seen in Fig. 27. This chapter focuses on continuous-time nonlinear dynamical systems whose dynamics are unknown but sampled trajectories are available. The approach constructs a *finite* abstract model of the system using only a finite number of sampled trajectories and a growth bound of the system. This approach formulates the computation of a growth bound as a *robust convex program* (RCP) that has infinite uncountable number of constraints. This then

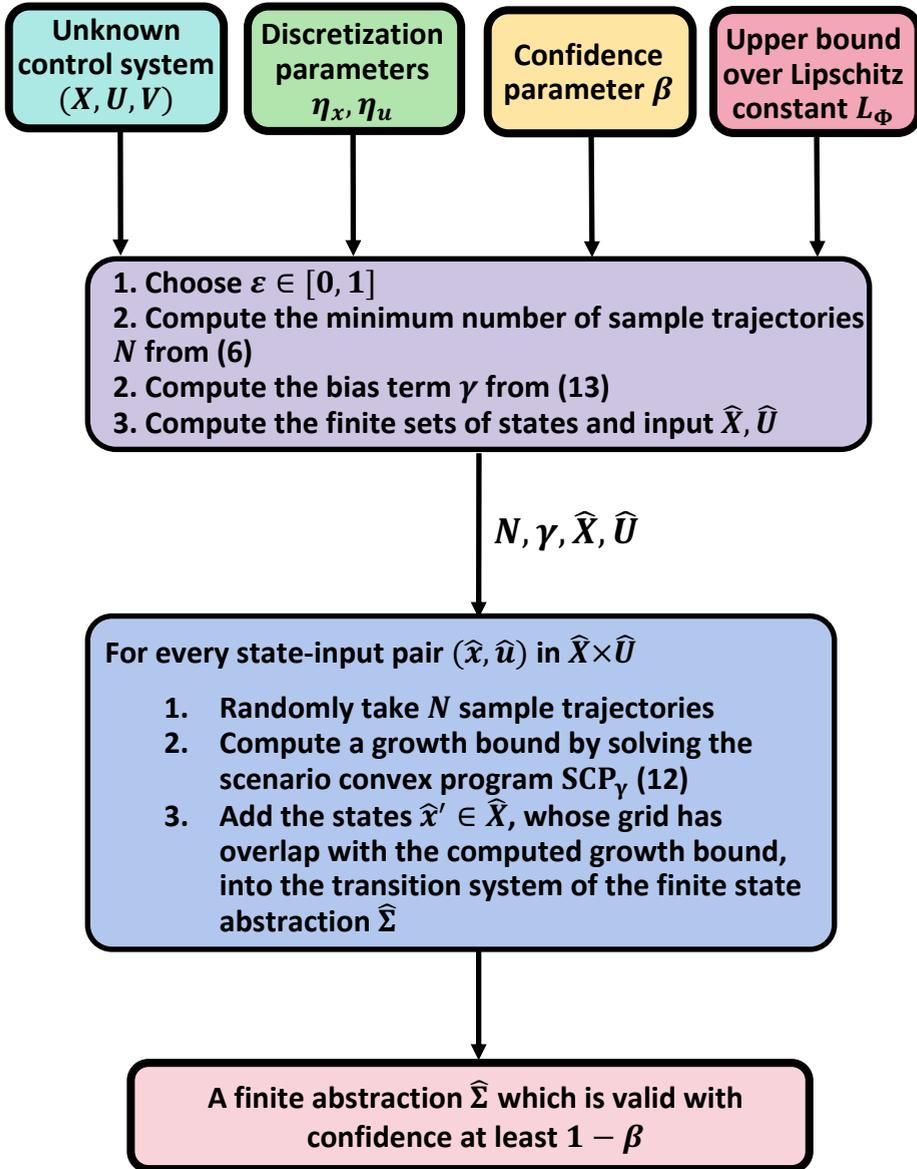


Figure 27: A diagram illustrating the steps of the proposed data-driven method for constructing finite abstractions.

approximates the solution of the RCP with a *scenario convex program* (SCP) that has a finite number of constraints and can be solved using only a finite set of sampled trajectories. This establishes a sample complexity result that gives a lower bound for the required number of trajectories to guarantee the correctness of the computed growth bound *over the whole state space* with a given confidence. A sample complexity result for the satisfaction of the specification on the system in closed loop with the designed controller for a given confidence is also provided. The result requires estimating a bound on the Lipschitz constant of the system with respect to the initial state, that is obtained using extreme value theory. As a last contribution, this chapter shows that the approach can be extended to a model-free abstraction refinement scheme by modifying the formulation of the system's growth bound and providing similar sample complexity results. The performance of the approach is demonstrated on two case studies.

The sample complexity result requires knowing a (possibly conservative) bound on the Lipschitz constant of the system. Algorithms founded on the extreme value theory can be utilised to estimate the true Lipschitz constant [202,204]. The estimated values will converge to a true Lipschitz constant when the number of samples used for this estimation goes to infinity. This sample complexity result for the abstraction and synthesis is valid under the assumption that these estimations give a correct upper bound on the Lipschitz constant.

The remainder of this chapter is organised as follows. After discussing the related work, Section 6.2 covers preliminaries on dynamical systems and finite abstractions and provides the problem statement. Section 6.3 presents the assumptions and theoretical results needed for connecting RCPs and their corresponding SCPs. Section 6.4 presents the approach on data-driven computation of a growth bound and the abstraction and prove the sample complexity result under the assumption of knowing a bound on the Lipschitz constant of the system. Estimation of the bound on the Lipschitz constant of the system for computing the number of samples is also discussed in this section. Section 6.5 discusses the extension of the approach to a data-driven abstraction refinement scheme. Several numerical examples are provided in Section 6.6 that support the theoretical findings. Finally, Section 6.7 concludes.

6.1.2 Related Work.

There is an extensive body of literature on *model-based* formal synthesis for both deterministic and probabilistic systems, such as the books [19,25,188], seminal chapters [3,64], and the survey chapter [99]. *Data-driven* approaches for analysis, verification, and synthesis of systems have received significant attention to improve efficiency and scalability of model-based approaches, and to study problems in which a model of the system is either not available or costly and time-consuming to construct. Given a prior inaccurate knowledge about the model of the system, a research line is to use data for refining the model and then synthesise a controller. Such approaches assume a class of models and improve the estimation of the uncertainty within the model class. These approaches range from using

Gaussian processes [20, 129], differential inclusions [46], rapidly-exploring random graphs [70], piecewise affine models [165], and model-based reinforcement learning algorithms [38]. A data-driven framework is proposed by Fan *et al.* [52] for verifying properties of hybrid systems when the continuous dynamics are unknown but the discrete transitions are known.

The work [37] considers the verification of stochastic neural network controllers for signal temporal logic specifications. Similarly, the works [111, 112] guarantees signal temporal logic specifications for control barrier functions using data. Signal temporal logic is studied for semi-supervised control design in [108]. Data-driven reachability analysis is studied in [109]. The work [152] discusses control synthesis using deep kernel learning. Safety guarantees for bayesian neural networks are discussed in [203]. Forward invariance in neural network controlled systems is examined in [74]. Interval reachability for nonlinear systems with neural network controllers is considered in [82].

Data-driven approaches for solutions of scenario convex programs are developed for switching systems by Wang and Jungers to establish stability [200] and by Berger *et al.* for invariant subspace identification [27]. Ahmad *et al.* [5] have developed an adaptive sampling-based approach for motion planning using deterministic nonlinear control systems and robust control barrier functions. Zhong *et al.* [217] have studied linear dynamical systems with bounded disturbances by proposing a data-driven method to compute state feedback controllers that enforce staying in safety invariant sets by using finite number of state-input data points. Cohen *et al.* [39] have developed a model-based reinforcement learning approach to satisfy linear temporal logic specifications on continuous-time nonlinear systems.

Data-driven model-free approaches compute the solution of the synthesis problem directly from data without constructing a model. Hsu *et al.* [76] provide a reach-avoid Q-learning algorithm with convergence guarantees for an arbitrarily tight conservative approximation of the reach-avoid set. Wang *et al.* [198] propose a falsification-based adversarial reinforcement learning algorithm for metric temporal logic specifications. Satisfying signal temporal logic specifications is studied by Verdier *et al.* [194] using counterexample-guided inductive synthesis on nonlinear systems, and using model-free reinforcement learning by Kalagarla *et al.* [83] to satisfy signal temporal logic specifications. A learning framework for synthesis of control-affine systems is provided by Sun *et al.* [183]. Watanabe *et al.* [201] study learning from demonstration while preventing the violation of safety under the learned policy. The recent chapters [103, 166] propose a data-driven approach to compute barrier certificates with correctness guarantees on satisfaction of safety specifications.

The research on data-driven constructions of abstract models is very limited. Legat *et al.* [107] provide an abstraction-based controller synthesis approach for hybrid systems by computing Lyapunov functions and Bellman-like Q-functions, and using a branch and bound algorithm to solve the optimal control problem. This differs from this approach where it is wanted to satisfy temporal specifications instead of solving optimal control problems. Makdesi *et al.* [124, 125] studied unknown monotone dynamical systems and sampled a set of trajectories generated by the system to find a minimal map overapproximating the dynamics of any sys-

tem that produces these transitions. Consequently, they calculate an abstraction of the system related to this map and prove that an alternating bisimulation relation exists between them. In contrast, the approach is not restricted to monotone systems and is applicable to any nonlinear dynamical system. Abstract models are also constructed for stochastic systems using sampled data.

Data-driven construction of abstract models for stochastic systems has also been studied recently. Badings *et al.* [17, 18] consider constructing abstract models in the form of interval Markov decision processes (IMDPs) by computing probably approximately correct (PAC) bounds on the transition probabilities of the system. This makes the approach applicable for satisfying infinite-horizon specifications and providing confidence bounds on the (probabilistic) satisfaction of the specification. The work by Lavaei *et al.* [104] constructs finite MDPs using data for general nonlinear stochastic systems utilising the concept of stochastic bisimulation functions. The focus of these works is on stochastic systems, but this chapter develops the results for non-probabilistic systems.

The closest works to the problem formulation is the work by Devonport *et al.* [45] and the work of Xue *et al.* [212], where data-driven abstraction techniques are provided for satisfying finite-horizon specifications. The results of this chapter are more general and provide stronger guarantees in two main aspects. First, the constructed abstraction can be used for synthesising a controller against any linear temporal logic (LTL) specification and is not restricted to a fragment of LTL specifications. The sample complexity result is independent of the horizon of the specification and does not limit using the approach on finite-horizon specifications. Second, the guarantee provided by Devonport *et al.* and by Xue *et al.* are based on PAC bounds, which means the constructed abstraction is always wrong on a small subset of the state space whose size can be made smaller at the cost of high computational efforts, and the approach will require infinite number of samples if the size of this subset is set to zero. The formulated guarantee ensures that the abstraction is valid on the *entire* state space with high confidence (i.e., confidence close to 1). The confidence is specified by $(1 - \beta)$ in this chapter and is interpreted from the frequentist view of probability: if the algorithm is run multiple times, a correct abstraction is always generated except for a small number of times reflected in the confidence value. Having such a confidence value is essential in the approach since it relies on data gathered from the system. Smaller values of β gives higher confidence on getting a correct abstraction. This in turn increases the computational complexity of the approach since β appears directly in the sample complexity results.

In this approach, the synthesis problem is formulated as a robust convex program and approximated with a scenario program. Such approximations have been studied for the past two decades. Calafiore and Campi [32] provide an approximately feasible solution for the associated chance-constrained program by solving a scenario program, and give a sample complexity result. Relaxing the convexity assumption is studied by Soudjani and Majumdar [179] by assuming additional properties of the underlying probability distributions. The results by Esfahani *et al.* [51] will be used, where the optimality of the robust program is connected directly to the scenario program for performing data-driven verification and

synthesis. Inspired by the works of Wood and Zhang [204], and Weng *et al.* [202], this chapter will use extreme value theory to estimate the Lipschitz constant needed for the sample complexity results.

6.2 Preliminaries and Problem Statement

6.2.1 Preliminaries

Control Systems. Similar to Def. 4.1, a continuous-time control system is a tuple $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$, where $X \subset \mathbb{R}^n$ is the state space, $\mathbf{x}_{\text{in}} \in X_0 \subset X$ is the initial state, $U \subset \mathbb{R}^m$ is the input space, and $V \subset \mathbb{R}^q$ is the disturbance space which is assumed to be a compact set containing the origin. The vector field $g : X \times U \rightarrow X$ is such that $g(\cdot, \mathbf{u})$ is locally Lipschitz for all $\mathbf{u} \in U$. The evolution of the state of Σ is characterised by the differential equation

$$\dot{\mathbf{x}}(t) = g(\mathbf{x}(t), \mathbf{u}(t)) + \mathbf{v}(t), \quad (6.1)$$

where $\mathbf{v}(t) \in V$ represents the additive disturbance.

Consider the class of input and disturbance signals $\mathbf{u} : \mathbb{R}_{\geq 0} \rightarrow U$ and $\mathbf{v} : \mathbb{R}_{\geq 0} \rightarrow V$ to be piecewise constant with respect to a *sampling time* $\tau > 0$, i.e., $\mathbf{u}(t) = \mathbf{u}(k\tau)$ and $\mathbf{v}(t) = \mathbf{v}(k\tau)$ for every $k\tau \leq t < (k+1)\tau$ and $k \in \mathbb{N}_{\geq 0}$.

Trajectories of Control Systems. This chapter considers control systems $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$ whose vector fields g are not known, but whose *time-sampled trajectories* can be observed. In order to define time-sampled trajectories, it is necessary to first define *continuous-time trajectories* of control systems. Given a sampling time $\tau > 0$, an initial state $\mathbf{x}_0 \in X$, a constant input $\mathbf{u} \in U$, and a constant disturbance $\mathbf{v} \in V$, define the continuous-time trajectory $\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}$ of the system on the time interval $[0, \tau]$ as an absolutely continuous function $\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}} : [0, \tau] \rightarrow X$ such that $\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}(0) = \mathbf{x}_0$, and $\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}$ satisfies the differential equation $\dot{\zeta}_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}(t) = g(\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}(t), \mathbf{u}) + \mathbf{v}$ for almost all $t \in [0, \tau]$. The solution of (6.1) from \mathbf{x}_0 for the constant control input \mathbf{u} with $\mathbf{v}(t) = 0$ for all $t \geq 0$ is called the *nominal trajectory* of the system. For a fixed τ , the operators are defined as

$$\begin{aligned} \varphi(\mathbf{x}, \mathbf{u}, \mathbf{v}) &:= \zeta_{\mathbf{x}, \mathbf{u}, \mathbf{v}}(\tau) \text{ and} \\ \Phi(\mathbf{x}, \mathbf{u}) &:= \{ \varphi(\mathbf{x}, \mathbf{u}, \mathbf{v}) \mid \mathbf{v} \in V \} \end{aligned}$$

respectively for the trajectory at time τ and the set of such trajectories starting from \mathbf{x} . A sequence $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ is a time-sampled trajectory of Σ if for each $i \geq 0$, then $\mathbf{x}_{i+1} \in \Phi(\mathbf{x}_i, \mathbf{u}_i)$ for some $\mathbf{u}_i \in U$.

Remark 6.1. *The effect of general disturbance signals $\{\mathbf{v} : \mathbb{R}_{\geq 0} \rightarrow V'\}$ that are not necessarily piecewise constant, on the time-sampled trajectory $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ can be approximated by a set of piecewise constant disturbances $\{\mathbf{v} : \mathbb{R}_{\geq 0} \rightarrow V\}$ such that $V' \subset V$. The derivation of V from V' depends on the sampling time τ and the continuity properties of the system.*

Note that selecting a sampling time τ and considering time-sampled trajectories is essential for defining linear temporal specifications in their full generality, as defined next.

LTL Specifications. The control tasks are defined using Linear Temporal Logic (LTL). From Def. 4.10, LTL specifications are considered with syntax

$$\psi := \top \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \bigcirc\psi \mid \psi_1 \mathbf{U} \psi_2,$$

where $p \in \mathbb{R}^n$ is an element of the set of atomic propositions AP . The reader is referred back to Chapter 4 for further details on temporal logic specifications.

Feedback Controller. A *feedback controller* for Σ is a function $\mathfrak{C}: X \rightarrow U$. The feedback composition of Σ and \mathfrak{C} is denoted by $\mathfrak{C} \parallel \Sigma$. The set of trajectories of the closed-loop system $\mathfrak{C} \parallel \Sigma$ consists of all finite trajectories $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots$ such that for all $i \in \mathbb{N}_{\geq 0}$, then $\mathbf{x}_{i+1} \in \Phi(\mathbf{x}_i, \mathfrak{C}(\mathbf{x}_i))$.

Finite Abstraction of Control Systems. Let $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$ be a control system, $\tau > 0$ be a fixed sampling time, and $\eta_x \in \mathbb{R}_{>0}^n$ and $\eta_u \in \mathbb{R}_{>0}^m$ denote the discretisation parameters for state and input spaces, respectively. Let $\hat{X} \subset X$ be a finite set of points that are equally spaced with respect to the state space discretisation parameter η_x , and \hat{U} be a finite subset of equally spaced points the state input discretisation parameter η_u . A *finite-state abstraction* of Σ is denoted by a tuple $\hat{\Sigma} = (\hat{X}, \hat{U}, \hat{g})$, where $\hat{\mathbf{x}}'$ in $\hat{g}(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ if there is a pair of states $\mathbf{x} \in \Omega_{\eta_x}(\hat{\mathbf{x}})$ and $\mathbf{x}' \in \Omega_{\eta_x}(\hat{\mathbf{x}}')$ such that $\mathbf{x}' \in \Phi(\mathbf{x}, \hat{\mathbf{u}})$, similar to Def. 4.8. Note that, the larger η_x is (where comparison is made dimension-wise), the smaller is the cardinality of \hat{X} resulting in a coarser abstraction. On the other hand, the smaller η_x is then the more precise the abstraction $\hat{\Sigma}$ will be, increasing the chance of a successful controller synthesis (e.g., [188] for more details on this construction).

Feedback Refinement Relation. Let Σ be a control system and $\hat{\Sigma}$ be its finite-state abstraction. A *feedback refinement relation* (FRR) from Σ to $\hat{\Sigma}$ is a relation $\mathcal{Q} \subseteq X \times \hat{X}$ s.t. for all $\mathbf{x} \in X$ there is some $\hat{\mathbf{x}} \in \hat{X}$ such that $\mathcal{Q}(\mathbf{x}, \hat{\mathbf{x}})$ and for all $(\mathbf{x}, \hat{\mathbf{x}}) \in \mathcal{Q}$, then (i) $\hat{U} \subseteq U$, and (ii) $\mathbf{u} \in \hat{U} \Rightarrow \mathcal{Q}(g(\mathbf{x}, \mathbf{u})) \subseteq \hat{g}(\hat{\mathbf{x}}, \mathbf{u})$. It is written $\Sigma \preceq_{\mathcal{Q}} \hat{\Sigma}$ if \mathcal{Q} is an FRR from Σ to $\hat{\Sigma}$.

Abstraction-based Controller Synthesis. The synthesis objective is expressed as LTL specifications. The abstraction-based controller design (ABCD) [153] is a 3-step method to find a robust controller for the control system Σ . First, a finite state abstraction $\hat{\Sigma}$ is computed s.t. $\Sigma \preceq_{\mathcal{Q}} \hat{\Sigma}$. Second, an abstract controller is synthesised of the form $\hat{\mathfrak{C}}: \hat{X} \rightarrow \hat{U}$ for $\hat{\Sigma}$ using methods from the reactive synthesis literature. Finally, the desired controller \mathfrak{C} is obtained as $\mathfrak{C} := \hat{\mathfrak{C}} \circ \mathcal{Q}$. It is known that this three step process produces a controller \mathfrak{C} such that $\mathfrak{C} \parallel \Sigma$ satisfies the specification [153]. For the details of the tool implementation using abstract models see [163].

6.2.2 Problem Statement

Abstraction-based control design (ABCD) for systems with *unknown* dynamics are studied using available data from the system such that a given specification is satisfied with high confidence on the closed-loop system.

Assumption 6.1. *The vector field g of the control system $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$ is unknown, but sampled trajectories of the system can be obtained in the form of $\mathbb{S}_N := \{(\mathbf{x}_k, \mathbf{u}_k, \mathbf{x}'_k) \mid \mathbf{x}'_k \in \Phi(\mathbf{x}_k, \mathbf{u}_k), k = 1, 2, \dots, N\}$.*

Problem Description 6.1 (Data-driven ABCD). *Inputs:* Control system $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$ with unknown vector field g , specification ψ , sampled trajectories \mathbb{S}_N , and confidence parameter $\beta \in (0, 1)$.

Outputs: Abstract model $\widehat{\Sigma}$, abstract controller $\widehat{\mathfrak{C}}$, and refined controller \mathfrak{C} for Σ , such that $\mathfrak{C} \parallel \Sigma$ satisfies ψ with confidence $(1 - \beta)$.

The first step of the ABCD is to compute a finite abstraction $\widehat{\Sigma}$ for Σ . Once such an abstraction is computed, synthesis of the controller $\widehat{\mathfrak{C}}$ and refining it to \mathfrak{C} follows the model-based ABCD scheme. Therefore, the main challenge is to provide a data-driven computation of the abstraction $\widehat{\Sigma}$ that is a true overapproximation of Σ with confidence $(1 - \beta)$.

Problem Description 6.2 (Data-driven Abstraction). *Inputs:* Control system $\Sigma = (X, \mathbf{x}_{\text{in}}, U, V, g)$ with unknown vector field g , sampled trajectories \mathbb{S}_N , discretisation parameters η_x and η_u , and confidence parameter $\beta \in (0, 1)$.

Outputs: Finite model $\widehat{\Sigma}$ that is an abstraction of Σ with confidence $(1 - \beta)$.

In this chapter, Problem 6.2 is tackled by showing how to construct $\widehat{\Sigma}$ from sampled trajectories \mathbb{S}_N , and providing a lower bound on the data size N in order to ensure correctness of the abstraction with confidence $(1 - \beta)$. The required theoretical concepts are presented in the next section.

6.3 Robust Convex Programs

This section describes *robust convex programs* (RCPs) and the data-driven approximation of their solution. In Sections 6.4 and 6.5, it is shown how such an approximation can be used for solving the data-driven abstraction in Problem 6.2.

Let $\mathcal{T} \subset \mathbb{R}^q$ be a compact convex set for some $q \in \mathbb{N}$ and $c \in \mathbb{R}^q$ be a constant vector. Let $(\mathcal{D}, \mathcal{F}_{\mathcal{D}}, \mathbb{P}_{\mathcal{D}})$ be the probability space of the *uncertainty* and $\Lambda: \mathcal{T} \times \mathcal{D} \rightarrow \mathbb{R}$ be a measurable function, which is convex in the first argument for each $d \in \mathcal{D}$, and bounded in the second argument for each $\theta \in \mathcal{T}$. The RCP is defined as

$$\text{RCP: } \begin{cases} \min_{\theta} c^{\top} \theta \\ \text{s.t. } \theta \in \mathcal{T} \text{ and } \Lambda(\theta, d) \leq 0 \quad \forall d \in \mathcal{D}. \end{cases} \quad (6.2)$$

Computationally tractable approximations of the optimal solution of the RCP (6.2) can be obtained using *scenario convex programs* (SCPs) that only require gathering finitely many samples from the uncertainty space [51]. Let $(d_i)_{i=1}^N$ be N independent and identically distributed (i.i.d.) samples drawn according to the probability measure $\mathbb{P}_{\mathcal{D}}$. The SCP corresponding to the RCP (6.2) strengthened with $\gamma \geq 0$ is defined as

$$SCP_{\gamma} : \begin{cases} \min_{\theta} c^{\top} \theta \\ \text{s.t. } \theta \in \mathcal{T}, \text{ and } \Lambda(\theta, d_i) + \gamma \leq 0 \quad \forall i \in \{1, 2, \dots, N\}. \end{cases} \quad (6.3)$$

The optimal solution of RCP (6.2) is denoted as θ_{RCP}^* and the optimal solution of SCP_{γ} (6.3) as θ_{SCP}^* . Note that θ_{RCP}^* is a single deterministic quantity but θ_{SCP}^* is a random quantity that depends on the i.i.d. samples $(d_i)_{i=1}^N$ drawn according to $\mathbb{P}_{\mathcal{D}}$. The RCP (6.2) is a challenging optimisation problem since the cardinality of \mathcal{D} is infinite and the optimisation has infinite number of constraints. In contrast, the SCP (6.3) is a convex optimisation with finite number of constraints for which efficient optimisation techniques are available [30]. The following theorem provides a sample complexity result for connecting the optimal solution of the SCP_{γ} to that of the RCP.

Theorem 6.1 ([51]). *Assume that the mapping $d \mapsto \Lambda(\theta, d)$ in (6.2) is Lipschitz continuous uniformly in $\theta \in \mathcal{T}$ with Lipschitz constant L_d and let $\Upsilon: [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ be a strictly increasing function such that*

$$\mathbb{P}_{\mathcal{D}}(\Omega_{\varepsilon}(d)) \geq \Upsilon(\varepsilon), \quad (6.4)$$

for every $d \in \mathcal{D}$ and $\varepsilon \in [0, 1]$. Let θ_{RCP}^* be the optimal solution of the RCP (6.2) and θ_{SCP}^* the optimal solution of SCP_{γ} (6.3) with

$$\gamma = L_d \Upsilon^{-1}(\varepsilon) \quad (6.5)$$

computed by taking N i.i.d. samples $(d_i)_{i=1}^N$ from $\mathbb{P}_{\mathcal{D}}$. Then θ_{SCP}^* is a feasible solution for the RCP with confidence $(1 - \beta)$ if the number of samples $N \geq N(\varepsilon, \beta)$, where

$$N(\varepsilon, \beta) := \min \left\{ N \in \mathbb{N} \mid \sum_{i=0}^{q-1} \binom{N}{i} \varepsilon^i (1 - \varepsilon)^{N-i} \leq \beta \right\}, \quad (6.6)$$

with q being the dimension of the decision vector $\theta \in \mathcal{T}$.

6.4 Data-Driven Abstraction

This section first discusses the steps required for model-based abstraction of control systems. Then it shows how this can be formulated as an RCP and presents its associated SCP. Finally, the connection between the RCPs and SCPs is used in Theorem 6.1 to provide a lower bound for number of required samples to certify a desired confidence. The simplifying assumption used in this section is that samples from the *nominal trajectories* of the system Σ are also available in the form

of $\{(\mathbf{x}_k, \mathbf{u}_k, \mathbf{x}'_k) \mid \mathbf{x}'_k = \varphi(\mathbf{x}_k, \mathbf{u}_k, 0), k = 1, 2, \dots, N\}$. Discussed in the next section is how this assumption can be relaxed by modifying the inequality of the growth bound.

6.4.1 Growth Bound for Reachable Sets

Consider a control system $\Sigma = (X, \mathbf{x}_{in}, U, V, g)$ with the disturbance set $V = [-\bar{\mathbf{v}}, \bar{\mathbf{v}}]$ for some vector $\bar{\mathbf{v}} \in \mathbb{R}_{\geq 0}^n$. Let η_x and η_u be discretisation parameters for the state and input spaces X and U used to construct \widehat{X} and \widehat{U} of sizes n_x and n_u , respectively. The first step of ABCD is to compute a finite abstraction $\widehat{\Sigma} = (\widehat{X}, \widehat{U}, \widehat{g})$ using overapproximations of the reachable sets for every pair of abstract state and input. The reachable set for every pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \widehat{X} \times \widehat{U}$ is defined as

$$Reach(\hat{\mathbf{x}}, \hat{\mathbf{u}}) := \{\mathbf{x}' \in \Phi(\mathbf{x}, \hat{\mathbf{u}}) \mid \mathbf{x} \in \Omega_{\eta_x}(\hat{\mathbf{x}})\}.$$

The set $Reach(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ is usually overapproximated using a growth bound of the system dynamics [153].

Definition 6.1. For a control system Σ with abstract state and input spaces \widehat{X}, \widehat{U} , a function $\kappa: \mathbb{R}_{\geq 0}^n \times \widehat{X} \times \widehat{U} \rightarrow \mathbb{R}_{\geq 0}$ is called a growth bound function for Σ if it satisfies

$$\begin{aligned} |\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)| &\leq \kappa(|\mathbf{x} - \hat{\mathbf{x}}|, \hat{\mathbf{x}}, \hat{\mathbf{u}}) \\ \forall \hat{\mathbf{x}} \in \widehat{X}, \forall \hat{\mathbf{u}} \in \widehat{U}, \forall \mathbf{x} \in \Omega_{\eta_x}(\hat{\mathbf{x}}), \forall \mathbf{v} \in V. \end{aligned} \quad (6.7)$$

Note that $\varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)$ is the nominal (disturbance-free) trajectory of the system. Using this definition, for every abstract state-input pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \widehat{X} \times \widehat{U}$, the reachable set $Reach(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ is overapproximated with a ball centred at $z(\hat{\mathbf{x}}, \hat{\mathbf{u}}) := \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)$ with radius $\lambda(\hat{\mathbf{x}}, \hat{\mathbf{u}}) := \kappa(\eta_x, \hat{\mathbf{x}}, \hat{\mathbf{u}})$.

When the system dynamics are known, it is shown by Reissig *et al.* [153] that a growth bound for the system can be computed as

$$\kappa(\mathbf{r}, \hat{\mathbf{x}}, \hat{\mathbf{u}}) = e^{L(\hat{\mathbf{u}})\tau} \mathbf{r} + \int_0^\tau e^{L(\hat{\mathbf{u}})s} \bar{\mathbf{v}} ds, \quad (6.8)$$

for all $\mathbf{r} \in \mathbb{R}_{\geq 0}^n$, $\hat{\mathbf{x}} \in \widehat{X}$, and $\hat{\mathbf{u}} \in \widehat{U}$, where $\bar{\mathbf{v}}$ is the upper bound of the disturbance and $L: \widehat{U} \rightarrow \mathbb{R}^{n \times n}$ is a matrix such that the entries of $L(\hat{\mathbf{u}})$ satisfy the following inequality for all $\mathbf{x} \in X$:

$$L_{i,j}(\hat{\mathbf{u}}) \geq \begin{cases} D_j g_i(\mathbf{x}, \hat{\mathbf{u}}) & i = j \\ |D_j g_i(\mathbf{x}, \hat{\mathbf{u}})| & i \neq j, \end{cases} \quad (6.9)$$

for all $i, j \in \{1, 2, \dots, n\}$, where $g_i(\mathbf{x}, \mathbf{u})$ is the i^{th} element of the vector field $g(\mathbf{x}, \mathbf{u})$ and $D_j g_i$ is its partial derivative with respect to the j^{th} element of \mathbf{x} .

6.4.2 SCP for the Computation of Growth Bound

When the model of the system is unknown, the growth bound in (6.8) is not available since the matrix $L(\hat{\mathbf{u}})$ defined using (6.9) is not computable. To tackle this bottleneck, the aim is to compute a growth bound for the system that has the following parameterised form

$$\kappa_\theta(\mathbf{r}, \hat{\mathbf{x}}, \hat{\mathbf{u}}) := \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\mathbf{r} + \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}}), \forall \mathbf{r} \in \mathbb{R}_{\geq 0}^n, \hat{\mathbf{x}} \in \hat{X}, \hat{\mathbf{u}} \in \hat{U}, \quad (6.10)$$

where $\theta_1 \in \mathbb{R}^{n \times n}$ and $\theta_2 \in \mathbb{R}^n$. The concatenation of columns of θ_1 and θ_2 is denoted by $\theta \in \mathbb{R}^{n^2+n}$.

Remark 6.2. *The parameterised growth bound in (6.10) is linear with respect to \mathbf{r} similar to (6.8), but is more general and less conservative by allowing θ_1, θ_2 to depend on $\hat{\mathbf{x}}$ (i.e., they are defined locally for each abstract state).*

Theorem 6.2. *The parameterised growth bound in (6.10) can be computed by solving the following robust convex program*

$$\begin{cases} \min_\theta c^\top \theta \\ \text{s.t. } 0 \leq \theta \leq \bar{\theta}, \text{ and } \forall \mathbf{x} \in \Omega_{\eta_x}(\hat{\mathbf{x}}), \forall \mathbf{v} \in V, \\ |\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)| - \kappa_\theta(|\mathbf{x} - \hat{\mathbf{x}}|, \hat{\mathbf{x}}, \hat{\mathbf{u}}) \leq 0, \end{cases} \quad (6.11)$$

where $c = [1, 1, \dots, 1] \in \mathbb{R}^{n^2+n}$ and $\bar{\theta}$ is a sufficiently large positive vector.

Proof. The inequality in (6.11) is a reformulation of (6.7) with a special choice of κ in (6.10). Therefore, the optimisation (6.11) is in fact a robust convex program. Let $\mathcal{D} = \Omega_{\eta_x}(\hat{\mathbf{x}}) \times V$ be the uncertainty space and

$$\Lambda(\theta, \mathbf{x}, \mathbf{v}) := |\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)| - \kappa_\theta(|\mathbf{x} - \hat{\mathbf{x}}|, \hat{\mathbf{x}}, \hat{\mathbf{u}})$$

for all $\mathbf{x} \in \Omega_{\eta_x}(\hat{\mathbf{x}})$ and $\mathbf{v} \in V$ and fixed $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \hat{X} \times \hat{U}$. It is necessary to show that Λ is convex in θ for each $(\mathbf{x}, \mathbf{v}) \in \mathcal{D}$ and bounded in (\mathbf{x}, \mathbf{v}) for every $\theta \in [0, \bar{\theta}]$. The convexity holds due to the parameterisation of κ_θ in (6.10) being linear with respect to the optimisation variables in θ . The boundedness holds due to the set \mathcal{D} being compact and trajectories of the system being continuous. \square

To construct the SCP_γ associated with the RCP (6.11), fix $\hat{\mathbf{x}} \in \hat{X}$ and $\hat{\mathbf{u}} \in \hat{U}$, consider a uniform distribution on the space $\mathcal{D} = \Omega_{\eta_x}(\hat{\mathbf{x}}) \times V$ and obtain N i.i.d. sample trajectories $\mathbb{S}_N = \{(\mathbf{x}_i, \hat{\mathbf{u}}, \mathbf{x}'_i) \mid \mathbf{x}'_i \in \Phi(\mathbf{x}_i, \hat{\mathbf{u}}), i = 1, 2, \dots, N\}$. Note that every \mathbf{x}'_i corresponds to a random disturbance $\mathbf{v}_i \in V$. The SCP_γ is

$$\begin{cases} \min_\theta c^\top \theta \\ \text{s.t. } 0 \leq \theta \leq \bar{\theta} \text{ and } \forall i \in \{1, \dots, N\}, \\ |\mathbf{x}'_i - \mathbf{x}'_{nom}| - \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})|\mathbf{x}_i - \hat{\mathbf{x}}| + \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}}) + \gamma \leq 0, \end{cases} \quad (6.12)$$

where $\mathbf{x}'_{nom} := \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)$ and $\gamma \in \mathbb{R}_{\geq 0}$.

Assumption 6.2. Let $L_\varphi(\hat{\mathbf{u}})$ be an upper bound for the Lipschitz constant of the system trajectories $\varphi(x, \hat{\mathbf{u}}, w)$ with respect to (x, w) , i.e., for all $x, x' \in \Omega_{\eta_x}(\hat{\mathbf{x}})$ and $w, w' \in W$,

$$\|\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{w}) - \varphi(\mathbf{x}', \hat{\mathbf{u}}, \mathbf{w}')\| \leq L_\varphi(\hat{\mathbf{u}})\|(\mathbf{x}, \mathbf{w}) - (\mathbf{x}', \mathbf{w}')\|. \quad (6.13)$$

We assume that a possibly conservative correct upper bound $L_\varphi(\hat{\mathbf{u}})$ is known.

Theorem 6.3. Let $|\hat{X}| = n_x$ and $|\hat{U}| = n_u$. Under Assumption 6.2, for any $\hat{\mathbf{x}} \in \hat{X}$ constructed with discretisation size η_x , any $\hat{\mathbf{u}} \in \hat{U}$, and the disturbance set $V = [-\bar{\mathbf{v}}, \bar{\mathbf{v}}]$, the optimal solution of (6.12) gives a growth bound for the system Σ corresponding to $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ with confidence $(1 - \beta/(n_x n_u))$, when the number of samples $N \geq N(\varepsilon, \beta/(n_x n_u))$ and

$$\gamma = 4L_\varphi(\hat{\mathbf{u}}) \sqrt[2n]{\varepsilon \prod_{i=1}^n \eta_x(i) \prod_{i=1}^n \bar{\mathbf{v}}(i)}, \quad (6.14)$$

where $\varepsilon \in [0, 1]$, n is the dimension of the state space.

Proof. Apply Theorem 6.2 to the RCP (6.11) for fixed $\hat{\mathbf{x}} \in \hat{X}$ and $\hat{\mathbf{u}} \in \hat{U}$. Define

$$\Lambda(\theta, \mathbf{x}, \mathbf{v}) := \max\{|\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\hat{\mathbf{x}}, \hat{\mathbf{u}}, 0)| - \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})|\mathbf{x} - \hat{\mathbf{x}}| - \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}})\}, \quad (6.15)$$

where the $\max\{\cdot\}$ is applied to the elements of its argument that belongs to \mathbb{R}^n . Since the distribution on $\mathcal{D} = \Omega_{\eta_x}(\hat{\mathbf{x}}) \times V$ is uniform, $\Upsilon(\varepsilon) = \mathbb{P}_{\mathcal{D}}(\Omega_\varepsilon(d)) = \frac{(\varepsilon/2)^{2n}}{\prod_{i=1}^n \eta_x(i) \prod_{i=1}^n \bar{\mathbf{v}}(i)}$ is chosen to satisfy the inequality (6.4). Note that $\Upsilon(\varepsilon)$ gives the probability of choosing a point within the $2n$ -ball $\Omega_\varepsilon(d)$ uniformly at random. Equation (6.5) is used as $\gamma = L_d \Upsilon^{-1}(\varepsilon)$ to get the value of γ in (6.14). It only remains to show that $\Lambda(\theta, \mathbf{x}, \mathbf{v})$ is Lipschitz continuous with constant $L_d = 2L_\varphi(\hat{\mathbf{u}})$. Note that $L_\varphi(\hat{\mathbf{u}})$ is the upper bound on the Lipschitz constant of $\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v})$ with respect to (\mathbf{x}, \mathbf{v}) , and satisfies (6.13) for all $\mathbf{x}, \mathbf{x}' \in \Omega_{\eta_x}(\hat{\mathbf{x}})$ and $\mathbf{w}, \mathbf{w}' \in W$. Since $\|\theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\|$ can be bounded by $L_\varphi(\hat{\mathbf{u}})$, it follows that

$$\begin{aligned} & \|\Lambda(\theta, \mathbf{x}, \mathbf{v}) - \Lambda(\theta, \mathbf{x}', \mathbf{v}')\| \\ & \leq \|\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\mathbf{x}', \hat{\mathbf{u}}, \mathbf{v}')\| + \|\theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\|\|\mathbf{x} - \mathbf{x}'\| \\ & \leq L_\varphi(\hat{\mathbf{u}})\|(\mathbf{x}, \mathbf{v}) - (\mathbf{x}', \mathbf{v}')\| + L_\varphi(\hat{\mathbf{u}})\|\mathbf{x} - \mathbf{x}'\| \\ & \leq 2L_\varphi(\hat{\mathbf{u}})\|(\mathbf{x}, \mathbf{v}) - (\mathbf{x}', \mathbf{v}')\|, \end{aligned}$$

Therefore, $\Lambda(\theta, \mathbf{x}, \mathbf{v})$ is Lipschitz continuous with constant $2L_\varphi(\hat{\mathbf{u}})$. This completes the proof. \square

Remark 6.3. The statement of Theorem 6.3 holds under Assumption 6.2 that requires knowing a correct (possibly conservative) upper bound $L_\varphi(\hat{\mathbf{u}})$ on the Lipschitz constant of the system trajectories. To compensate for conservative values of $L_\varphi(\hat{\mathbf{u}})$, smaller values of ε is chosen, which will require taking higher number of samples N .

Remark 6.4. An algorithm is provided in the next subsection for estimating L_φ using sampled trajectories of the system. The approach of the next subsection gives only an

“estimate” of L_φ . The results of Theorem 6.3 remains valid under the assumption that such estimation methods return a correct upper bound for L_φ .

Corollary 6.1. *Assume the SCP_γ is solved, under Assumption 6.2, for each state-input pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \hat{X} \times \hat{U}$ with confidence $(1 - \beta)/(n_x n_u)$. Then the abstract model constructed using the obtained growth bounds is a valid abstract model for Σ with confidence at least $(1 - \beta)$.*

Proof. Denote the optimal solution of SCP_γ in (6.12) by θ^* . The ball centred at $z(\hat{\mathbf{x}}, \hat{\mathbf{u}}) := \mathbf{x}'_{nom}$ with radius $\lambda(\hat{\mathbf{x}}, \hat{\mathbf{u}}) = \kappa_{\theta^*}(\eta_x, \hat{\mathbf{x}}, \hat{\mathbf{u}}) + \gamma$ is a valid overapproximation of the reachable set from the state-input pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ with confidence at least $1 - \beta/(n_x n_u)$. Since the number of pairs $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ is $n_x n_u$, the chance of getting an invalid growth bound in at least one instance of SCP_γ is bounded by β . Therefore, a sound abstraction is acquired that truly overapproximates the behaviour of the system with confidence $(1 - \beta)$. \square

Remark 6.5. *The parameter $\varepsilon \in [0, 1]$ gives a trade off between the required number of samples and the level of conservativeness applied to the SCP. Smaller ε results in a larger number of sample trajectories, but reduces the value of γ in (6.14) (less conservative constraints in the SCP and higher chance of finding a feasible solution). In contrast, larger ε results in a smaller number of sample trajectories but increases the value of γ .*

Remark 6.6. *The quantity $2n$ used in (6.14) is in fact the dimension of the sample space $\mathcal{D} = \Omega_{\eta_x}(\hat{\mathbf{x}}) \times V$. If the system does not have any disturbance (i.e., the system can be modelled as an ODE having deterministic trajectories), the sample space will be $\mathcal{D} = \Omega_{\eta_x}(\hat{\mathbf{x}})$ and its dimension n can be used in (6.14): $\gamma = 4L_\varphi(\hat{\mathbf{u}}) \sqrt[n]{\varepsilon \prod_{i=1}^n \eta_x(i)}$. This will substantially reduce the number of required sample trajectories. Similarly, if the disturbance does not affect some of the state equations, $2n$ can be replaced by $(n + q)$ where q is the dimension of the disturbance set considered as a non-zero measure set.*

Algorithm 1 uses the result of Corollary 6.1 to provide an algorithmic solution for Problem 6.2. This algorithm receives a confidence parameters β , divides it by the cardinality of $\hat{X} \times \hat{U}$ (i.e., $n_x n_u$), computes the growth bounds for each pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \hat{X} \times \hat{U}$ using the SCP_γ in (6.12) with confidence $1 - \beta/(n_x n_u)$, and constructs the abstraction using these growth bounds.

The finite abstraction $\hat{\Sigma}$ constructed by Algorithm 1 is a valid abstraction for Σ with confidence $(1 - \beta)$. This means any controller $\hat{\mathbf{C}}$ synthesised on $\hat{\Sigma}$ and refined to a controller \mathbf{C} for Σ will satisfy the desired specification with confidence $(1 - \beta)$ on the closed loop system $\Sigma \parallel C$. In the next section, the approach is extended to make it suitable for abstraction refinement in the case that there is no controller $\hat{\mathbf{C}}$ satisfying the specification due to the conservatism of the approach.

6.4.3 Lipschitz Constant Estimation

For estimating L_φ in (6.13), it is desired to find an estimate of an upper bound for the fraction

$$\Delta(\hat{\mathbf{u}}) := \frac{\|\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v}) - \varphi(\mathbf{x}', \hat{\mathbf{u}}, \mathbf{v}')\|}{\|(\mathbf{x}, \mathbf{v}) - (\mathbf{x}', \mathbf{v}')\|}$$

Algorithm 1: Data-Driven Abstraction

Data: (X, U, V) of a control system Σ , confidence β , discretisation parameters η_x, η_u

- 1 Compute the finite state and input sets \hat{X} and \hat{U} using η_x, η_u ;
- 2 Define n_x and n_u as cardinalities of \hat{X} and \hat{U} ;
- 3 Choose $\varepsilon \in [0, 1]$;
- 4 Set $N = N(\varepsilon, \frac{\beta}{n_x n_u})$ using Eq. (6.6);
- 5 Compute γ using Eq. (6.14);
- 6 **for** $\hat{x} \in \hat{X}$ **do**
- 7 **for** $\hat{u} \in \hat{U}$ **do**
- 8 $\hat{g}(\hat{x}, \hat{u}) = \emptyset$;
- 9 Consider the uncertainty space $\mathcal{D} = \Omega_{\eta_x}(\hat{x}) \times V$;
- 10 Select N i.i.d sample trajectories using uniform distribution over \mathcal{D} ;
- 11 Simulate the nominal trajectory $(\hat{x}, \hat{u}, \mathbf{x}'_{nom})$;
- 12 Solve the SCP $_{\gamma}$ (6.12) to get the optimiser $\theta^*(\hat{x}, \hat{u})$;
- 13 $z \leftarrow \mathbf{x}'_{nom}$;
- 14 $\lambda \leftarrow \kappa_{\theta^*}(\eta_x, \hat{x}, \hat{u}) + \gamma$;
- 15 Find all states $\hat{x}' \in \hat{X}$ for which $\Omega_{\eta_x}(\hat{x}') \cap \Omega_{\lambda}(z) \neq \emptyset$ and add them to $\hat{g}(\hat{x}, \hat{u})$;
- 16 **end**
- 17 **end**

Result: $\hat{\Sigma} = (\hat{X}, \hat{U}, \hat{g})$ as a finite abstraction of Σ with confidence $(1 - \beta)$, $\theta^*(\hat{x}, \hat{u})$ as a growth bound for $\hat{x} \in \hat{X}, \hat{u} \in \hat{U}$

that holds for all $\mathbf{x}, \mathbf{x}' \in X$ and $\mathbf{v}, \mathbf{v}' \in V$. Following the line of reasoning in [202, 204], extreme value theory is used for the estimation.

Let us fix a $\delta > 0$ and assign uniform distribution to the pairs (\mathbf{x}, \mathbf{v}) and $(\mathbf{x}', \mathbf{v}')$ over the domain

$$\{\mathbf{x}, \mathbf{x}' \in X, \mathbf{v}, \mathbf{v}' \in V \text{ with } \|(\mathbf{x}, \mathbf{v}) - (\mathbf{x}', \mathbf{v}')\| \leq \delta\}. \quad (6.16)$$

Then $\Delta(\hat{\mathbf{u}})$ is a random variable with an unknown cumulative distribution function (CDF). Based on the assumption of Lipschitz continuity of the system, the support of the distribution of $\Delta(\hat{\mathbf{u}})$ is bounded from above, and it is wanted to estimate an upper bound for its support. n sample pairs (\mathbf{x}, \mathbf{v}) and $(\mathbf{x}', \mathbf{v}')$ should be taken, to compute n samples $\Delta_1, \Delta_2, \dots, \Delta_n$ for $\Delta(\hat{\mathbf{u}})$. The CDF of $\max\{\Delta_1, \Delta_2, \dots, \Delta_n\}$ is called the limit distribution of $\Delta(\hat{\mathbf{u}})$. Fisher-Tippett-Gnedenko theorem [71, 130] says that if the limit distribution exists, it can only be one of the three family of extreme value distributions – the Gumbel class, the Fréchet class, and the reverse Weibull class. These CDF's have the following forms:

$$\text{Gumbel class: } G(s) = \exp \left[-\exp \left[\frac{s-a}{b} \right] \right], s \in \mathbb{R}$$

$$\text{Fréchet class: } G(s) = \begin{cases} 0 & \text{if } s < a \\ \exp \left[-\left[\frac{s-a}{b} \right]^{-c} \right] & \text{if } s \leq a \end{cases}$$

$$\text{Reverse Weibull class: } G(s) = \begin{cases} \exp \left[-\left[\frac{a-s}{b} \right]^c \right] & \text{if } s < a \\ 1 & \text{if } s \leq a \end{cases}$$

where $a \in \mathbb{R}, b > 0, c > 0$ are respectively the location, scale and shape parameters of the distributions.

Among the above three distributions, only the reverse Weibull class has a support bounded from above. Therefore, the limit distribution of $\Delta(\hat{\mathbf{u}})$ will be from this class and the location parameter a is such an upper bound. As a result, the location parameter of the limit distribution of $\Delta(\hat{\mathbf{u}})$ can be estimated to get an estimation of the Lipschitz constant.

The approach is summarised in Algorithm 2. The most inner loop computes samples of $\Delta(\hat{\mathbf{u}})$. The middle loop computes samples of $\max\{\Delta_1, \dots, \Delta_n\}$. The outer loop estimates the Lipschitz constant for each $\hat{\mathbf{u}}$ by fitting a reverse Weibull distribution.

Remark 6.7. *The approach presented in this section can only be used for estimating the Lipschitz constant, which can then be enlarged by a factor greater than one to account for the effect of estimation using finite number of samples. Note that this factor can be selected depending on the system under study by fitting the Reverse Weibull distribution to datasets of varying size and observing its convergence behaviour. The results are valid under the assumption that the (corrected) estimation provided by Algorithm 6.2 gives a valid upper bound on the Lipschitz constant of the system.*

Algorithm 2: Lipschitz Constant Estimation

Data: (X, U, V) of a control system Σ , abstract input space \widehat{U}

- 1 Select number of samples n and m for the estimation
- 2 Select $\delta > 0$
- 3 **for** $\hat{u} \in \widehat{U}$ **do**
- 4 **for** $j = 1 : m$ **do**
- 5 **for** $i = 1 : n$ **do**
- 6 Sample pairs $(\mathbf{x}, \mathbf{v}), (\mathbf{x}', \mathbf{v}')$ uniformly from the domain in (6.16)
- 7 Run Σ to get trajectories $\varphi(\mathbf{x}, \hat{u}, \mathbf{v})$ and $\varphi(\mathbf{x}', \hat{u}, \mathbf{v}')$
- 8 Compute $\Delta_i := \frac{\|\varphi(\mathbf{x}, \hat{u}, \mathbf{v}) - \varphi(\mathbf{x}', \hat{u}, \mathbf{v}')\|}{\|(\mathbf{x}, \mathbf{v}) - (\mathbf{x}', \mathbf{v}')\|}$
- 9 **end**
- 10 $\Gamma_j := \max\{\Delta_1, \dots, \Delta_n\}$
- 11 **end**
- 12 Fit a reverse Weibull distribution to the sample set $\{\Gamma_1, \Gamma_2, \dots, \Gamma_m\}$
- 13 $L_\varphi(\hat{u})$ is the location parameter of the fitted distribution
- 14 **end**

Result: Estimated value of $L_\varphi(\hat{u})$ for all $\hat{u} \in \widehat{U}$

6.5 Synthesis via Abstraction Refinement

The data-driven synthesis discussed in Section 6.4 inherits the soundness property from the ABCD approach: they both work with overapproximations of the dynamics and may not return a controller despite one may exist. Therefore, there is a need for refining the abstraction in order to check for controllers using less conservative abstractions. While the method of Section 6.4 is good for a given fixed discretisation parameter η_x , it is not suitable for reducing η_x , which requires re-computing all local parameters of the growth bounds $\theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}}), \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}})$. Another shortcoming of the method is related to the data collection: the nominal trajectories of the system should be available and are used in the constraints of the SCP. This section discusses an extension of the approach of Section 6.4, in order to

- enable reducing η_x without the need for re-computing a growth bound, and
- relaxing the assumption of having access to the nominal trajectories of the system.

Let us define a modified growth bound as a function $\kappa^e : \mathbb{R}_{\geq 0}^n \times \widehat{X} \times \widehat{U} \rightarrow \mathbb{R}_{\geq 0}^n$ that is strictly increasing in its first argument and satisfies

$$\begin{aligned}
 |\varphi(\mathbf{x}_1, \hat{\mathbf{u}}, \mathbf{v}_1) - \varphi(\mathbf{x}_2, \hat{\mathbf{u}}, \mathbf{v}_2)| &\leq \kappa^e(|\mathbf{x}_1 - \mathbf{x}_2|, \hat{\mathbf{x}}, \hat{\mathbf{u}}) \\
 \forall \hat{\mathbf{x}} \in \widehat{X}, \forall \hat{\mathbf{u}} \in \widehat{U}, \forall \mathbf{x}_1, \mathbf{x}_2 \in \Omega_{\eta_x}(\hat{\mathbf{x}}), \forall \mathbf{v}_1, \mathbf{v}_2 \in V.
 \end{aligned} \tag{6.17}$$

This definition is more conservative than (6.7) in comparing trajectories under two arbitrary disturbances, and κ^e always satisfies (6.7). Using this new definition, for

every pair of abstract state and input $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$, the corresponding overapproximation of the reach set can be computed as a ball centred at *any* $z(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \Phi(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ with radius $\lambda(\hat{\mathbf{x}}, \hat{\mathbf{u}}) = \kappa^e(\eta_x, \hat{\mathbf{x}}, \hat{\mathbf{u}})$.

A parametrisation for κ^e is chosen similar to (6.10), i.e.,

$$\kappa_{\theta}^e(\mathbf{r}, \hat{\mathbf{x}}, \hat{\mathbf{u}}) = \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\mathbf{r} + \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}}), \quad (6.18)$$

where $\mathbf{r} \in \mathbb{R}_{\geq 0}$, $\theta_1 \in \mathbb{R}^{n \times n}$, $\theta_2 \in \mathbb{R}^n$, and $\theta \in \mathbb{R}^{n^2+n}$ is constructed by concatenating columns of θ_1 and θ_2 . The SCP associated with this growth bound is constructed by considering a uniform distribution over $\Omega_{\eta_x}(\hat{\mathbf{x}}) \times V$ and obtain $2N$ i.i.d. sample trajectories $\mathbb{S}_{2N} = \{(\mathbf{x}_i, \hat{\mathbf{u}}, \mathbf{x}'_i) \mid \mathbf{x}'_i \in \Phi(\mathbf{x}_i, \hat{\mathbf{u}}), i = 1, 2, \dots, 2N\}$ so that every \mathbf{x}'_i corresponds to a random disturbance $\mathbf{v}_i \in V$. The modified SCP $_{\gamma}$ is defined as

$$\begin{cases} \min c^T \theta \\ \text{s.t. } 0 \leq \theta \leq \bar{\theta} \text{ and } \forall i \in \{1, \dots, N\} \\ |\mathbf{x}'_{2i-1} - \mathbf{x}'_{2i}| - \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})|\mathbf{x}_{2i-1} - \mathbf{x}_{2i}| - \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}}) + \gamma \leq 0 \end{cases}$$

where $c = [1, 1, \dots, 1] \in \mathbb{R}^{n^2+n}$ is a constant vector, $\bar{\theta} \in \mathbb{R}_{>0}^{n^2+n}$ is sufficiently large, and $\gamma \geq 0$.

Theorem 6.4. *For any $\hat{\mathbf{x}} \in \hat{X}$ constructed with the discretisation size η_x , any $\hat{\mathbf{u}} \in \hat{U}$, and the disturbance set $V = [-\bar{\mathbf{v}}, \bar{\mathbf{v}}]$, the optimal solution of (6.19) gives a growth bound for the system Σ corresponding to $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ that satisfies (6.17) with confidence $(1 - \beta)$, when the number of samples $2N \geq N(\varepsilon, \beta)$ and*

$$\gamma = 8L_{\varphi} \sqrt[4n]{\varepsilon \left[\prod_{i=1}^n \eta_x(i) \prod_{i=1}^n \bar{\mathbf{v}}(i) \right]^2}, \quad (6.19)$$

where $\varepsilon \in [0, 1]$, n is the dimension of the state space, and $L_{\varphi}(\hat{\mathbf{u}})$ is an upper bound on the Lipschitz constant of the system trajectories $\varphi(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{v})$ with respect to (\mathbf{x}, \mathbf{v}) .

Proof. The proof of this theorem is similar to that of Theorem 6.3. Define

$$\Lambda(\theta, \mathbf{x}_1, \mathbf{v}_1, \mathbf{x}_2, \mathbf{v}_2) := \max\{|\varphi(\mathbf{x}_1, \hat{\mathbf{u}}, \mathbf{v}_1) - \varphi(\mathbf{x}_2, \hat{\mathbf{u}}, \mathbf{v}_2)| - \theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})|\mathbf{x}_1 - \mathbf{x}_2| - \theta_2(\hat{\mathbf{x}}, \hat{\mathbf{u}})\}.$$

To satisfy the inequality (6.4), one may choose

$$\Upsilon(\varepsilon) = \mathbb{P}_{\mathcal{D}}(\Omega_{\varepsilon}(d)) = \frac{(\varepsilon/2)^{4n}}{\left[\prod_{i=1}^n \eta_x(i) \prod_{i=1}^n \bar{\mathbf{v}}(i) \right]^2},$$

since the distribution on $(\Omega_{\eta_x}(\hat{\mathbf{x}}) \times V)^2$ is uniform. Using Equation (6.5), one has $\gamma = L_d \Upsilon^{-1}(\varepsilon)$. In order to prove that γ takes the value in (6.19), it must be shown that Λ is Lipschitz continuous with constant $L_d = 4L_{\varphi}(\hat{\mathbf{u}})$. Bounding $\|\theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\|$

by L_φ , for all $(\mathbf{x}_1, \mathbf{v}_1, \mathbf{x}_2, \mathbf{v}_2)$ and $(\mathbf{x}'_1, \mathbf{v}'_1, \mathbf{x}'_2, \mathbf{v}'_2)$ then

$$\begin{aligned} & \|\Lambda(\theta, \mathbf{x}_1, \mathbf{v}_1, \mathbf{x}_2, \mathbf{v}_2) - \Lambda(\theta, \mathbf{x}'_1, \mathbf{v}'_1, \mathbf{x}'_2, \mathbf{v}'_2)\| \\ & \leq \|\varphi(\mathbf{x}_1, \hat{\mathbf{u}}, \mathbf{v}_1) - \varphi(\mathbf{x}'_1, \hat{\mathbf{u}}, \mathbf{v}'_1)\| \\ & \quad + \|\varphi(\mathbf{x}_2, \hat{\mathbf{u}}, \mathbf{v}_2) - \varphi(\mathbf{x}'_2, \hat{\mathbf{u}}, \mathbf{v}'_2)\| \\ & \quad + \|\theta_1(\hat{\mathbf{x}}, \hat{\mathbf{u}})\|(\|\mathbf{x}_1 - \mathbf{x}'_1\| + \|\mathbf{x}_2 - \mathbf{x}'_2\|) \\ & \leq 4L_\varphi(\hat{\mathbf{u}})\|(\mathbf{x}_1, \mathbf{v}_1, \mathbf{x}_2, \mathbf{v}_2) - (\mathbf{x}'_1, \mathbf{v}'_1, \mathbf{x}'_2, \mathbf{v}'_2)\|. \end{aligned}$$

Therefore, Λ is Lipschitz continuous with constant $4L_\varphi(\hat{\mathbf{u}})$. This completes the proof. \square

A statement similar to Corollary 6.1 holds for the growth bound computed using (6.19).

6.6 Experimental Evaluation

To demonstrate the approach, it is applied to a DC-DC boost converter. This case study is taken from Girard *et al.* [66] and will be used as a black-box model to generate sample trajectories. Another case study is provided from power systems based on the work of Ma and Fan [117], that is implemented in the Power System Toolbox (PST) [36]. Trajectories from the black-box reduced model of the 30 state power system model will be used. The approach is applied to construct finite abstractions of these systems and employ SCOTS [163] to design controllers. The algorithms are implemented in C++ on a 64-bit Linux cluster machine with two Intel Xeon E5 v2 CPUs, 1866 MHz, and 50GB RAM. Additional case studies can be found in the original paper [86].

6.6.1 DC-DC Boost Converter

The objective in the DC-DC boost converter problem is to design a controller to enforce a reach and stay specification. The DC-DC boost converter can be modelled as a two dimensional linear switching system with two functional modes. The state vector of the system at time $t \in \mathbb{R}_{\geq 0}$ is $\mathbf{x}(t) = (i_l(t), v_c(t))$, where i_l is the inductor current and v_c is the capacitor voltage. The system's evolution can be controlled by selecting the appropriate mode $\mathbf{u}(t) \in \{1, 2\}$ at every time $t \in \mathbb{R}_{\geq 0}$. The system's dynamics under the two modes can be represented as $\dot{\mathbf{x}} = A_{\mathbf{u}(t)}\mathbf{x}(t) + b + c\mathbf{v}(t)$, $\mathbf{u} \in \{1, 2\}$, with matrices A_1, A_2, b, c as reported in [66]. The state and input spaces are $X = [0.65, 1.65] \times [4.95, 5.95]$ and $U = [1, 2]$. The initial state is $(i_{l_0}(t), v_{c_0}(t)) = (0.7, 5.4)$ and the target set is $[1.1, 1.6] \times [5.4, 5.9]$. The target set is shown in red colour in Figure 28.

The implementation results are reported in Table 3 for the system without disturbance ($\bar{\mathbf{v}} = (0, 0)$) and with disturbance bound $\bar{\mathbf{v}} = (0.01, 0)$. These results are obtained with discretisation parameters $\eta_x = (0.005, 0.005)$ and $\eta_u = 1$, confidence parameter $\beta = 0.01$, $\varepsilon = 0.01$ and estimation for $L_\varphi = 0.9935$. The resulted

Table 3: Results for the DC-DC boost converter.

Case-study	Dimension		Disturbance V	Fixed Discretisation		
	X	U		N	time (min)	$ \mathcal{V} $
DC-DC boost converter	2	1	$\{0\}$	1,807	22.2	37,783
			$[-0.01, 0.01]$	2,285	30.6	37,414

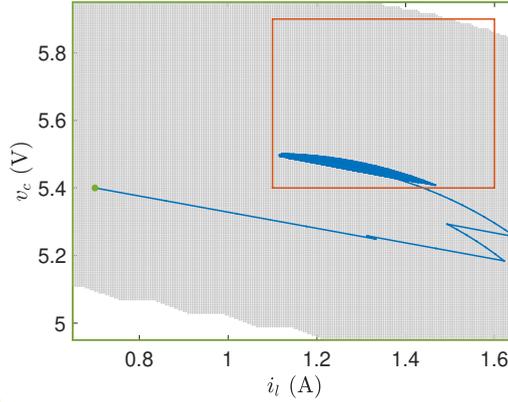


Figure 28: The closed-loop trajectory of the DC-DC boost converter with $\bar{\mathbf{v}} = (0, 0)$ under the controller designed by the data-driven abstraction approach. The rectangle in red colour represents the target region and the area in grey shows the winning region of the controller.

finite abstraction has cardinalities $n_x = 40,000$ and $n_u = 2$. The required number of sample trajectories, N , for each $(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \hat{X} \times \hat{U}$ is computed using equation (6.6). Runtimes and the resulting winning region sizes, $|\mathcal{V}|$, for the DC-DC boost converter are given in Table 3.

Algorithm 1 is used to compute the finite-state abstraction by collecting sample trajectories of the system. Subsequently, SCOTS is used for designing the controller. The performance of the controller is shown in Figures 28 and 29 for the system without and with the disturbance. These figures show one sample closed-loop trajectory of the system under the controllers designed by the data-driven ABCD approach. In both cases, without and with disturbance, it can be noticed from Figures 28 and 29 that the approach has been successful in finding controllers satisfying the given reach and stay specification, despite the dynamics being unknown.

6.6.2 Three Area Three Machine Power System

Consider a three area three machine (3A3M) power system adapted from Ma and Fan [117], and shown in Figure 30. The system consists of three buses, which

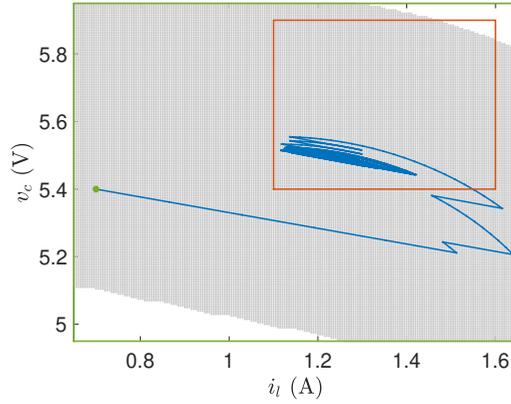


Figure 29: The closed-loop trajectory of the DC-DC boost converter with $\bar{v} = (0.01, 0)$ under the controller designed by the data-driven abstraction approach. The rectangle in red colour represents the target region and the area in grey shows the winning region of the controller.

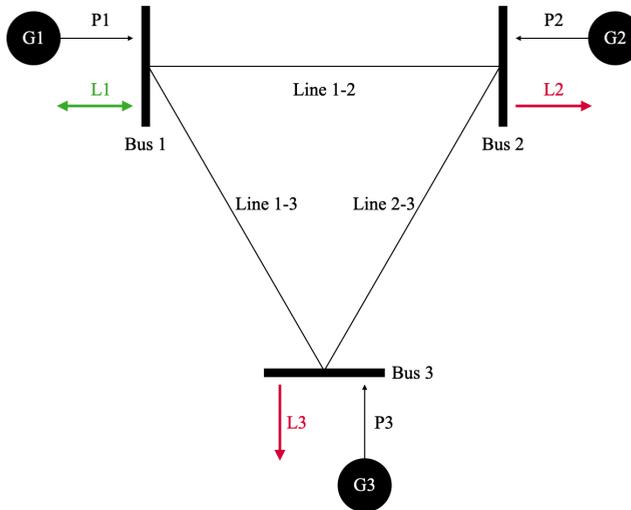


Figure 30: 3A3M power system with generators (G) and loads (L). L1 represents a bidirectional load such as Electric Vehicles or Energy Storage Systems.

are each connected to a power source (generator) and a load. At bus 1 consider a load which is bidirectional, meaning it can both draw power and inject power into the system. The loads at buses 2 and 3 can only draw power from the system; when these loads increase, more power will be drawn from the system, causing an imbalance between generation and consumption which may result in a reduction of the network frequency. The nominal frequency of the network is set to 60 Hz.

A worst-case scenario is considered where a sudden increase occurs in the loads at buses 2 and 3 by 0.2 and 0.3 pu, respectively. The control task is for the load at bus 1 to balance the load increase at buses 2 and 3 by either reducing its load or injecting power into the network. The simulation is run using PST on a 30 state model of this power system. Balanced realisation, the model-order reduction technique from Sec. 4.8, of the system reduces its dynamics to three states. To compute the data-driven finite abstraction, sample trajectories are gathered using a black-box approach of the reduced system representation for the original model. Note that this model-order reduction involves some information loss, this will be dealt with in Chapter 8 and Chapter 9. The dynamics of the reduced system are given by

$$\begin{aligned}\dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{D}\mathbf{v} \\ \mathbf{y} &= \mathbf{C}\mathbf{x},\end{aligned}\tag{6.20}$$

where

$$\mathbf{A} = \begin{bmatrix} 0.00027563 & 0 & 0 \\ 0 & -0.3951 & 0.687 \\ 0 & -0.6869 & -0.016 \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} 0.00031166 \\ 0.1359 \\ 0.0230 \end{bmatrix}$$

$$\mathbf{D} = \begin{bmatrix} 0.00033103 & 0.00031244 \\ 0.1309 & 0.1308 \\ 0.0250 & 0.0233 \end{bmatrix}$$

$$\mathbf{C} = [-0.0115 \quad -0.2296 \quad 0.0412].\tag{6.21}$$

The state and input spaces are $X = [-0.02, 0.02] \times [-0.05, 0.05] \times [-0.12, 0.12]$ and $U = [0, 0.5]$. Further, $V = [-0.2, 0.2] \times [-0.3, 0.3]$, $\eta_u = 0.025$, $\tau = 0.4$, $\eta_x = (0.0015, 0.0015, 0.0015)$, $\beta = 0.01$ and $\varepsilon = 0.01$ are set. The resulted abstraction has $n_x = 228, 480$ and $n_u = 20$. The estimated Lipschitz constant is $L_\varphi = 1.5715$. The target set is given by $-0.008 < \mathbf{y} < 0.008$ and the avoid set is given by $\mathbf{y} < -0.015$. Multiplying by the nominal frequency to get the specification in Hertz, the target region is $[59.52, 60.48]$ and the avoid region is $(-\infty, 59.1)$. Figure 31 shows that the specification is violated when no control is applied.

The data-driven approaches of Section 6.4 (fixed discretisation) and Section 6.5 (abstraction refinement) are now applied. Both controllers are synthesised with disturbance $V = [-0.2, 0.2] \times [-0.3, 0.3]$. A comparison of the two control approaches is shown in Table 4. The required number of sample trajectories for each

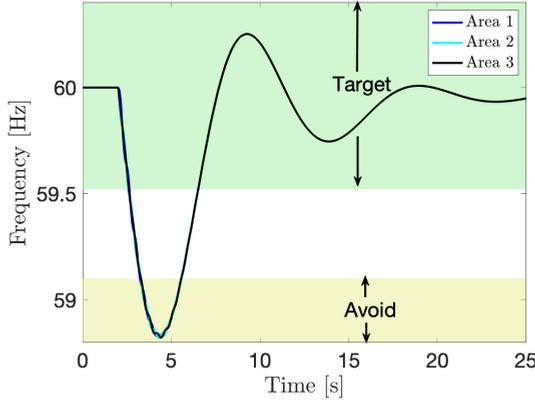


Figure 31: 3A3M power system frequency without applying any control input. The frequency falls below 59.1 Hz thus violates the specification.

(\hat{x}, \hat{u}) is computed using equation (6.6) and marked with N in the table. The abstraction refinement starts with $\eta_x = 0.012$ and refines the discretisation iteratively with a factor of two. The algorithm successfully finds a controller after five iterations. The runtimes and the resulting winning region sizes $|\mathcal{V}|$ are also given in Table 4. The abstraction refinement synthesises the controller a factor of 100 times faster than the fixed discretisation by iteratively decreasing the value of η_x .

Table 4: Results for the 3A3M power system.

Control Approach	Dimension		Disturbance \bar{v}	N	time (min)	$ \mathcal{V} $
	X	U				
Fixed Discretisation	3	1	(0.2, 0.3)	3, 290	5, 253	230, 760
Adaptive Refinement	3	1	(0.2, 0.3)	4, 460	50.25	314, 802

The data-driven control approach with fixed discretisation is simulated in PST and is reported in Figures 32 and 33. The controlled system successfully keeps the frequencies of the three areas outside of the avoid set (i.e., always above 59.1 Hz) and bring them back to the target set (i.e., above 59.52 Hz). Figure 33 shows the load changes in the system. Load at bus 1 is able to maintain the frequencies of the three areas above the avoid region and facilitate the system returning to the target set for the maximum disturbances applied at buses 2, 3. Figures 34 and 35 show the results of simulating the system in PST with the control obtained from the abstraction refinement approach. The controlled system has the same performance in satisfying the specification.

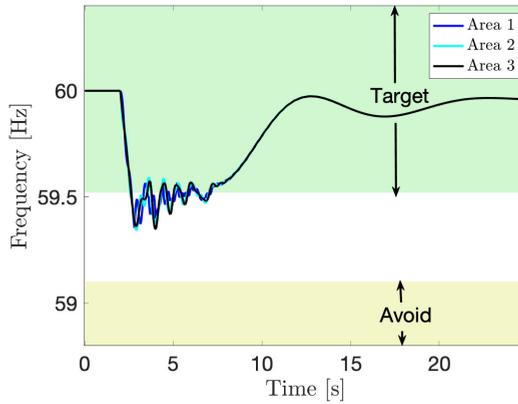


Figure 32: 3A3M power system frequencies for the three areas, with the frequency of an area is measured at the corresponding bus in that area. The control synthesised by the fixed discretisation approach successfully keeps the frequencies of the three areas outside of the avoid set. The frequencies leave the target set for around 4.4 seconds before staying in the target set.

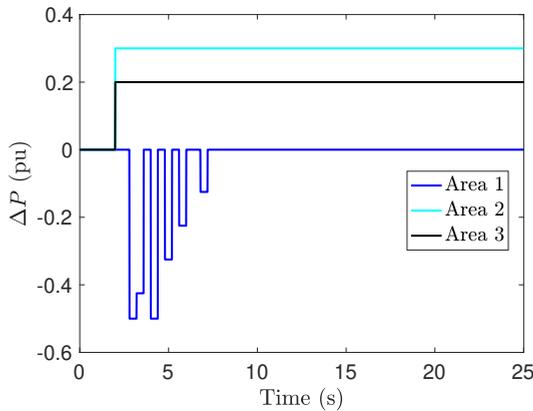


Figure 33: 3A3M power system load changes for the three areas. Loads at buses 2 and 3 increase by 0.3 and 0.2 pu, respectively. Load at bus 1 is used to control the frequency using the data-driven approach with fixed discretisation.

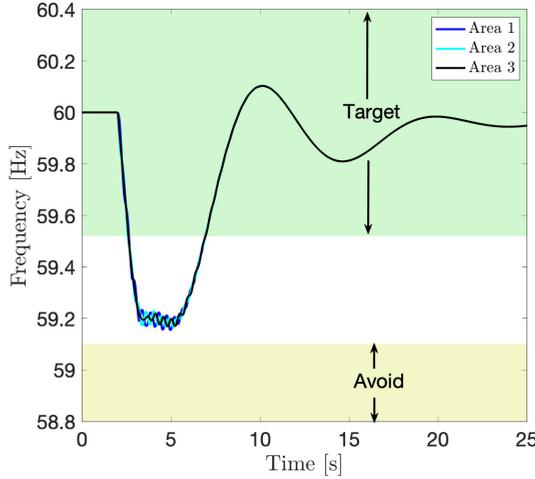


Figure 34: 3A3M power system frequencies for the three areas, with the frequency of an area is measured at the corresponding bus in that area. The control synthesised by the abstraction refinement approach successfully satisfies the specification. The frequencies leave the target set for around 4.2 seconds before staying in the target set.

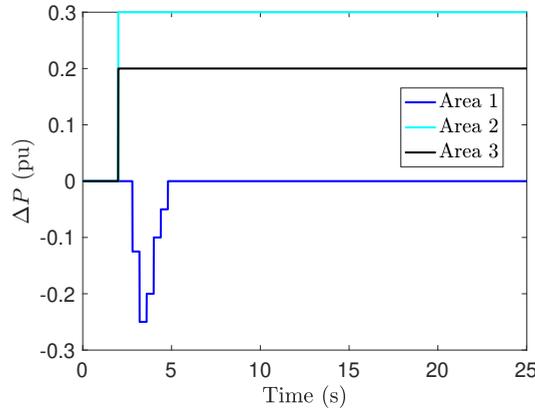


Figure 35: 3A3M power system load changes for the three areas. Loads at buses 2 and 3 increase by 0.3 and 0.2 pu, respectively. Load at bus 1 is used to control the frequency using the data-driven approach with abstraction refinement.

Table 5: Comparing the winning domain of controllers obtained from the RSA method, PAC method of [212], and the model-based approach of [153]. The pairwise comparison is made by computing the intersections (\cap) and set differences (row \setminus column). The results are reported both in cardinalities and percentages.

Winning Domain	RSA		PAC		Model-based	
	\cap	\setminus	\cap	\setminus	\cap	\setminus
RSA	230,760	0	230,760	0	230,760	0
%	100.00%	0.00%	100.00%	0.00%	100.00%	0.00%
PAC	230,760	15,664	246,424	0	245,345	1,079
%	93.64%	6.36%	100.00%	0.00%	99.56%	0.44%
Model-based	230,760	22,216	245,345	7,631	252,976	0
%	91.22%	8.78%	96.98%	3.02%	100.00%	0.00%

6.6.3 Comparison with PAC Learning

In this subsection, the approach is compared with the results provided by Xue *et al.* [212] that is based on probably approximately correct (PAC) bounds on the 3A3M power system case study. The abstraction approach of Xue *et al.* has no bias term γ and requires the number of samples

$$N \geq \frac{2}{\nu} (\ln \frac{1}{\beta} + q), \quad (6.22)$$

where $\beta \in (0, 1)$ is the confidence parameter, $\nu \in (0, 1)$ is the error threshold, and $q \in \mathbb{N}$ is the cardinality of the parameter vector θ . The error threshold allows the constructed abstraction to hold for the entire state space except a subset measured by parameter $\nu > 0$. Note that setting $\nu = 0$ in (6.22) results in requiring an infinite dataset, which is not practical. The approach provides an abstraction that is correct on the entire state space, i.e., $\nu = 0$, with a finite sample size.

The data-driven robust scenario approach (RSA), the PAC approach [212] with parameters $\beta = 0.01$ and $\nu = 0.01$, and the model-based approach [153] are now implemented. Table 5 compares the winning domain of the controllers by reporting the intersections (\cap) and set differences (row \setminus column). It can be seen that the winning domain obtained by the RSA method is a subset of the ones computed by PAC and the model-based approaches. This shows that the approach is more conservative than the model-based approach but correctly finds a subset of the winning domain. In contrast, the PAC approach gives a winning domain that includes states not identified as winning by the model-based approach. It includes 1079 states outside of the winning domain obtained by the model-based approach. Due to the nature of the PAC learning, some of these states are incorrectly identified as winning. The main reason is that the PAC method may miss capturing some of the transitions and does not always generate an overapproximation of the system behaviours. Among these 1079 states, a counter-example

can be found, demonstrating a lack of guarantee provided by the PAC method. At state $(0.0187, 0.0262, -0.1163)$ the PAC controller calculates $\mathbf{u} = -0.075$ to be an input which will transition to a safe state under any disturbances. However, the system under disturbances $V_1 = 0.2$ and $V_2 = 0.3$ will lead to the state $(0.0188, 0.0131, -0.1167)$ that is outside of the winning domain of the controller. In comparison, the winning domain provided by the RSA method is a subset of the one from the model-based method and provides full guarantees of the satisfaction of the specification and correctness of the controller. This guarantee is obtained at the cost of an increased number of samples and a bias term included in the growth bound calculations, which makes the controller more conservative.

As a final point on this case study, note that the sampling approach uses the Lipschitz constant estimated using sample trajectories. This Lipschitz constant can in turn be used to construct the abstraction. The direct use of the estimated Lipschitz constant does not provide a formal guarantee as it is an estimated value that converges to the true value only in the limit (i.e., the number of samples goes to infinity), and is likely to provide an overly conservative controller. To account for a finite sample size, the upper bound on the Lipschitz constant needs to be corrected by multiplying it with a factor greater than one after observing the convergence behaviour of the distribution fitting for different sizes of the dataset. In this particular case study, the direct use of the upper bound on the Lipschitz constant (without correction) gives a controller that covers only 78.8% of the winning domain of the model-based approach.

6.6.4 Parameter Optimisation

In this subsection, it is discussed how a selection of different parameters can affect the sample complexity and conservativeness of the method. This is based on a path planning case study from the paper [86], with the estimated Lipschitz constant of 1.46. Figures 36 and 37 illustrate the effect of changing parameters ε, β on the number of samples N required for each pair $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$ in order to compute the growth bound with confidence $(1 - \beta)$. Figure 36 illustrates the effect of increasing the confidence parameter β on reducing the sample complexity, for a fixed $\varepsilon = 0.01$. Figure 37 shows that for a fixed $\beta = 0.01$, increasing ε leads to a rapid drop in N . In both Figures 36 and 37, the sample complexity increases in the presence of disturbance as the dimension of the sample space becomes larger.

Figure 38 demonstrates the effect of changing ε on the value of the bias term γ that makes the inequalities of the SCP more conservative. The bias term γ increases for larger values of ε . Therefore, increasing ε can decrease the sample complexity while increasing γ . Finally, it can be observed that the value of γ is larger in the presence of disturbance.

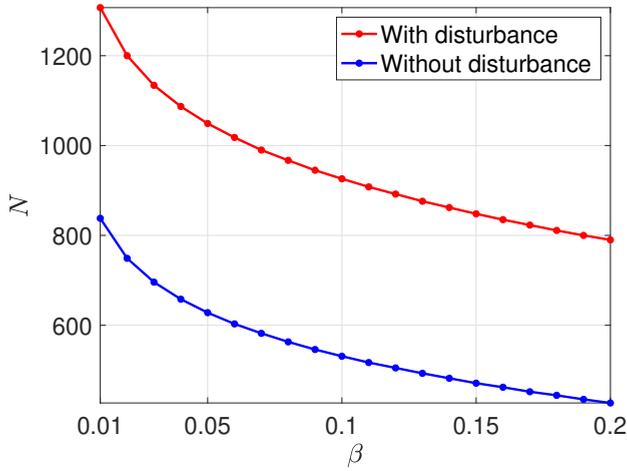


Figure 36: Required number of samples for the approach as a function of β for a fixed $\varepsilon = 0.01$.

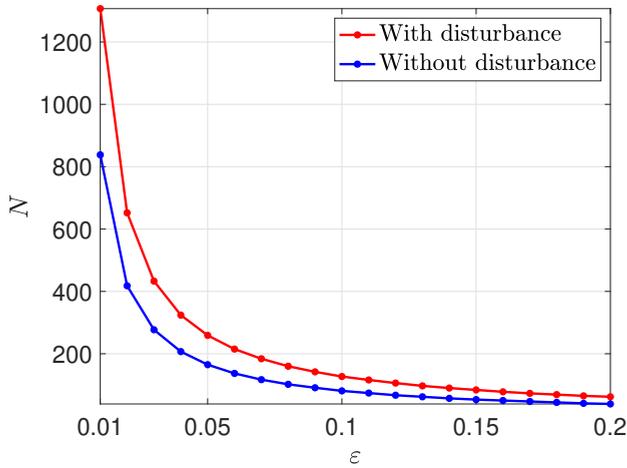


Figure 37: Required number of samples for the approach as a function of ε for a fixed $\beta = 0.01$.

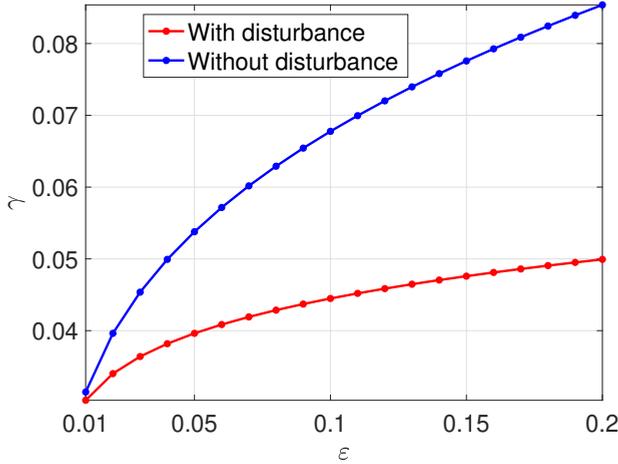


Figure 38: The bias term γ as a function of ϵ .

6.7 Conclusion

In conclusion, this chapter proposed a new data-driven method for computing finite abstractions of continuous systems with unknown dynamics. In particular:

- The approach casts the computation of an overapproximation of reachable sets as a robust convex program (RCP). A feasible solution for the RCP is then obtained with a given confidence by solving a corresponding scenario convex program (SCP). The SCP does not need the dynamics of the system and requires only a finite set of sample trajectories.
- A sample complexity result was provided that gives a lower bound on the number of trajectories to achieve a certain confidence. The sample complexity results require knowing a bound on the Lipschitz constant of the system, which was estimated using extreme value theory.
- Guarantees were given that with high confidence the computed abstraction is a valid abstraction of the system that overapproximates its behaviours on its entire state space. It was shown that the data-driven approach can be embedded into abstraction refinement schemes for designing a controller and enlarging the winning region of the controller with respect to satisfaction of temporal properties.
- The approach has an exponential complexity with respect to the dimension in computing the growth bound. This exponential sample complexity is due to the lack of knowledge of the dynamics and requiring guarantees on satisfaction of robust convex programs. I am not aware of any result in the literature that does not have this exponential complexity while providing the type of guarantees in this chapter.

- The approach is highly parallelisable, and computations can be done individually for each state-action pair in parallel. The abstraction can also be refined depending on the local computation of the growth bound, thus improving the performance of the model-based approaches depending on the required confidence.
- Finally, the approach was evaluated on two case studies; a DC-DC boost converter and the 3A3M power system.

The next chapter, will improve the scalability of the approaches being used, in particular to consider a 9-state area of the well-known benchmark New England 39-bus Test System (NETS) and use approximate simulation functions to complete model-reduction with robustness to the information lost during the reduction.

Robust Simulation Functions with Disturbance Refinement

The results of this chapter is motivated by the lack of formal guarantees for model-order reduction, particularly guarantees of the closeness between the trajectories of the original system and its reduced order model over time, e.g. the reduction step taken for the 3A3M model used in the previous chapter. This chapter, based on the work [207], approximates concrete systems with abstractions of lower dimension (reduced-order models) and develops *robust* simulation functions further to consider the perturbation in the abstract system by designing an interface function for the disturbance. Simulation functions are Lyapunov-like functions defined over the Cartesian product of state spaces of two (un)perturbed systems, *a.k.a., concrete and abstract systems*, to relate output trajectories of abstract systems to those of concrete ones while the mismatch between two systems remains within some guaranteed error bounds. The proposed approach allows concrete systems to have large disturbances, which is the case in many real-life applications, while noticeably reducing the closeness error between the two systems. Accordingly, this enables controller design using a reduced-order form of the concrete system and reducing the computational efforts required for formal synthesis. The efficacy of the approach is demonstrated on linear control systems by synthesising a formal controller for a 9-state area of the known New England 39-Bus Test System, using only a 3-state abstract system.

Notation. The following notation overrides any previous definitions and uses in prior chapters. The notation $\|a\|$ is used for the Euclidean norm of a vector a , and $\|a\|_\infty$ for taking the Euclidean norm followed by a maximisation over the bounded domain of a . Primarily in this chapter, original system and concrete system are used interchangeably to describe Σ_1 , and abstract system or reduced-order model are used to describe Σ_2 , although Σ_2 does not have to necessarily be of reduced-order, the examples considered treat Σ_2 as lower dimension than Σ_1 .

7.1 Introduction

7.1.1 Motivations and State of the Art.

Cyber-physical systems (CPS) are complex networked models combining both cyber (computation and communication) and physical components, which tightly interact with each other in a feedback loop [106]. In the past few years, CPS have gained remarkable attentions as an important modelling tool for engineering systems spanning a wide range of real-life applications of which power systems are the primary focus of this thesis. The interconnection of CPS components in the models often results in high-dimensional systems with complex behaviour specifications that are safety critical in nature.

Providing safety and reliability guarantees on the behaviour of these complex systems is therefore essential but also incredibly challenging as formal methods, which can achieve such guarantees, often suffer from the *curse of dimensionality* and cannot handle high-dimensional models [77]. In particular, formal methods give a strong mathematical framework to provide guarantees over CPS, whether that is verifying the behaviour of a system or synthesising a controller to create (or enforce) system behaviour [99, 146].

To alleviate the encountered computational complexity, symbolic control is one of the promising techniques, proposed in the relevant literature, for formal analysis of CPS [188]. In this regard, symbolic abstractions replace concrete systems to provide a more appropriate medium for formal verification or controller synthesis of CPS. Since the mismatch between outputs of concrete systems and those of their symbolic abstractions are well-quantified, one can guarantee that concrete systems also satisfy the same property of interest as abstract ones with some quantified error bounds.

In order to relate output trajectories of abstract systems to those of concrete ones, *simulation* and *bisimulation functions* (where both systems can simulate each other) are powerful techniques, proposed in the related literature [188]. If concrete and abstract systems are (bi)similar, one can consider the abstract system as an appropriate substitute in the controller design process with reduced computational loads while still preserving closeness guarantees between the two systems. For underlying systems where expecting the same output may be too strict, *approximate* (bi)simulation functions have been developed in the literature [63]. The reader may refer to Chapter 4 for more details.

Approximate (bi)simulation functions aim at establishing a formal relation between the abstract system which is similar to the concrete one, while bounding the closeness between the outputs of two systems over time by some maximal threshold ϵ , known as the simulation relation error. An interface function is then designed to map the control inputs from the abstract system to the concrete domain enforcing the ϵ -closeness. This notion is extended in [96] to *robust simulation functions* (RSFs), which considers small disturbances inside the concrete system, while the abstract system remains unperturbed, to establish an approximate simulation relation between the two systems.

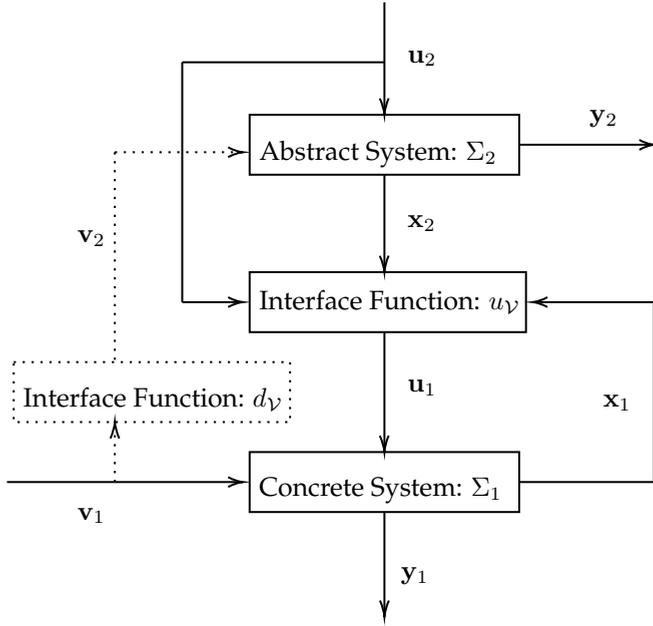


Figure 39: Hierarchical control system architecture employed in this chapter. The dashed part considers the need for the disturbance in the low-dimensional abstract system Σ_2 for the sake of control over large measurable disturbances.

7.1.2 Original Contributions.

The main contribution in this chapter is to extend the notion of simulation functions to its robust versions by incorporating the disturbance in the abstract system via designing an interface function for the disturbance, see Fig. 39. This reduces the simulation relation error ϵ , particularly when one is dealing with concrete systems with large disturbances. Incorporating the disturbance in the abstract system enables formal controller synthesis design for the concrete system using the abstract system where ϵ is included in the controller process. Consequently, formal controllers designed on a low-dimensional abstract system can be refined back to control any high-dimensional concrete systems models. The efficacy of the approach is demonstrated on a case study of the *New England 39-Bus Test System* (NETS).

7.1.3 Related Work.

There have been some results, proposed in the past two decades, on establishing (bi)simulation functions for dynamical systems. In this respect, the work [96] extends the approaches of simulation functions to consider small disturbances in the concrete domain providing robustness in the simulation relation. However, in most real-world scenarios, the proposed approach may not be practical given that the simulation relation error increases for larger disturbances. The work [147]

demonstrates systems that are approximately equivalent (bisimilar) to their symbolic models. The results in [65] provide an approximation framework that applies to both discrete and continuous systems.

The approach in [181] employs approximate bisimulation in transient power systems, which is mainly used for model order reductions: they consider differential-algebraic equations as their model of NETS with bounded disturbances. Reachability and formal analysis of power systems have been studied in [10, 110]. A controller designed based on abstract models for frequency regulation of smart grids is studied in [209]. Compositional abstraction-based techniques to construct symbolic models for an interconnected system based on symbolic models of individual smaller subsystems are studied in [100, 101, 148, 190]. Data-driven construction of finite abstractions has been studied in Chapter 6 for continuous-time systems [86], discrete-time incrementally input-to-state stable systems [98], and monotone systems [124].

The rest of this chapter structured as follows. Preliminaries and the formal definition of underlying systems are presented in Section 7.2. Section 7.3 contains the solution methodologies while considering the disturbance refinement. The approach is demonstrated over NETS in Section 7.4 and conclusions are given in Section 7.5.

7.2 Preliminaries

Class of Systems. Consider two general dynamical systems, as in Def. 4.1, Σ_1 and Σ_2 , modelled as:

$$\Sigma_i : \begin{cases} \dot{\mathbf{x}}_i = g_i(\mathbf{x}_i, \mathbf{u}_i, \mathbf{v}_i), \\ \mathbf{y}_i = h_i(\mathbf{x}_i), \end{cases} \quad i \in \{1, 2\}, \quad (7.1)$$

where $\mathbf{x}_i \in \mathbb{R}^{n_i}$ are system states, $\mathbf{u}_i \in \mathbb{R}^{p_i}$ are control inputs, $\mathbf{y}_i \in \mathbb{R}^{m_i}$ are system outputs, $\mathbf{v}_1 \in \mathbb{R}^{q_1}$ is a measurable large disturbance in Σ_1 and \mathbf{v}_2 is derived from \mathbf{v}_1 with an interface function d_γ . Similarly, \mathbf{u}_1 can be derived from \mathbf{u}_2 using an interface function u_γ . Without loss of generality, consider Σ_1 as the original system and Σ_2 as the (possibly) lower-dimensional abstraction. It can then be taken that $n_2 \leq n_1$.

Linear Temporal Logic Specifications. For the dynamical systems in (7.1), linear temporal logic (LTL) specifications are considered with syntax

$$\psi := \text{true} \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \bigcirc\psi \mid \psi_1 \cup \psi_2,$$

where p is the element of an atomic proposition, see Def. 4.10. We refer the reader back to Chapter 4 for further details.

7.3 Solution Methodologies

The main contribution of this chapter is to extend the notion of simulation functions to its robust versions by considering disturbance refinement using an interface function for the disturbance in the concrete system to be visible in the abstract domain. The proposed approach enables the controller synthesis for systems with large disturbances.

The following section shows how incorporating the disturbance of the concrete system into the abstract one through the interface function $d_{\mathcal{V}}$ can further reduce the simulation relation error ϵ between Σ_1 and Σ_2 . This enables one to perform controller synthesis on the abstract domain and refine it back over the high-dimensional original system while improving the scalability of the control scheme.

7.3.1 Robust Approximate Simulation with Disturbance Refinement

Given the systems in (7.1), a robust approximate simulation with disturbance refinement is defined with a robust simulation function \mathcal{V} and two interface functions $u_{\mathcal{V}}$ and $d_{\mathcal{V}}$. The function \mathcal{V} has the following Lyapunov-like properties:

Definition 7.1. Consider the two systems in (7.1). Let $\mathcal{V} : \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightarrow \mathbb{R}^+$ be a differentiable function, $u_{\mathcal{V}} : \mathbb{R}^{p_2} \times \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightarrow \mathbb{R}^{p_1}$ and $d_{\mathcal{V}} : \mathbb{R}^{q_1} \times \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightarrow \mathbb{R}^{q_2}$ be continuous functions. Then the function \mathcal{V} is called a robust simulation function from Σ_2 to Σ_1 with the associated interface functions $u_{\mathcal{V}}$ and $d_{\mathcal{V}}$ if there exists class- κ functions ϱ_1 and ϱ_2 such that for all $\mathbf{x}_1 \in \mathbb{R}^{n_1}$ and $\mathbf{x}_2 \in \mathbb{R}^{n_2}$,

$$\|h_1(\mathbf{x}_1) - h_2(\mathbf{x}_2)\| \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2), \quad (7.2)$$

and for any $\mathbf{u}_2 \in \mathbb{R}^{p_2}$ and $\mathbf{v}_1 \in \mathbb{R}^{q_1}$ satisfying $\varrho_1(\|\mathbf{v}_1\|) + \varrho_2(\|\mathbf{u}_2\|) \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2)$, then

$$\frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} g_2(\mathbf{x}_2, \mathbf{u}_2, d_{\mathcal{V}}(\mathbf{v}_1, \mathbf{x}_1, \mathbf{x}_2)) + \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} g_1(\mathbf{x}_1, u_{\mathcal{V}}(\mathbf{u}_2, \mathbf{x}_1, \mathbf{x}_2), \mathbf{v}_1) \leq 0. \quad (7.3)$$

Σ_1 robustly approximately simulates Σ_2 if there exists a robust simulation function \mathcal{V} of Σ_2 by Σ_1 .

Remark 7.1. Definition 7.1 is a generalisation of the robust approximate simulation notation proposed in the literature [96]. In particular, when $d_{\mathcal{V}} = 0$, then the existing robust approximate simulation is recovered.

The next subsection focuses on the class of linear control systems with potentially large measurable disturbances and proposes an approach to construct its reduced-dimensional abstractions together with a robust simulation function as presented in Definition 7.1.

7.3.2 Linear Systems under Large Measurable Disturbance

Here, the focus is on the class of linear control systems with (potentially large) measurable disturbances, defined as follows:

$$\Sigma_i : \begin{cases} \dot{\mathbf{x}}_i = A_i \mathbf{x}_i + B_i \mathbf{u}_i + D_i \mathbf{v}_i, \\ \mathbf{y}_i = C_i \mathbf{x}_i, \end{cases} \quad i \in \{1, 2\}, \quad (7.4)$$

where A_i, B_i, C_i , and D_i are matrices of appropriate dimensions, and \mathbf{v}_1 is the measured disturbance having some known bound $\|\mathbf{v}_1\|_\infty \leq v_{\max}$. The main problem to solve in this chapter is now stated.

Problem Description 7.1. *Given a linear system Σ_1 as in (7.4) under (potentially large) measurable disturbances and an LTL specification ψ , construct its reduced-dimensional abstraction Σ_2 together with robust simulation functions according to Definition 7.1. Employ the constructed abstraction Σ_2 and design a formal controller through robust simulation relations with disturbance refinement such that the specification is satisfied over the original system.*

In order to address Problem 7.1, the following lemma and theorems need to be raised. Note that the next lemma is similar to the one presented in [63] but it is adapted here to this setting by incorporating the measurable disturbance inside the dynamics.

Lemma 7.1. *If Σ_1 is stabilisable, there are matrices K_1, K_2, P, D_2, Q_1 such that $(A_1 - PD_2(K_1 + Q_1) + B_1K_2)$ is Hurwitz, and there exist a positive definite matrix M and positive scalar constant λ such that the following matrix inequalities hold:*

$$C_1^T C_1 \leq M, \quad (7.5a)$$

$$\begin{aligned} & (A_1 - PD_2(K_1 + Q_1) + B_1K_2)^T M + \\ & M(A_1 - PD_2(K_1 + Q_1) + B_1K_2) \leq -2\lambda M. \end{aligned} \quad (7.5b)$$

Remark 7.2. *The matrices M and K_2 in Lemma 7.1 can be computed using semi-definite programming by letting $\bar{K} = K_2 M^{-1}$ and $\bar{M} = M^{-1}$. This gives the equivalent matrix inequality conditions:*

$$\begin{aligned} & \begin{bmatrix} \bar{M} & \bar{M}C_1^T \\ C_1\bar{M} & \mathbb{I} \end{bmatrix} \geq 0, \text{ and} \\ & \bar{M}A_1^T - \bar{M}(Q_1^T + K_1^T)D_2^T P^T + \bar{K}^T B_1^T \\ & \quad + A_1\bar{M} - PD_2(K_1 + Q_1)\bar{M} + B_1\bar{K} \leq -2\lambda\bar{M}. \end{aligned}$$

Under Lemma 7.1, the next theorem is proposed to construct the robust simulation function \mathcal{V} .

Theorem 7.1. *Consider two systems of the form (7.4). Assume that Σ_1 is stabilisable, a feedback gain K_1 exists for Σ_2 and that there exist matrices P, K_2, Q_1 and Q_2 such that*

$(A_1 + B_1K_2 - PD_2Q_1)$ is Hurwitz, and the following matrix equalities hold:

$$A_1P + B_1Q_2 = PA_2 + PD_2Q_1P, \quad (7.6a)$$

$$C_2 = C_1P. \quad (7.6b)$$

Then \mathcal{V} in the form of

$$\mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{(\mathbf{x}_1 - P\mathbf{x}_2)^T M (\mathbf{x}_1 - P\mathbf{x}_2)}$$

is a robust simulation function from Σ_2 to Σ_1 with its associated interfaces

$$u_{\mathcal{V}} = R_2\mathbf{u}_2 + Q_2\mathbf{x}_2 + K_2(\mathbf{x}_1 - P\mathbf{x}_2), \quad (7.7a)$$

$$d_{\mathcal{V}} = R_1\mathbf{v}_1 + Q_1\mathbf{x}_1 + K_1(\mathbf{x}_1 - P\mathbf{x}_2). \quad (7.7b)$$

In addition, the class- κ functions ϱ_1 and ϱ_2 are designed as

$$\varrho_1(\nu) = \frac{\|\sqrt{M}(D_1 - PD_2R_1)\|}{\lambda} \nu, \quad (7.8)$$

$$\varrho_2(\nu) = \frac{\|\sqrt{M}(B_1R_2 - PB_2)\|}{\lambda} \nu, \quad (7.9)$$

where R_1 and R_2 are some arbitrary matrices of appropriate dimensions and M, λ are such that (7.5) holds.

Proof. From (7.5a) and (7.6b), then

$$\mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) \geq \sqrt{(\mathbf{x}_1 - P\mathbf{x}_2)^T C_1^T C_1 (\mathbf{x}_1 - P\mathbf{x}_2)} = \|C_1\mathbf{x}_1 - C_2\mathbf{x}_2\|,$$

so condition (7.2) holds. To prove condition (7.3) by using conditions (7.5b) and (7.6a), one has

$$\begin{aligned} & \frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} (A_2\mathbf{x}_2 + B_2\mathbf{u}_2 + D_2d_{\mathcal{V}}) + \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} (A_1\mathbf{x}_1 + B_1u_{\mathcal{V}} + D_1\mathbf{d}_1) \\ & \leq -\lambda \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) + \|\sqrt{M}(D_1 - PD_2R_1)\mathbf{d}_1 + \sqrt{M}(B_1R_2 + PB_2)\mathbf{u}_2\| \\ & \leq -\lambda \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) + \|\sqrt{M}(D_1 - PD_2R_1)\| \|\mathbf{d}_1\| + \|\sqrt{M}(B_1R_2 + PB_2)\| \|\mathbf{u}_2\|. \end{aligned}$$

Therefore, for all \mathbf{d}_1 and \mathbf{u}_2 satisfying

$$\frac{\|\sqrt{M}(D_1 - PD_2R_1)\|}{\lambda} \|\mathbf{d}_1\| + \frac{\|\sqrt{M}(B_1R_2 - PB_2)\|}{\lambda} \|\mathbf{u}_2\| \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2),$$

then

$$\frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} (A_2\mathbf{x}_2 + B_2\mathbf{u}_2 + D_2d_{\mathcal{V}}) + \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} (A_1\mathbf{x}_1 + B_1u_{\mathcal{V}} + D_1\mathbf{d}_1) \leq 0.$$

□

The simulation relation error is reduced with the proposed disturbance refinement

method since P and R_1 can be designed to reduce ϱ_1 . This is an improvement on robust simulation function where ϱ_1 depends only on D_1 and λ (as $D_2 = 0$).

Remark 7.3. *The multiplication of unknown matrices, e.g. P with Q_1 , gives a bilinear matrix inequality. To resolve the bilinearity, setting $Q_1 = K_1 = 0$ results in a simpler interface function only considering the disturbance \mathbf{v}_1 . Using LMI solvers, other parameters can then be optimised. The original form (7.7b) is left in the proofs for generality.*

Remark 7.4. *Another approach to resolve the bilinearity is to fix one term and solve for the other term, then iterate by swapping which term is fixed. This may converge toward an optimal solution.*

The constructed \mathcal{V} in Theorem 7.1 can be leveraged to quantify the mismatch between output trajectories of Σ_1 and Σ_2 with measurable disturbances as presented in the next theorem.

Theorem 7.2. *Consider two systems of the form (7.4). Let \mathcal{V} be a robust simulation function from Σ_2 to Σ_1 with its associated interface functions $u_{\mathcal{V}}$ and $d_{\mathcal{V}}$. Let $\mathbf{u}_2(t)$ be an admissible input of Σ_2 and $\mathbf{x}_1(t)$ be a state trajectory of Σ_1 satisfying*

$$\dot{\mathbf{x}}_1 = A_1 \mathbf{x}_1 + B_1 u_{\mathcal{V}} + D_1 \mathbf{v}_1. \quad (7.10)$$

Then

$$\|\mathbf{y}_1(t) - \mathbf{y}_2(t)\| \leq \max\{\mathcal{V}(\mathbf{x}_1(0), \mathbf{x}_2(0)), \varrho_1(\|\mathbf{v}_1\|_{\infty}) + \varrho_2(\|\mathbf{u}_2\|_{\infty})\}.$$

Proof. For the sake of an easier presentation, notation is slightly abused to denote $\mathcal{V}(\mathbf{x}_1(t), \mathbf{x}_2(t))$ by $\mathcal{V}(t)$. Let

$$\epsilon = \max\{\mathcal{V}(0), \gamma_1(\|\mathbf{d}_1\|_{\infty}) + \gamma_2(\|\mathbf{u}_2\|_{\infty})\}.$$

First, show $\mathcal{V}(t) \leq \epsilon$ for all t . As (7.10) involves a feedback composition, it is assumed the composition is well-defined and for any initial state there exists a unique solution defined on the interval $t \subseteq \mathbb{R}^+$. Showing $\mathcal{V}(0) \leq \epsilon$ is straightforward due to the definition of ϵ . Assume there exists $\tau > 0$ such that $\mathcal{V}(\tau) > \epsilon$. Then there also exists some $0 \leq \tau' < \tau$ such that $\mathcal{V}(\tau') = \epsilon$ and $\forall t \in (\tau', \tau], \mathcal{V}(t) > \epsilon$. Note that one has, $\forall t \in (\tau', \tau]$,

$$\gamma_1(\|\mathbf{d}_1\|) + \gamma_2(\|\mathbf{u}_2\|) \leq \gamma_1(\|\mathbf{d}_1\|_{\infty}) + \gamma_2(\|\mathbf{u}_2\|_{\infty}) \leq \epsilon < \mathcal{V}(t).$$

From (7.3), one then has $\frac{\partial \mathcal{V}(t)}{\partial t} \leq 0$ for all $t \in (\tau', \tau]$, which implies

$$\mathcal{V}(\tau) - \mathcal{V}(\tau') = \int_{\tau'}^{\tau} \frac{\partial \mathcal{V}(t)}{\partial t} dt \leq 0.$$

This contradicts $\mathcal{V}(\tau) > \epsilon = \mathcal{V}(\tau')$. Therefore, $\mathcal{V}(t) \leq \epsilon, \forall t$. Finally from (7.2) one has:

$$\mathcal{V}(\mathbf{x}_1(t), \mathbf{x}_2(t)) \leq \epsilon \implies \|\mathbf{y}_1(t) - \mathbf{y}_2(t)\| \leq \epsilon.$$

□

7.4 Case Study

To show the efficacy of the proposed approach, a model of the *New England 39-Bus Test System* (NETS) is employed which is similar in design to the three-control area power system in [16,144]. NETS has 10 machines, 39 buses and three areas. In this chapter, just one area of this model is considered, containing 9 states with one input and one disturbance. The single line diagram for this system is depicted in Fig. 40. A linear model for Area 1 of NETS is acquired using the Simulink Model Linearizer on the closed-loop system.

The large disturbance \mathbf{v}_1 is assumed to be measurable in the power system domain as the disturbance may represent changes in the behaviour of generation and load components, e.g., generators, plug-in electric vehicles (EVs) and energy storage systems (ESSs). The generation or load values of these components may be known to operators and the connection and disconnection of these components could be tracked through sensors in a smart grid. It is assumed there is access to a fleet of EVs which can connect/disconnect from the power grid almost instantaneously. Such responsive loads are flexible and can be used for load shedding and frequency regulation of smart grids [209].

The dynamics of the model can be presented as a linear system Σ_1 equivalent to (7.4). A power loss disturbance of 1 per unit (100 MW, equivalent to a typical generator or 35,000 households) is applied to Σ_1 in all the scenarios of this case study. The abstract system Σ_2 is constructed using MATLAB's *balreal* function by truncating the matrices to a reduced-state order of 3. YALMIP [116] and MOSEK [132] are employed to solve LMIs and optimisations in MATLAB on a macOS machine with 8 GB RAM and Intel Core i5 Processor. Simulations are run over a time horizon of 6 seconds, with a time step of 0.005 seconds.

7.4.1 System Specification

For this system, consider a specification for primary frequency control. The frequency f can deviate away from its steady state value $f_0 = 50\text{Hz}$, this deviation is denoted by $\Delta f = f - f_0$. Two regions are bounded that the frequency deviation should never transition into; $\mathcal{A}_{ub} = (0.5, +\infty)$ and $\mathcal{A}_{lb} = (-\infty, -0.35)$. Additionally, whenever there are deviations, it should come back the target range $\mathcal{T} = [-0.3, 0.5]$. Therefore the desired system behaviour can be described by the LTL formula:

$$\psi = \square(\psi_1 \wedge \psi_2) \text{ with } \psi_1 = \diamond\mathcal{T}, \psi_2 = \neg(\mathcal{A}_{ub} \vee \mathcal{A}_{lb}). \quad (7.11)$$

The specification is modified in (7.11) appropriately with the error ϵ of the robust simulation function to get a more conservative specification $\hat{\psi}$ on Σ_2 . This modification ensures that whenever Σ_2 satisfies $\hat{\psi}$, then Σ_1 satisfies ψ by applying the appropriate input and disturbance interface functions for refining the controller. Then, one has the modified specification

$$\hat{\psi} = \square(\hat{\psi}_1 \wedge \hat{\psi}_2) \text{ with } \hat{\psi}_1 = \diamond\hat{\mathcal{T}}, \hat{\psi}_2 = \neg(\hat{\mathcal{A}}_{ub} \vee \hat{\mathcal{A}}_{lb}), \quad (7.12)$$

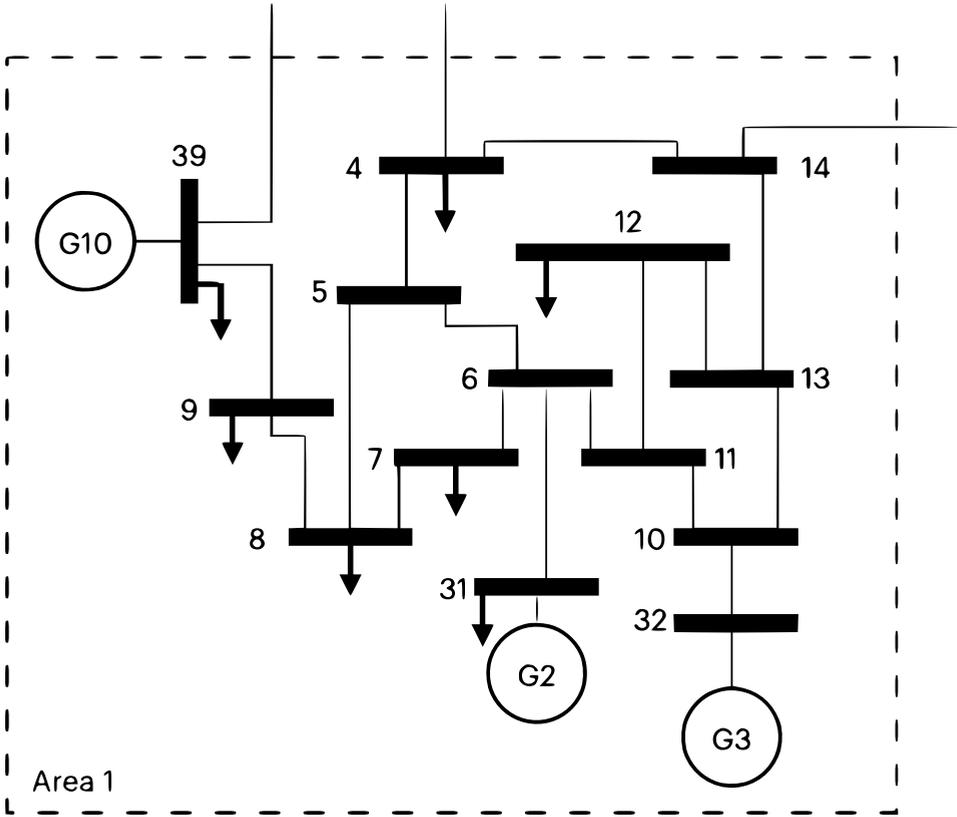


Figure 40: A single line diagram of Area 1 of the New England 39-Bus Test System.

with $\hat{\mathcal{T}} = [-0.3 + \epsilon, 0.5 - \epsilon]$, $\hat{\mathcal{A}}_{ub} = (0.5 - \epsilon, +\infty)$ and $\hat{\mathcal{A}}_{lb} = (-\infty, -0.35 + \epsilon)$.

7.4.2 Simulation Relation Error

The primary goal of employing robust simulation functions is to construct an abstract system Σ_2 which is ϵ -close to the concrete system Σ_1 , where ϵ remains small enough. Note that in the modified specification (7.12), any value $\epsilon \geq 0.4$ results in $\hat{\mathcal{T}} = \emptyset$ and the set of controllers enforcing the specification becomes empty. Therefore, the approximation approach must provide error thresholds small enough to give a feasible controller on the abstract system.

7.4.3 Uncontrolled system.

If the response of EVs is not included in the system ($\mathbf{u}_1 = 0$), the open-loop Σ_1 has the maximum frequency deviation of $\Delta f = -0.6872\text{Hz}$, which clearly violates the specification ψ . Therefore, the contribution of EVs is essential to satisfy the specification on the frequency.

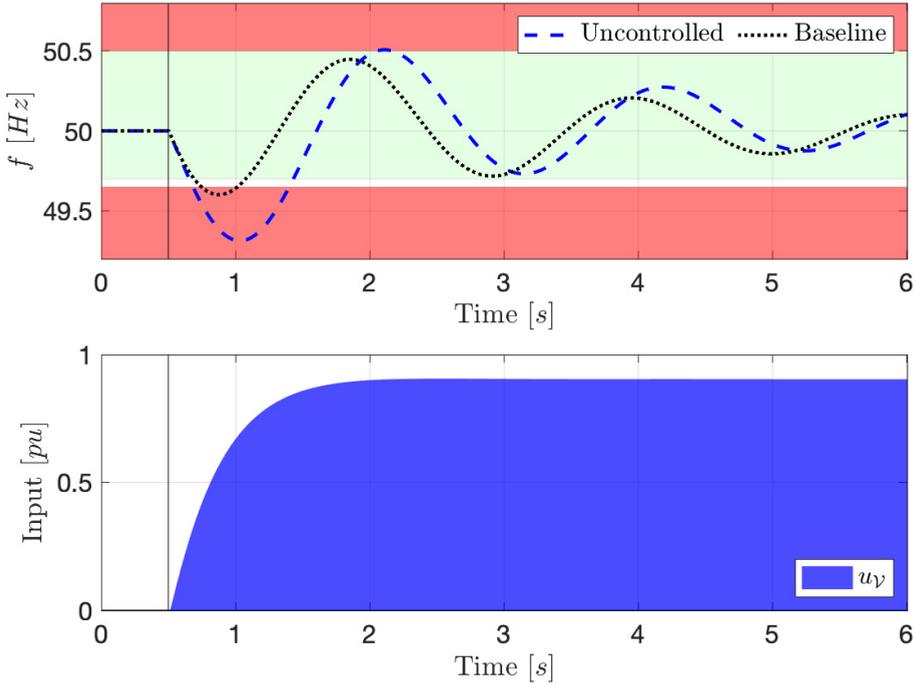


Figure 41: **Top.** Target range \mathcal{T} is shown in green, \mathcal{A}_{ub} and \mathcal{A}_{lb} are shown in red as two regions that the system should never transition into. The baseline controller notably improves the frequency response of the system in compare with the uncontrolled system. However, both curves still fall into the red unsafe region. **Bottom.** The input u_γ is a byproduct of the simulation relation interface keeping Σ_1 and Σ_2 ϵ -close. No controller is synthesised over Σ_2 , so $\mathbf{u}_2 = 0$.

7.4.4 Abstraction without disturbance refinement.

The error threshold ϵ is minimised under the assumption of no disturbance refinement ($D_2 = 0$), $\lambda = 1.7$, $\|\mathbf{u}_2\|_\infty = 0.5$, and $0.01 \mathbb{I}_9 \leq \bar{M} \leq 120 \mathbb{I}_9$. This gives the value $\epsilon_{\min} = 3.9156$, which makes the specification $\hat{\psi}$ unsatisfiable.

7.4.5 Abstraction with disturbance refinement.

Now the approach from Theorems 7.1–7.2 is used with the proposed disturbance interface function. It is assumed λ and the bounds on \bar{M} and $\|\mathbf{u}_2\|$ are selected as before, $R_1 = 1$, and $Q_1 = K_1 = 0$. D_2 and B_2 are optimised to minimise (7.8) and (7.9), respectively. Accordingly, the value $\epsilon_{\min} = 0.1019$.

In both cases of the approach with and without disturbance refinement, the same matrices for Σ_2 are constructed. The only difference is that $D_2 = 0$ for the case without disturbance refinement.

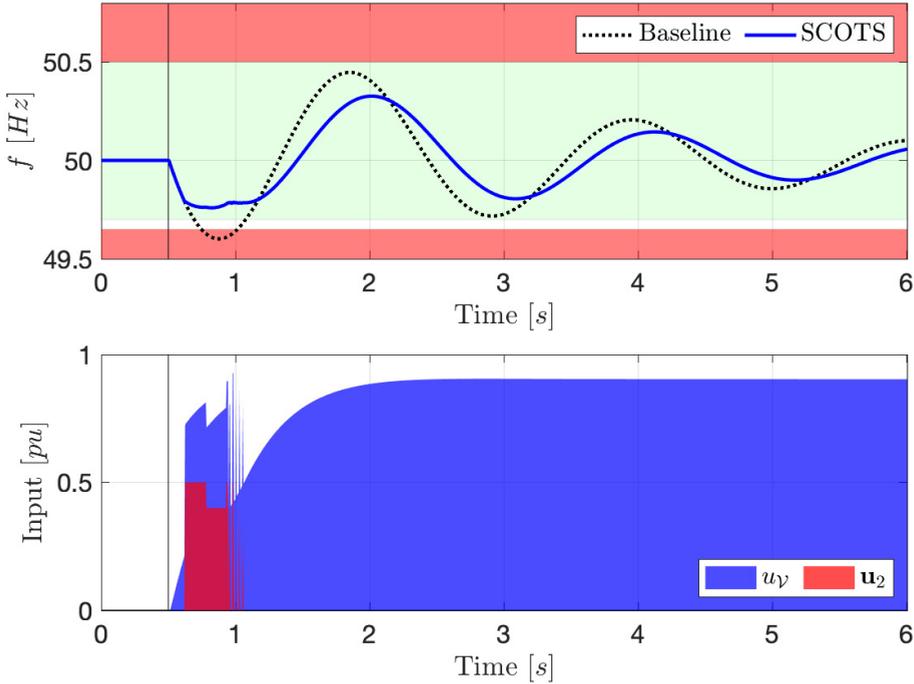


Figure 42: **Top.** Target range \mathcal{T} is shown in green, unsafe regions \mathcal{A}_{ub} and \mathcal{A}_{lb} are shown in red. The controller designed using SCOTS and the robust simulation function with disturbance refinement successfully satisfy ψ , compared with the baseline controller which violates the specification. **Bottom.** The control input u_2 designed using SCOTS for Σ_2 and the refined control input u_1 for Σ_1 using the robust simulation function.

7.4.6 Controller Synthesis Process

7.4.7 Baseline controller.

Consider the robust simulation function with the designed abstract system Σ_2 and the interface functions (7.7) but with $u_2 = 0$ in (7.7a). As Q_2 and K_2 are non-zero in (7.7a), control inputs are chosen automatically based on the current states of Σ_1 and Σ_2 to maintain the outputs of the two systems within distance ϵ . When the power system frequency moves away from its steady-state value, the input interface function u_γ generates a control input for Σ_1 , which is considered as the *baseline controller*. Fig. 41 shows the frequency response in Σ_1 without EV participation (uncontrolled system with $u_1 = 0$) against the baseline controller. Although the baseline controller reduces the frequency deviations, it is still unable to satisfy the required specification ψ .

7.4.8 Controller using robust simulation functions.

The constructed abstraction Σ_2 is employed as an appropriate substitute in the controller synthesis process. In particular, by knowing ϵ as the maximum error between outputs of Σ_1 and Σ_2 , a symbolic controller can be first designed for the reduced-order model Σ_2 to satisfy $\hat{\psi}$ and then be refined back to Σ_1 with the guarantee on satisfying ψ . To do so, the tool SCOTS [163] is used for the synthesis of the symbolic controller using a high-performance computer with 2 nodes and 11 GB memory per core, taking 55 minutes. Note that applying such a symbolic design directly to the 9-dimensional system Σ_1 is infeasible due the required exponentially large computational time and memory space.

Fig. 42 compares the baseline controller against the controller designed by combining the robust simulation function with SCOTS. The input u_2 designed by SCOTS is taken as the minimum value that guarantees satisfaction of the specification ψ (to use participation of EVs only if needed). Successful synthesis of the controller over Σ_2 by SCOTS proves formally that ψ holds on Σ_1 . Fig. 42 (bottom) shows that over the time interval $t \in [0.5, 1]$, the controller designed on Σ_2 takes non-zero values to bring back the frequency to the intended target region, thus enabling Σ_1 to satisfy ψ .

Overall, formal guarantees are provided using symbolic control over a 9-dimensional system while only requiring the computational load of a 3-dimensional system. Verifying Theorem 7.2, the maximum mismatch between the output trajectories of Σ_1 and Σ_2 is calculated from simulations. The value 0.6872 is acquired for the approach without disturbance refinement and 0.0449 for the approach with disturbance refinement. This confirms the theoretical error bounds ϵ for both cases.

7.5 Conclusion

This chapter extended the notion of simulation functions to its robust version by considering large disturbances in the dynamics and introducing an interface function for the disturbance refinement. In particular:

- Concrete systems were approximated with abstractions of lower dimensions (reduced-order models) and *robust* simulation functions were developed to consider the perturbation in the abstract system.
- The proposed approach enables controller design using a reduced-order form of the concrete system and reduces the computational load required for formal synthesis.
- The applicability of the approach is illustrated by synthesising a formal controller for a linear 9-state area of the known New England 39-Bus Test System, using only a 3-state abstract system.
- The method with disturbance refinement is compared against the one without disturbance refinement to validate the results.

The next chapter will further extend the concepts of this chapter to consider a class of non-linear systems and also compositionality of the different power system areas combining to control the full NETS system.

Assume-Guarantee Contracts for Compositional Control of Power Systems

This chapter is concerned with an assume-guarantee contract approach to compositionally control the *New England 39-bus Test System* (NETS). The proposed scheme is based on the work [206], and the results of the previous chapter looking at *robust simulation functions* (RSFs) with disturbance refinement. The composition of multiple subsystems can tackle difficulties associated with scalability, also known as the *curse of dimensionality*. In the proposed setting, concrete subsystems are approximated with abstractions with lower dimensions (*a.k.a.* reduced-order models), while providing mathematical guarantees over the controller synthesis. This is through the quantification of an upper bound on the closeness between output trajectories of the original systems and their reduced-order models. Two control methods are proposed to provide guarantees for NETS: one using the principle of interconnected synchronous machines and another considering the power flows in the network between neighbouring subsystems.

Notation. The notation $\|a\|$ is used for the Euclidean norm of a vector a , and $\|a\|_\infty$ for taking the Euclidean norm followed by a maximisation over the bounded domain of a . Intervals as subsets of real numbers are denoted by $\mathcal{B} = [\underline{\mathcal{B}}, \overline{\mathcal{B}}]$ where $\underline{\mathcal{B}}$ and $\overline{\mathcal{B}}$ are used for the lower and upper boundaries of the interval. Specifically, \mathcal{B} will denote the safe set and \mathcal{T} the target set. *Area* and *subsystem* are used interchangeably throughout this chapter.

8.1 Introduction

Cyber-physical systems (CPS) combine both cyber and physical components in interconnected models with interactions through feedback loops [106]. They are an important modelling framework for engineering real-life systems such as the power systems focused on in this thesis. The interconnection of these components in the models often results in high-dimensional systems with complex behaviour specifications that are generally safety critical in nature. Providing guarantees on the behaviour of these systems is therefore essential but also incredibly challenging. To tackle this difficulty, formal methods have been introduced in the relevant literature as a strong mathematical framework to provide guarantees on either verification or controller synthesis of CPS [99, 146].

Symbolic control is one of the promising techniques for formal control synthesis of CPS [188]. In particular, symbolic models (*a.k.a.* finite abstractions), see Def. 4.8, replace concrete systems to provide an easier medium for synthesis of a formal controller. In abstraction-based techniques, each discrete state and input in the finite abstraction maps to an aggregate collection of continuous states in the original (concrete) model. By establishing a similarity relation between original systems and their symbolic models, one can consider the abstract system as an appropriate substitute in the controller design process with lower computational complexity while still preserving closeness guarantees between the two systems.

Simulation and bisimulation functions are powerful techniques to relate output trajectories of abstract systems to those of concrete ones [19, 188], see also Chapter 4. If a concrete system is (bi)similar to an abstract system, only the abstract system needs to be considered in the formal synthesis process, while guarantees are still provided. For control systems where output trajectories of two systems may not be identical, *approximate* (bi)simulation functions [67] have been developed in which output trajectories of two systems are only required to remain measurably close. In this case, the closeness between output trajectories through time can be bounded by some maximal ϵ , known as the simulation relation error. Given an ϵ -closeness, interface functions can be used to map the synthesised controller from the abstract system back to the concrete one. In [96], this type of relation is extended to *robust* simulation functions (RSF) with small disturbances inside the concrete system, but with an unperturbed abstract system.

Abstraction-based techniques often suffer severely from the *curse of dimensionality* while dealing with high-dimensional systems [77]. To alleviate this computational complexity, one potential approach is to use compositional techniques: decompose a large-scale system into multiple subsystems and provide analysis over the high-dimensional system via its smaller subsystems [88]. Assume-guarantee contracts have been explored extensively in the literature to provide control techniques over a network of continuous-time dynamical systems [169]. Compositional approaches have also been used for the construction of (in)finite abstractions for interconnected systems based on abstractions of smaller subsystems [100, 101, 148, 190].

Power networks are a demanding application of CPS that have received remarkable attention in the past decade. In particular, as the contribution of renewable

energy rises, power networks are becoming increasingly intermittent. To ensure stability and functionality of power networks, demand-side control techniques are required [28]. In this respect, smart grid control involves the demand-side of a power grid responding to events in order to reduce the strain on the network, while also optimising consumer satisfaction and other specialist requirements [95]. Smart grids contain sensors and information-based technical devices, so it assumes that the current frequency, power generation or load values applied in different locations of the system can be accurately measured.

Formal methods play a significant role in power systems to provide formal analysis over this type of demanding systems. In this regard, the work [181] proposes approximate bisimulations in transient power systems and employs differential-algebraic equations (DAEs) to model the New England 39-Bus Test System. In [9], DAEs are utilised as models of the IEEE 14-Bus System and the IEEE 30-Bus System to provide reachability analysis for transient stability without performing any controller synthesis. The work [110], studies formal analysis of power systems via reachable sets of microgrids with distributed energy resources. The results of [218] use contract-based symbolic controller design for voltage stability in DC microgrids.

8.2 Original Contributions.

In this chapter, the notion of *robust simulation functions* (RSFs) with disturbance refinement is generalised from linear systems to a class of nonlinear systems. An assume-guarantee contracts approach with RSF for the control of an interconnected network composed of several subsystems is provided. Given the employed assume-guarantee contracts with RSF, the efficacy of the results are demonstrated on the *New England 39-bus Test System* (NETS), as a large closely-coupled benchmark test system, composed of three 9-dimensional subsystems (totally 27 dimensions). Model-order reduction techniques are leveraged to construct a 3-state reduced-order model for each subsystem (totally 9 dimensions) to further mitigate the curse of dimensionality. A set of temporal logic specifications are provided for the GB power network. The results for primary frequency control are demonstrated using two scenarios: (i) leveraging the principle of interconnected synchronous machines to control isolated subsystems, and (ii) considering internal disturbances in the network between different subsystems to provide accurate controls using shared information of neighbouring frequencies. For the sake of better illustrations of the results, the complex case study is presented as a running example throughout this chapter.

A limited subset of the proposed results of this chapter are presented in Chapter 7. The results of the previous chapter are extended in three main directions. First and foremost, instead of considering only a single subsystem of NETS, an interconnected network of these subsystems is studied that are then controlled compositionally. Secondly, the theoretical results of Chapter 7 are generalised from simple linear systems to a class of nonlinear control systems. Finally, the results are applied to the New England 39-bus Test System, as a highly challenging large-scale

closely-coupled system, which is significantly more complex than the case study in Chapter 7. The approach of Chapter 7 cannot cope with the case study in this chapter due to the scalability limitations caused by the curse of dimensionality.

The chapter is organised as follows. Preliminaries and the class of systems are provided in Section 8.3, this section also introduces NETS as a running case study through this chapter. The GB power network frequency specifications are defined, expressed in linear temporal logic (LTL) in Section 8.4. The notion of RSF with disturbance refinement is generalised to a class of nonlinear systems in Section 8.5 and a proof of concept for the proposed technique is provided in Section 8.6. The interconnection of subsystems is presented in Section 8.7 and the methodology of assume-guarantee contracts in Section 8.8. Demonstrations of the proposed approaches for isolated areas and for compositional techniques with internal disturbances is provided in Sections 8.9 and 8.10, respectively. Finally, concluding remarks are provided in Section 8.11.

8.3 Preliminaries

Subsystems. Consider a network of N subsystems, where each subsystem i can be modelled by $\Sigma_z^i = (X_z^i, U_z^i, V_z^i, W_z^i, g_z^i, Y_{z_1}^i, Y_{z_2}^i, h_{z_1}^i, h_{z_2}^i)$, $z \in \{1, 2\}$, and $i \in \{1, \dots, N\}$, as in Def. 4.15. The evolution of subsystems can be characterised by

$$\Sigma_z^i: \begin{cases} \dot{\mathbf{x}}_z^i = g_z^i(\mathbf{x}_z^i, \mathbf{u}_z^i, \mathbf{v}_z^i, \mathbf{w}_z^i), \\ \mathbf{y}_{z_1}^i = h_{z_1}^i(\mathbf{x}_z^i), \\ \mathbf{y}_{z_2}^i = h_{z_2}^i(\mathbf{x}_z^i), \end{cases} \quad z \in \{1, 2\}, i \in \{1, \dots, N\}. \quad (8.1)$$

where $\mathbf{x}_z^i \in X_z^i$, $\mathbf{y}_{z_1}^i \in Y_{z_1}^i$, $\mathbf{y}_{z_2}^i \in Y_{z_2}^i$, $\mathbf{u}_z^i \in U_z^i$, $\mathbf{w}_z^i \in W_z^i$, and $\mathbf{v}_z^i \in V_z^i$. It is assumed \mathbf{v}_z^i are measurable and potentially large.

Without loss of generality, consider Σ_1^i as the original (concrete) subsystem and Σ_2^i as its (possibly) lower-dimensional abstraction (with $n_2^i \leq n_1^i$). In the following, the definition of interconnected systems is presented in which subsystems Σ_z^i are connected with each other via internal disturbances \mathbf{w}_z^i .

Using the definition of interconnected systems, Def. 4.16, the internal disturbances are constrained by

$$[\mathbf{w}_z^1; \dots; \mathbf{w}_z^N] = \mathcal{M}[\mathbf{y}_{z_2}^1; \dots; \mathbf{y}_{z_2}^N].$$

The evolution of the interconnected system is therefore characterised by

$$\Sigma_z: \begin{cases} \dot{\mathbf{x}}_z = g_z(\mathbf{x}_z, \mathbf{u}_z, \mathbf{v}_z), \\ \mathbf{y}_z = h_z(\mathbf{x}_z) \end{cases} \quad z \in \{1, 2\}.$$

Linear Temporal Logic Specifications. For dynamical systems in (8.1), consider linear temporal logic (LTL) specifications with syntax

$$\psi := \text{true} \mid p \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \bigcirc\psi \mid \psi_1 \cup \psi_2,$$

where p is the element of an atomic proposition, see also Def. 4.10. I refer the reader back to Chapter 4 for more details.

Running Case Study. The developed approach in this chapter is applied mainly to a model of New England 39-bus Test System (NETS) as a highly challenging and demanding power system. This model is similar to the three-control area power system in [144], where here it is presented as a running case study throughout this chapter for the sake of better illustration. NETS has 10 machines, 39 buses, 46 lines and three areas. One of the generators is used to represent the connection between the NETS power system and the wider American power network. This provides an especially challenging case study for the techniques proposed in this chapter.

NETS can be decomposed into three smaller areas (*a.k.a.* subsystems), each of which contains three generators. The interconnected system consists of 27 states, with 9 states in each area. In addition, possible external disturbances are considered together with inputs (one per area) that can be used for control synthesis purposes, *e.g.*, Energy Storage Systems (ESSs) or Plug-in Electric Vehicles (EVs). Internal disturbances are defined here as power dynamics affecting a local area, *i.e.*, own subsystem, caused by neighbouring areas *i.e.*, other subsystems.

In this running case study, the main goal is to formally control NETS. Given that formal control approaches often struggle with scalability, model order reduction techniques are employed together with RSF with disturbance refinement to reduce the number of states while providing mathematical guarantees for the system behaviour. Reducing NETS from 27 states to lower dimensions introduces a reduction error ϵ which is generally very large. Therefore, a compositional technique is employed to first decompose the NETS into three 9-dimensional areas and then reduce the dimension of each area via constructing reduced-order abstractions.

It is worth highlighting that although each decomposed area of NETS has 9 states, this is still an intractable problem due to the *curse of dimensionality* during the synthesis procedure. To resolve this issue, the first aim is building a reduced-model abstraction with 3 states for each area and then constructing an RSF with disturbance refinement as a relation between each concrete area and its reduced-order model. To maintain the interconnection of all areas, the frequency of neighbouring areas are used as internal disturbances of the local area. The following relationship between frequencies is used to connect these areas

$$\frac{2\pi}{s} \sum_{j=1}^N T^{ij} (f^i - f^j), \quad (8.2)$$

where f^i is the local area frequency and f^j are neighbouring area frequencies, T^{ij} are constants related to the power interchange between the respective neighbours. By leveraging the principle of interconnected synchronous machines [95], one can assume that for all neighbours, $f^i = f^j$. This assumption results in (8.2) to be zero so that each area is simplified to 9 states with no internal disturbances. The linear dynamics of NETS are acquired using the Simulink Model Linearizer on the closed-loop system.

8.4 Frequency Specifications

In this section, the requirements on frequency regulation of the GB power grid are defined using LTL specifications. The requirements are compiled from the collection of references [69, 138–140, 142, 173] as follows. The nominal frequency of the GB power grid is $f_0 = 50 \text{ Hz}$. The frequency f should remain between the statutory limits $\mathcal{S} = [\underline{\mathcal{S}}, \overline{\mathcal{S}}]$ with $\underline{\mathcal{S}} = 49.5 \text{ Hz}$ and $\overline{\mathcal{S}} = 50.5 \text{ Hz}$, for all losses up to the maximum *normal infeed loss* ($\mathcal{L} = 1320 \text{ MW}$):

$$\psi_{normal} := [loss \leq \mathcal{L}] \implies \square[f \in \mathcal{S}].$$

Losses greater than 1320 MW are considered *infrequent infeed losses* and may fall below the statutory limits briefly, but no lower than the containment zone value $\mathcal{Z} = 49.2 \text{ Hz}$. Within a time constraint of 60 seconds, the frequency should return to the statutory limits under the following conditions:

$$\begin{aligned} \psi_{infrequent} := [loss \geq \mathcal{L} \wedge f < \underline{\mathcal{S}}] \implies \\ [\diamond^{60}(f \in \mathcal{S}) \wedge \square(f \geq \mathcal{Z})]. \end{aligned}$$

If the frequency rises above 52 Hz or falls below 47 Hz , *i.e.*, there is system shutdown, which should be avoided at all costs:

$$\psi_{shutdown} := \square[47 \leq f \leq 52].$$

Additionally, the GB power grid specifies certain minimum time constraints for the devices contributing to primary, secondary and high frequency response services, as discussed next.

Firm Frequency Response (FFR). For devices contributing to *primary frequency response*, it is necessary to inject power ($\mathcal{P}_p > 0 \text{ MW}$) within 2 seconds of a *low frequency event* (\mathcal{E}_{low}), and provide maximum power (\mathcal{P}_p^{max}) by 10 seconds. This maximum power must be at least 1 MW per response device or aggregated load. This delivery should be maintained for 30 seconds:

$$\begin{aligned} \psi_{p1} &:= \mathcal{E}_{low} \implies [\diamond_2(\mathcal{P}_p > 0) \wedge \diamond^{10}(\mathcal{P}_p = \mathcal{P}_p^{max})], \\ \psi_{p2} &:= [\mathcal{P}_p = \mathcal{P}_p^{max}] \implies \square^{30}[\mathcal{P}_p = \mathcal{P}_p^{max}], \\ \psi_p &:= \psi_{p1} \wedge \psi_{p2}. \end{aligned} \tag{8.3}$$

For devices contributing to *secondary frequency response*, it is essential to begin injecting maximum power (\mathcal{P}_s^{max}) within 30 seconds of a low frequency event. Similarly, the delivery (\mathcal{P}_s) should be maintained for 30 minutes:

$$\begin{aligned} \psi_{s1} &:= \mathcal{E}_{low} \implies \diamond^{30}[\mathcal{P}_s = \mathcal{P}_s^{max}], \\ \psi_{s2} &:= [\mathcal{P}_s = \mathcal{P}_s^{max}] \implies \square^{1800}[\mathcal{P}_s = \mathcal{P}_s^{max}], \\ \psi_s &:= \psi_{s1} \wedge \psi_{s2}. \end{aligned} \tag{8.4}$$

It is possible for devices to perform both primary and secondary response services

in the power grid with the following specification:

$$\psi_{ps} := \psi_p \wedge \psi_s.$$

Equivalent *high frequency response* specifications for any *high frequency event* (\mathcal{E}_{high}) are similar to both (8.3) and (8.4) but without a fixed delivery duration.

Enhanced Frequency Response (EFR). Taking advantage of the fast response capabilities of *energy storage systems* (ESSs), enhanced frequency response (EFR) is designed to allow *state-of-charge* (SoC) management which is not possible with FFR. ESS should respond within 1 second of the frequency crossing the deadband threshold which can be set at $db = [49.95, 50.05]$ Hz for a wide deadband and $db = [49.985, 50.015]$ Hz for a narrow deadband. The EFR service provided must be bidirectional, *i.e.*, both exported and imported to/from the grid. It must be possible for the EFR service to be provided at 100% capacity (\mathcal{P}_{EFR}^{max}) for a minimum of 15 minutes. To avoid short-term frequency instability issues from the fast response, ramp-rate limitations have been included in the specification when the frequency is inside the envelope but outside of the deadband. The ramp-rate limitations are included to limit short-term stability problems [69]. The maximum change in output is limited as a proportion of the rate of change of the frequency (ROCOF or $\frac{\partial f}{\partial t}$). The ramping constant k is 0.45 for the wide deadband and 0.485 for the narrow deadband:

$$\begin{aligned} \mathcal{P}_{EFR}^{max} \left(-\frac{1}{k} \frac{\partial f}{\partial t} - 0.01 \right) &< \frac{\partial \mathcal{P}}{\partial t} < \mathcal{P}_{EFR}^{max} \left(-\frac{1}{k} \frac{\partial f}{\partial t} + 0.01 \right), \\ \psi_{efr}^1 &:= [f \notin db] \implies \diamond_1 [\mathcal{P} = \mathcal{P}_{EFR}^{max}], \\ \psi_{efr}^2 &:= [\mathcal{P} = \mathcal{P}_{EFR}^{max}] \implies \square_{900} [\mathcal{P} = \mathcal{P}_{EFR}^{max}], \\ \psi_{efr} &:= \psi_{efr}^1 \wedge \psi_{efr}^2. \end{aligned}$$

Running case study (continued): Consider a stricter primary frequency specification, in which the frequency f can deviate away from its steady state value f_0 , the deviation is denoted by $\Delta f = f - f_0$. Two regions can be bounded that the frequency deviation should never transition into, $\mathcal{A}_{ub} = (\bar{B}, +\infty)$ and $\mathcal{A}_{lb} = (-\infty, \underline{B})$. Additionally, whenever there are deviations, the frequency should return to the target range $\mathcal{T} = [\underline{T}, \bar{T}]$. The desired system behaviour can be described by the following LTL formulae:

$$\psi = \square(\psi_1 \wedge \psi_2) \text{ with } \psi_1 = \diamond \mathcal{T}, \psi_2 = \neg(\mathcal{A}_{ub} \vee \mathcal{A}_{lb}).$$

This specification is modified appropriately with the simulation relation error ϵ to acquire a conservative specification $\hat{\psi}$ over Σ_2 as:

$$\hat{\psi} = \square(\hat{\psi}_1 \wedge \hat{\psi}_2) \text{ with } \hat{\psi}_1 = \diamond \hat{\mathcal{T}}, \hat{\psi}_2 = \neg(\hat{\mathcal{A}}_{ub} \vee \hat{\mathcal{A}}_{lb}), \quad (8.5)$$

with $\hat{\mathcal{T}} = [\underline{T} + \epsilon, \bar{T} - \epsilon]$, $\hat{\mathcal{A}}_{ub} = (\bar{B} - \epsilon, +\infty)$ and $\hat{\mathcal{A}}_{lb} = (-\infty, \underline{B} + \epsilon)$. This modification ensures that whenever the abstract system Σ_2 satisfies $\hat{\psi}$, the concrete system Σ_1 satisfies the original specification ψ by applying appropriate input and disturbance interface functions for refining the controller.

8.5 Simulation Functions

In this work, the notion of robust simulation functions is leveraged to construct an abstract system which is ϵ -close to the concrete one, where ϵ remains small enough. In the following subsection, it is shown how incorporating the disturbance of the concrete system into the abstract one, through an interface function $d_{\mathcal{V}}$, can further reduce the simulation relation error ϵ between Σ_1 and Σ_2 . This enables one to perform controller synthesis on the abstract domain and refine it back over potentially high-dimensional original system while improving the scalability of the control scheme.

8.5.1 Robust Simulation Function with Disturbance Refinement

Given the system in (8.1), the definition of a robust simulation function \mathcal{V} with two interface functions $u_{\mathcal{V}}$ and $d_{\mathcal{V}}$ is formalised as the following.

Definition 8.1 (Robust Simulation Functions). *Consider two systems of the form (8.1). Let $\mathcal{V} : X_1 \times X_2 \rightarrow \mathbb{R}^+$ be a smooth function, $u_{\mathcal{V}} : U_2 \times X_1 \times X_2 \rightarrow U_1$ and $d_{\mathcal{V}} : V_1 \times W_1 \times X_1 \times X_2 \rightarrow V_2 \times W_2$ be continuous functions. Then the function \mathcal{V} is called a robust simulation function (RSF) from Σ_2 to Σ_1 and $u_{\mathcal{V}}$, $d_{\mathcal{V}}$ are its associated interface functions if there exist class- κ functions ϱ_1 and ϱ_2 such that for all $\mathbf{x}_1 \in X_1$, $\mathbf{x}_2 \in X_2$,*

$$\|h_1(\mathbf{x}_1) - h_2(\mathbf{x}_2)\| \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2), \quad (8.6)$$

for any $\mathbf{u}_2 \in U_2$ and $\mathbf{d}_1 \in [V_1 \ W_1]^T$ satisfying $\varrho_1(\|\mathbf{d}_1\|) + \varrho_2(\|\mathbf{u}_2\|) \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2)$, then

$$\begin{aligned} & \frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} f_2(\mathbf{x}_2, \mathbf{u}_2, d_{\mathcal{V}}(\mathbf{d}_1, \mathbf{x}_1, \mathbf{x}_2)) + \\ & \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} f_1(\mathbf{x}_1, u_{\mathcal{V}}(\mathbf{u}_2, \mathbf{x}_1, \mathbf{x}_2), \mathbf{d}_1) \leq 0. \end{aligned} \quad (8.7)$$

So, Σ_1 robustly approximately simulates Σ_2 if there exists an RSF \mathcal{V} from Σ_2 to Σ_1 .

The next subsection focuses on a class of nonlinear control systems with potentially large measurable disturbances and proposes an approach to construct its reduced-dimensional abstractions together with an RSF as presented in Definition 8.1.

8.5.2 Class of Nonlinear Systems under Large Measurable Disturbance

Here, the focus is on a class of nonlinear control systems with (potentially large) measurable disturbances. A model in this class and its abstraction are specified by

$$\Sigma_z : \begin{cases} \dot{\mathbf{x}}_z = & A_z \mathbf{x}_z + B_z \mathbf{u}_z + G_z \mathbf{v}_z \\ & + S_z \mathbf{w}_z + E_z \phi(F_z \mathbf{x}_z), \quad z \in \{1, 2\}, \\ \mathbf{y}_z = & C_z \mathbf{x}_z, \end{cases} \quad (8.8)$$

where $A_z \in \mathbb{R}^{n_z \times n_z}$, $B_z \in \mathbb{R}^{n_z \times p}$, $C_z \in \mathbb{R}^{m \times n_z}$, $G_z \in \mathbb{R}^{n_z \times q}$, $S_z \in \mathbb{R}^{n_z \times r}$, $E_z \in \mathbb{R}^{n_z \times 1}$, $F_z \in \mathbb{R}^{1 \times n_z}$. In addition, $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear term satisfying the following slop restriction:

$$a \leq \frac{\phi(c) - \phi(d)}{c - d} \leq b, \quad \forall c, d \in \mathbb{R}, c \neq d. \quad (8.9)$$

Remark 8.1. Note that if $E_1 = 0$ and/or $F_1 = 0$ in (8.8), the proposed approach simplifies to the one provided in Chapter 7 for the class of linear control systems with potentially large measurable disturbances.

Let $\mathbf{d}_z = [\mathbf{v}_z \ \mathbf{w}_z]^T \in [V_z \ W_z]^T$ and $D_z = [G_z \ S_z] \in \mathbb{R}^{n_z \times (q+r)}$ be the concatenation of external and internal disturbances. It is assumed \mathbf{d}_1 is a measured disturbance having some known bound $\|\mathbf{d}_1\|_\infty \leq d_{\max}$. Moreover, \mathbf{d}_2 is derived from \mathbf{d}_1 with the interface function $d_\mathcal{V}$ (cf. (8.12b)). The main to solve in this chapter problem is now presented.

Problem Description 8.1. Given a nonlinear system Σ_1 under large measurable disturbances and an LTL specification ψ , construct its reduced-dimensional abstraction Σ_2 together with an RSF as presented in Definition 8.1. Leverage the constructed abstraction Σ_2 and design a formal controller through simulation relations with disturbance refinement such that the specification is satisfied over the original system. Assume that the (potentially) large disturbance \mathbf{v}_1 is measurable in the power system.

In order to address Problem 8.1, the following lemma and theorems are raised.

Lemma 8.1. If Σ_1 is stabilisable, there are matrices $K_1, K_2, P, D_2, Q_1, L_{11}, L_{21}$ such that H is Hurwitz, and there exist a positive-definite matrix M and a positive constant λ such that the following matrix inequalities hold:

$$C_1^T C_1 \leq M, \quad (8.10a)$$

$$H^T M + M H \leq -2\lambda M, \quad (8.10b)$$

where

$$H = (A_1 - P D_2 (K_1 + Q_1) + B_1 K_2) + \bar{\delta} (E_1 + B_1 L_{21} - P D_2 L_{11}) F_1.$$

Here $\bar{\delta}$ is an upper bound of δ , where δ is a scalar in the interval $[a, b]$, in $\phi(F_1 \mathbf{x}_1) - \phi(F_1 P \mathbf{x}_2) = \delta F_1 (\mathbf{x}_1 - P \mathbf{x}_2)$ obtained from the slope restriction (8.9).

Utilising Lemma 8.1, the next theorem to construct an RSF is proposed.

Theorem 8.1. Consider two systems of the form (8.8). Assume that Σ_1 is stabilisable, a feedback gain K_1 exists for Σ_2 and that there exist matrices $P, K_2, Q_1, Q_2, L_{11}, L_{12}, L_{21}$ and L_{22} such that the following matrix equalities hold:

$$C_2 = C_1 P, \quad (8.11a)$$

$$F_2 = F_1 P, \quad (8.11b)$$

$$A_1 P + B_1 Q_2 = P A_2 + P D_2 Q_1 P, \quad (8.11c)$$

$$E_1 = PE_2 - B_1(L_{21} - L_{22}) + PD_2(L_{11} - L_{12}). \quad (8.11d)$$

Then \mathcal{V} in the form of

$$\mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) = \sqrt{(\mathbf{x}_1 - P\mathbf{x}_2)^T M (\mathbf{x}_1 - P\mathbf{x}_2)}$$

is an RSF from Σ_2 to Σ_1 with its associated interfaces

$$\begin{aligned} u_{\mathcal{V}} &= K_2(\mathbf{x}_1 - P\mathbf{x}_2) + Q_2\mathbf{x}_2 + R_2\mathbf{u}_2 \\ &\quad + L_{21}\phi(F_1\mathbf{x}_1) - L_{22}\phi(F_1P\mathbf{x}_2), \end{aligned} \quad (8.12a)$$

$$\begin{aligned} d_{\mathcal{V}} &= K_1(\mathbf{x}_1 - P\mathbf{x}_2) + Q_1\mathbf{x}_1 + R_1\mathbf{d}_1 \\ &\quad + L_{11}\phi(F_1\mathbf{x}_1) - L_{12}\phi(F_1P\mathbf{x}_2). \end{aligned} \quad (8.12b)$$

In addition, the class- κ functions ϱ_1 and ϱ_2 are designed as

$$\varrho_1(\nu) = \frac{\|\sqrt{M}(D_1 - PD_2R_1)\|}{\lambda} \nu, \quad (8.13)$$

$$\varrho_2(\nu) = \frac{\|\sqrt{M}(B_1R_2 - PB_2)\|}{\lambda} \nu, \quad (8.14)$$

where R_1 and R_2 are some arbitrary matrices of appropriate dimensions, and M, λ are matrices satisfying (8.10).

Proof. From (8.10a) and (8.11a), then

$$\mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) \geq \sqrt{(\mathbf{x}_1 - P\mathbf{x}_2)^T C_1^T C_1 (\mathbf{x}_1 - P\mathbf{x}_2)} = \|C_1\mathbf{x}_1 - C_2\mathbf{x}_2\|,$$

implying that condition (8.6) holds. Now, to show condition (8.7) holds, as well. Using (8.10b) and (8.11b)-(8.11d), one has

$$\begin{aligned} &\frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} g_2(A_2\mathbf{x}_2 + B_2\mathbf{u}_2 + D_2d_{\mathcal{V}} + E_2\phi(F_2\mathbf{x}_2)) \\ &\quad + \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} g_1(A_1\mathbf{x}_1 + B_1u_{\mathcal{V}} + D_1\mathbf{d}_1 + E_1\phi(F_1\mathbf{x}_1)) \\ &\leq -\lambda \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) + \|\sqrt{M}(D_1 - PD_2R_1)\mathbf{d}_1 + \sqrt{M}(B_1R_2 - PB_2)\mathbf{u}_2\| \\ &\leq -\lambda \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2) + \|\sqrt{M}(D_1 - PD_2R_1)\| \|\mathbf{d}_1\| + \|\sqrt{M}(B_1R_2 - PB_2)\| \|\mathbf{u}_2\|. \end{aligned}$$

Therefore, given that

$$\frac{\|\sqrt{M}(D_1 - PD_2R_1)\|}{\lambda} \|\mathbf{d}_1\| + \frac{\|\sqrt{M}(B_1R_2 - PB_2)\|}{\lambda} \|\mathbf{u}_2\| \leq \mathcal{V}(\mathbf{x}_1, \mathbf{x}_2),$$

then

$$\begin{aligned} \frac{\partial \mathcal{V}}{\partial \mathbf{x}_2} g_2(A_2 \mathbf{x}_2 + B_2 \mathbf{u}_2 + D_2 d_\gamma + E_2 \phi(F_2 \mathbf{x}_2)) + \\ \frac{\partial \mathcal{V}}{\partial \mathbf{x}_1} g_1(A_1 \mathbf{x}_1 + B_1 u_\gamma + D_1 \mathbf{d}_1 + E_1 \phi(F_1 \mathbf{x}_1)) \leq 0, \end{aligned}$$

which completes the proof. \square

The constructed \mathcal{V} in Theorem 8.1 is leveraged to quantify the mismatch between output trajectories of Σ_1 and Σ_2 with measurable disturbances, as presented in the next theorem.

Theorem 8.2. *Consider two systems of the form (8.8). Let \mathcal{V} be an RSF from Σ_2 to Σ_1 with its associated interface function u_γ . Let $\mathbf{u}_2(t)$ be an admissible input of Σ_2 and $\mathbf{x}_1(t)$ be a state trajectory of Σ_1 satisfying*

$$\dot{\mathbf{x}}_1 = A_1 \mathbf{x}_1 + B_1 u_\gamma + D_1 \mathbf{d}_1 + E_1 \phi(F_1 \mathbf{x}_1). \quad (8.15)$$

Then

$$\|\mathbf{y}_1(t) - \mathbf{y}_2(t)\| \leq \max\{\mathcal{V}(\mathbf{x}_1(0), \mathbf{x}_2(0)), \varrho_1(\|\mathbf{d}_1\|_\infty) + \varrho_2(\|\mathbf{u}_2\|_\infty)\}.$$

Proof. For the sake of an easier presentation, notation is slightly abused to denote $\mathcal{V}(\mathbf{x}_1(t), \mathbf{x}_2(t))$ by $\mathcal{V}(t)$. Let

$$\epsilon = \max\{\mathcal{V}(0), \varrho_1(\|\mathbf{d}_1\|_\infty) + \varrho_2(\|\mathbf{u}_2\|_\infty)\}.$$

Now, show $\mathcal{V}(t) \leq \epsilon$ for all t . Showing $\mathcal{V}(0) \leq \epsilon$ is straightforward due to the definition of ϵ . The rest of the proof is shown based on contradiction. Assume there exists $\tau > 0$ such that $\mathcal{V}(\tau) > \epsilon$. Then there also exists some $0 \leq \tau' < \tau$ such that $\mathcal{V}(\tau') = \epsilon$ and $\forall t \in (\tau', \tau], \mathcal{V}(t) > \epsilon$. Note that, $\forall t \in (\tau', \tau]$,

$$\varrho_1(\|\mathbf{d}_1\|) + \varrho_2(\|\mathbf{u}_2\|) \leq \varrho_1(\|\mathbf{d}_1\|_\infty) + \varrho_2(\|\mathbf{u}_2\|_\infty) \leq \epsilon < \mathcal{V}(t).$$

From (8.7), then $\frac{\partial \mathcal{V}(t)}{\partial t} \leq 0$ for all $t \in (\tau', \tau]$, which implies

$$\mathcal{V}(\tau) - \mathcal{V}(\tau') = \int_{\tau'}^{\tau} \frac{\partial \mathcal{V}(t)}{\partial t} dt \leq 0.$$

This contradicts $\mathcal{V}(\tau) > \epsilon = \mathcal{V}(\tau')$. Therefore, $\mathcal{V}(t) \leq \epsilon$ for all t . Finally from (8.6), then:

$$\mathcal{V}(\mathbf{x}_1(t), \mathbf{x}_2(t)) \leq \epsilon \implies \|\mathbf{y}_1(t) - \mathbf{y}_2(t)\| \leq \epsilon. \quad \square$$

The primary goal of employing RSF is to construct an abstract system Σ_2 which is ϵ -close to the concrete system Σ_1 , where ϵ remains small enough. Note that in the modified specification (8.5), any value ϵ such that $\hat{\mathcal{T}} = \emptyset$ causes the set of controllers enforcing the specification to also be empty. Therefore, the approximation approach must provide error thresholds small enough to give a feasible controller on the abstract system.

8.6 Proof of Concept

Consider Σ_1 as the interconnected NETS, Σ_1^i as the 9-state decomposed subsystem i and Σ_2^i as the 3-state reduced-order subsystem i . In all scenarios of this case study, consider a power loss disturbance of $\mathbf{v}_1^i = 1$ per unit (100 MW, equivalent to a typical generator or 35,000 households) as the default external disturbance. Construct abstract systems Σ_2^i using MATLAB's *balreal* function by truncating matrices to a reduced-state order of 3. YALMIP [116] and MOSEK [132] are employed for solving LMIs on a macOS machine with 8 GB RAM and Intel Core i5 Processor. The tool SCOTS [163] is also used for the synthesis of the symbolic controller using a high-performance computer with 2 nodes and 11 GB memory per core. Simulations are run over a time horizon of 6 seconds, with a time step of 0.005 seconds. The values of the interconnected NETS model as well as the different subsystems can be found in the Appendix.

Running case study (cont.): To demonstrate the proposed RSF with disturbance refinement, consider just one area of NETS, containing 9 states with one input, one external disturbance and no internal disturbance. The single line diagram for this system is depicted in Fig. 40. A linear model for Area 1 of NETS is acquired using the Simulink Model Linearizer on the closed-loop system.

To add nonlinear parts to this model, a collection of energy storage systems (ESSs) are considered which provide feedback control to the system depending on the current frequency. The power output of these ESSs is limited by a saturation function. It is assumed one has knowledge of dynamics of this feedback but no control over its power output. The dynamics of this output are adapted and simplified from aggregate battery charger models appeared in [81, 133]:

$$\Delta P_{ESS} = N_{ESS} \times \text{sat}\left(\frac{k_{ESS}}{R_{ESS}} \Delta f\right),$$

where P_{ESS} is the power contribution of ESSs, f is the system frequency, N_{ESS} is the number of participating ESSs, k_{ESS} is the average participation factor, and R_{ESS} is the droop constant. The saturation function $\text{sat}(x)$ is also defined as

$$\text{sat}(x) := \begin{cases} ESS_{max}, & x \geq ESS_{max}, \\ x, & ESS_{min} \leq x \leq ESS_{max}, \\ ESS_{min}, & x \leq ESS_{min}. \end{cases}$$

Assume that the (potentially) large disturbance \mathbf{v}_1 is measurable in the power system domain given that the disturbance may represent changes in the behaviour of generation and load components, e.g., generators, plug-in electric vehicles (EVs) and ESSs. The generation or load values of these components may be known to operators and the connection/disconnection of these components could be tracked through sensors in a smart grid. It is assumed one has access to a penetration of ESSs which can connect/disconnect from the power grid almost instantaneously. Such responsive loads are flexible and can be used for load shedding [208] and frequency regulation of smart grids [209]. The dynamics of the model are therefore a nonlinear system Σ_1^i equivalent to (8.8).

8.6.1 Simulation Relation Error

Uncontrolled system. If the response of EVs is not included in the system ($\mathbf{u}_1 = 0$), the open-loop nonlinear system Σ_1 has the maximum frequency deviation of $\Delta f = -0.6872\text{Hz}$, which clearly violates the specification ψ . Therefore, the contribution of EVs is essential to satisfy the specification on the frequency.

Abstraction with disturbance refinement. Using the proposed approaches from Theorems 8.1–8.2 with the proposed disturbance interface function. The safe and target sets are defined as $\underline{B} = -0.35$, $\overline{B} = 0.5$, $\underline{T} = -0.3$ and $\overline{T} = 0.5$. Assume, $\lambda = 1.7$, $\|\mathbf{u}_2\|_\infty = 0.5$, and $0.01\mathbb{I}_9 \leq \overline{M} \leq 120\mathbb{I}_9$, $L_{22} = 1$, $L_{21} = 0$, $D_2 = E_2$ and $Q_1 = K_1 = 0$. R_1 , R_2 and B_2 are optimised to minimise (8.13) and (8.14), respectively. Accordingly, the value $\epsilon = 0.1019$.

8.6.2 Controller Synthesis

Baseline controller. Consider the RSF with the constructed abstract system Σ_2 and the interface functions (8.12) but with $\mathbf{u}_2 = 0$ in (8.12a). As Q_2 and K_2 are non-zero in (8.12a), control inputs are chosen automatically based on the current states of Σ_1 and Σ_2 to maintain outputs of the two systems within distance ϵ . When the power system frequency moves away from its steady-state value, the input interface function u_V generates a control input for Σ_1 , which is considered here as a *baseline controller*. The frequency response in Σ_1 without EV participation (uncontrolled system with $\mathbf{u}_1 = 0$) against the baseline controller is depicted in Fig. 43. Although the baseline controller reduces the frequency deviations, it is still unable to satisfy the required specification ψ .

Controller using RSF. The constructed abstraction Σ_2 is employed as an appropriate substitute in the controller synthesis process. In particular, by knowing ϵ as the maximum error between outputs of Σ_1 and Σ_2 , a symbolic controller can be first designed for the reduced-order model Σ_2 to satisfy $\hat{\psi}$ and then be refined back to Σ_1 while providing a guarantee on the satisfaction of ψ . The synthesis of the symbolic controller takes 77 minutes and 34 seconds.

Remark 8.2. Note that synthesising such a symbolic controller directly from any 9-dimensional system is impossible due the required exponentially large computational time and memory space.

A comparison between the baseline controller and the synthesised one is provided in Fig. 44. The input \mathbf{u}_2 , synthesised by SCOTS, is chosen to be the minimum \mathbf{u}_2 that guarantees satisfaction of the specification ψ . Successful synthesis of the controller over Σ_2 formally shows that ψ also holds on Σ_1 . Fig. 44 (bottom) shows that over the time interval $t \in [0.5, 1]$, the synthesised controller over Σ_2 takes non-zero values to bring back the frequency to the intended target region, thus enabling Σ_1 to satisfy ψ .

As it can be observed, formal guarantees were provided using symbolic control over a 9-dimensional system while only requiring the computational load of a

3-dimensional system. To verify Theorem 8.2, quantify the maximum mismatch between output trajectories of Σ_1 and Σ_2 from simulations as

$$\max_t \|\mathbf{y}_1(t) - \mathbf{y}_2(t)\| = 0.0541.$$

Since this value is less than ϵ , the controller is demonstrably formally robust.

Remark 8.3. Note that the input profile is different to previous chapters due to the interface function attempting to keep the frequencies of the concrete system and the abstract system close.

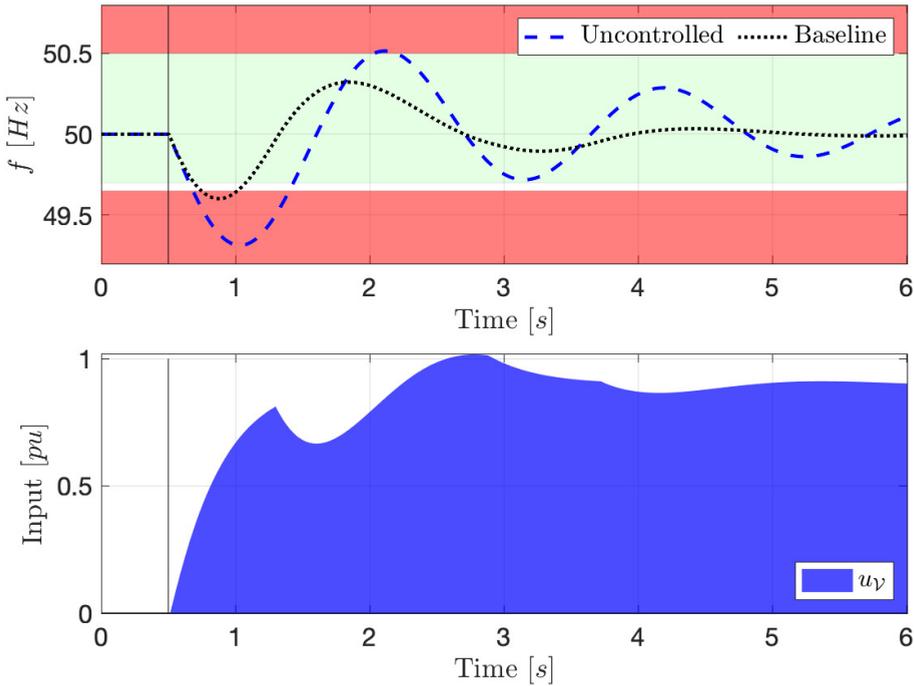


Figure 43: **Top.** Target range \mathcal{T} is shown in green, \mathcal{A}_{ub} and \mathcal{A}_{lb} are shown in red as two regions that the system should never transition into (unsafe regions). The baseline controller notably improves the frequency response of the system in compare with the uncontrolled system. However, both curves still fall into the red unsafe region. **Bottom.** The input u_γ is a byproduct of the simulation relation interface keeping Σ_1 and Σ_2 ϵ -close. Since no controller is synthesised over Σ_2 , then $\mathbf{u}_2 = 0$.

8.7 System and Specification Interconnection

In the previous section, it was seen, through a proof of concept, how to mitigate the *curse of dimensionality* using simulation functions with disturbance refinement

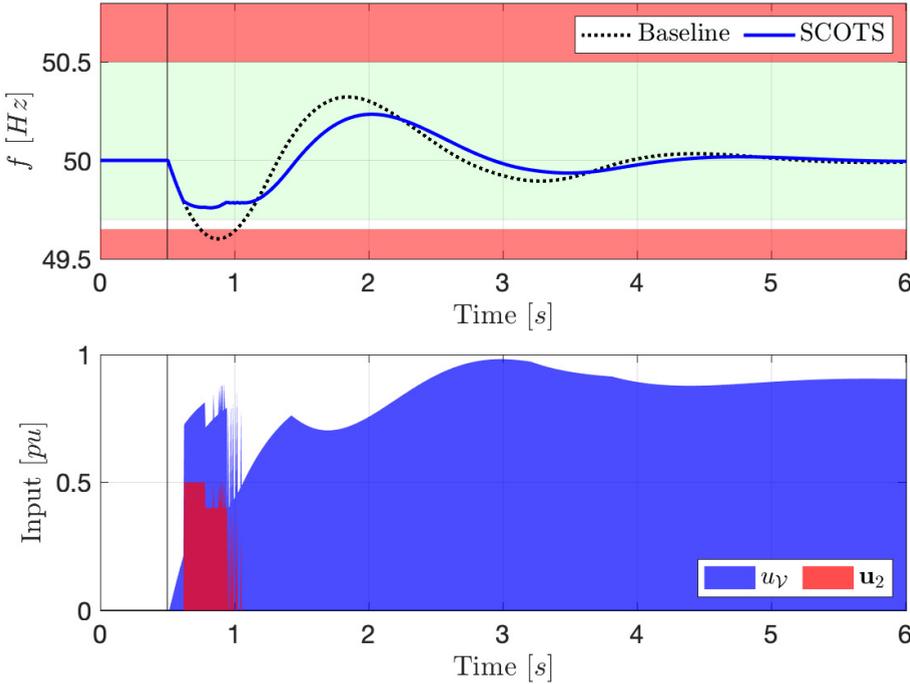


Figure 44: **Top.** Target range \mathcal{T} is shown in green, and unsafe regions \mathcal{A}_{ub} and \mathcal{A}_{lb} are shown in red. The controller designed using SCOTS and the RSF with disturbance refinement successfully satisfy ψ , compared with the baseline controller which violates the specification. **Bottom.** The control input u_2 designed using SCOTS for Σ_2 and the refined control input u_1 for Σ_1 using the RSF.

with a reduced-order model. The proposed approach considered a system model which was computationally intractable to synthesise and reduced it down to a system model with lower dimensions. The trade-off required a buffer ϵ , which is included in the synthesis to make sure the controller is robust to the loss of information due to the reduction. The more the system is reduced, the larger ϵ will be and the greater the challenge of synthesising a robust controller.

This approach is conservative when dealing with large-scale systems. To improve this technique for high-dimensional systems, the RSF with disturbance refinement can be combined with compositionality techniques from the literature, particularly assume-guarantee contracts. In particular, consider the large-scale system as an interconnected network composed of several smaller subsystems. Subsystems are worked on by constructing a reduced-order model abstraction for each subsystem. Under assume-guarantee contracts, the results from subsystems can be lifted to the interconnected system by providing formal guarantees on the satisfaction of the overall specification over the interconnected system.

Running case study (cont.): Here, consider NETS to be a composition of three areas, connected via the interface function on the frequency in (8.2). A graphical

representation of interconnections of NETS is provided in Fig. 45, in which each subsystem is labelled with its input, disturbance, and frequency.

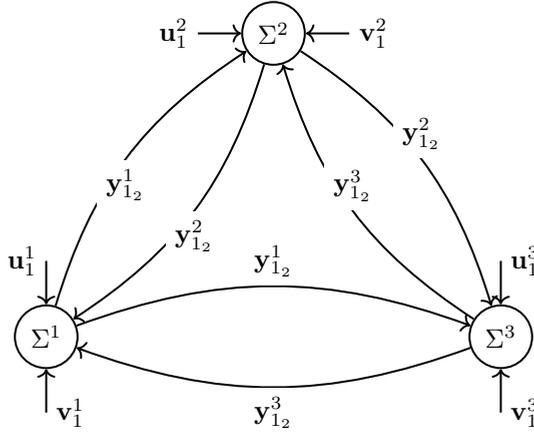


Figure 45: A graphical representation of NETS composed of 3 subsystems as vertices and interconnections with neighbours as edges. Including subsystems Σ^i , inputs u_1^i , external disturbances v_1^i , and internal disturbances derived as $w_1^i = y_{1_2}^i$.

The next subsection describes how local formal specifications of each subsystem can be combined to give a formal specification globally.

8.7.1 Specification Composition

The particular LTL properties useful for composition are conjunction (\wedge) and disjunction (\vee). Two separate LTL specifications can be combined to create either a stricter or looser specification. Therefore, in the case study, one can define local specifications on subsystems and combine them with conjunction to form a global LTL specification for the overall NETS.

Running case study (cont.): In NETS, the most important requirement is safety (or invariance), where globally the frequency must never fall beneath the containment zone. Additionally, the reachability specifications can be complementary as all subsystems are considered to have similar frequencies from the principle of interconnected systems. Therefore, all subsystems should simultaneously move toward their target area, if those targets are in a similar location.

The LTL specification ψ for the global system can be written using local specifications ψ^i for all subsystems i as

$$\psi^i = \square[f^i \geq \mathcal{Z}], \quad \forall i \in \{1, \dots, N\},$$

$$\psi = \bigwedge_{i=1}^N \psi^i.$$

Since subsystems use (8.2) to transfer power between the networks, this needs to be considered in local specifications. Concisely, when synthesising controllers to guarantee the specification for a local area, the worst-case scenarios of its neighbours' actions should be considered. This is done in two different ways in the next sections: the first is to consider that the frequency in all areas is always the same (*i.e.*, isolated subsystems) or to include the frequency of the other areas as an internal disturbance to the local area (*i.e.*, compositionality with internal disturbances).

Remark 8.4. *For emphasis, I remark again that the global specification for all the areas is a simple safety specification, but each local area has a more detailed reach-avoid specification.*

8.8 Assume Guarantee Contracts

Under assume-guarantee contracts, described in Chapter 4, subsystems can be controlled independently and combined to provide interconnected system guarantees. Satisfaction of the contract is acquired when individual subsystem contracts have a refinement relation. Controllers designed on these subsystems then provide a decentralised approach to acquiring guarantees.

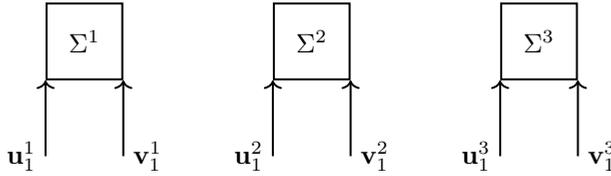


Figure 46: Design of NETS using isolated subsystems from the principle of interconnected synchronous machines.

8.9 Controllers for Subsystems

Running case study (cont.): Here, subsystems are isolated from each other using the assumption that the frequency of the NETS is the same for all subsystems, *i.e.*, $f^1 = f^2 = f^3$. This assumption follows the principle of interconnected synchronous machines. Under this assumption, the internal disturbance from neighbouring subsystems is removed, and accordingly, (8.2) equates to zero. An updated visual depiction of NETS is provided in Fig. 46.

Controllers are synthesised to satisfy the specification (or contract) of the interconnected system using contract composition of each subsystem. If the frequencies of each area remain close to one another and a controller is designed to satisfy the specification, then the contract is satisfied for the subsystem. When all subsystem contracts are satisfied and a refinement relation holds for the contract composition, then the specification/contract for the interconnected power network is satisfied.

Decentralised NETS Control

Using a similar technique to Section 8.6 with the linear case ($E_1^i = F_1^i = 0$ in (8.8)), each area's controller can be synthesised independently with no internal disturbances between neighbouring areas due to their isolation. By combining the guarantees that the controllers provide on each area, the global specification can be ensured on NETS. Under the assumption that the area frequencies remain close to one another, each subsystem can be disturbed independently by some v_1^i , and no area should violate its specification.

The global specification is defined as

$$\psi = \psi^1 \wedge \psi^2 \wedge \psi^3,$$

where ψ^i is the specification of subsystem i . Isolating the areas also uncouples the reachability specifications providing a higher likelihood of formal guarantees over the interconnected system. For the safety guarantees, if a system can guarantee Δf^i never falls to -0.35 Hz, then it implicitly guarantees Δf^i never falls to -0.6 Hz. So the overall guarantee provided for all areas would match the weakest guarantee of a single area (*i.e.*, worst-case scenario). In Fig. 47, subsystems are disturbed by each $v_1^i = 1$ per unit ($u_2^i = 0$) and the frequency of multiple areas violates the safety specification.

By deploying the formal synthesised controllers where $u_2^i = 0.5$, the assume-guarantee contract approach shows all subsystems satisfy both the safety guarantees of ψ and also the reachability guarantees. The frequency of each area also remains close to the neighbouring regions.

Single Area Control

Consider Area 3 of NETS. When designing the controller, consider a large measurable disturbance v_1^3 of 1 per unit. In addition, $v_1^1 = v_1^2 = 0$, $\underline{B} = -0.6$, $\underline{T} = -0.35$, $\overline{B} = \overline{T} = \infty$, $u_2^3 = 0.5$, and ϵ is calculated as 0.1016. The synthesised controller is depicted in Fig. 49. It is worth remarking that for Area 3, it was not possible to find as tight of a reach-avoid bound as in Σ^1 (Fig. 51) and Σ^2 (Fig. 52), where $\underline{B} = -0.35$, $\underline{T} = -0.3$, $\overline{B} = \overline{T} = \infty$.

8.10 Compositionality with Internal Disturbances

Now, compositional techniques are employed exclusively to capture internal disturbances from neighbouring subsystems. Under assume-guarantee contracts, the aim is to strengthen the guarantees on the behaviour of the system using shared information.

Running case study (cont.): For NETS, additional knowledge of the frequency of neighbouring regions is considered which impacts the frequency of the local

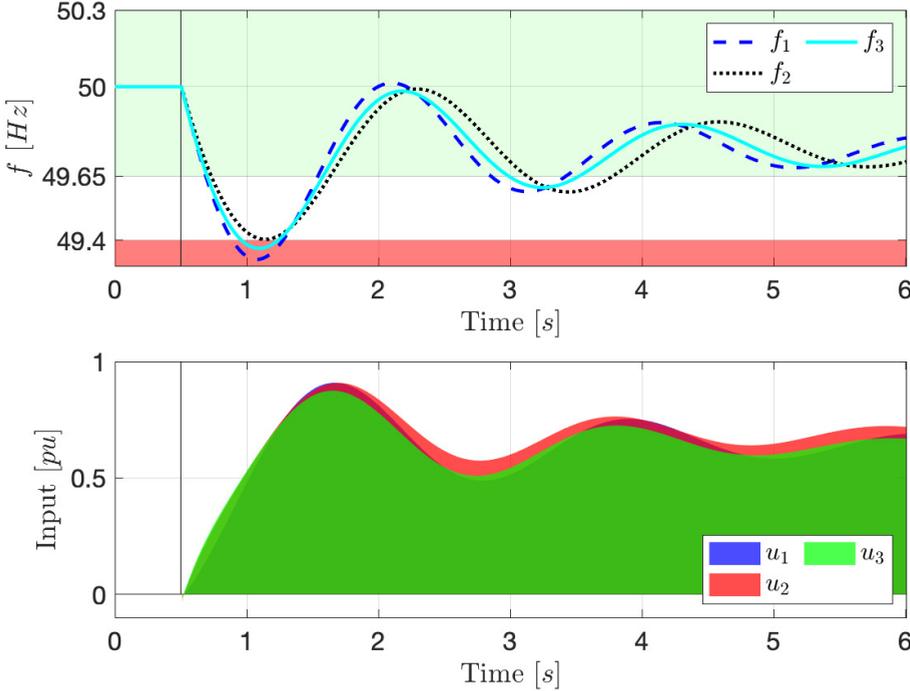


Figure 47: *Interconnected system without control.* **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. For the baseline controller which keeps Σ_1 and Σ_2 ϵ -close, it can be seen that f^1 falls into the red unsafe region. **Bottom.** The inputs are a byproduct of the simulation relation interface keeping Σ_1 and Σ_2 ϵ -close. No additional controller is synthesised over Σ_2 .

subsystem. Consider the frequency of neighbouring subsystems as internal disturbances defined in (8.2).

When synthesising a controller, it is important to include the neighbouring frequency information in the synthesis procedure. Here, a reach-avoid specification is considered where each subsystem should avoid a region of the state space while trying to return a safe region after being disturbed. The disturbance of a neighbouring area should never cause the local area to violate the specification. Therefore, in the control synthesis problem, the controller should be robust to the worst-case neighbours' frequencies. Given the reach-avoid specification for each area, the boundary of the avoid region can be used to define the worst-case disturbance acting on a local subsystem from its neighbour. Fig. 50 shows how this looks for Σ^3 with input \mathbf{u}_1^3 , external disturbance \mathbf{v}_1^3 , and internal disturbances $\mathbf{w}_1^1 = \mathbf{y}_{1_2}^1$ and $\mathbf{w}_1^1 = \mathbf{y}_{1_2}^2$.

Using the compositional approach, subsystems have two internal disturbances coming from the neighbouring areas. For the RSF with disturbance refinement, these two disturbances have a significant impact on the value of ϵ . For Σ^1 , the isolated systems approach has $\epsilon = 0.1016$ while for compositionality $\epsilon = 0.1896$.

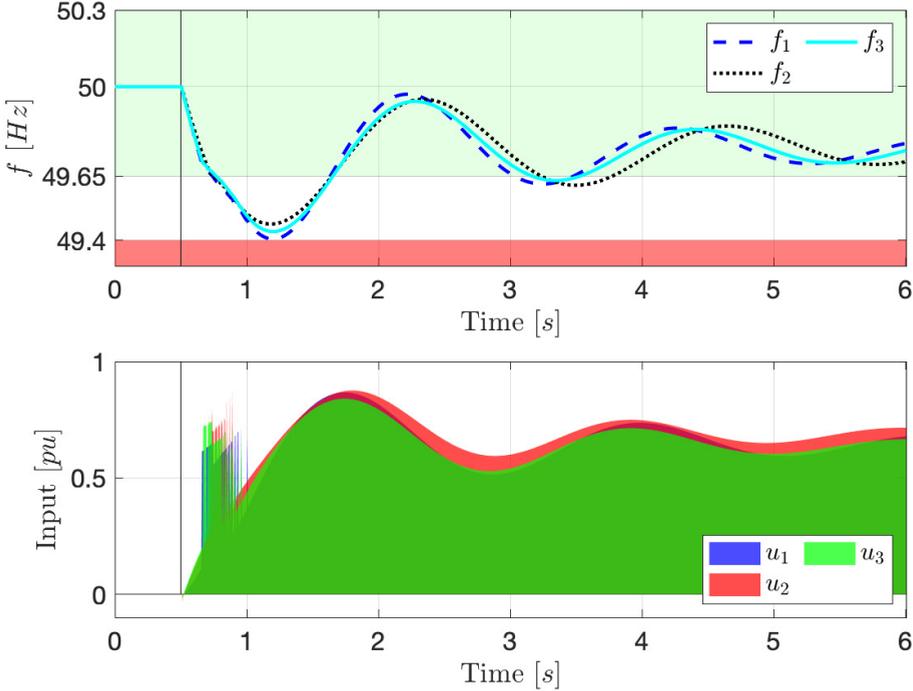


Figure 48: *Interconnected system with formal control.* **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. The controller, designed using SCOTS, and the RSF with disturbance refinement successfully satisfy ψ . **Bottom.** The synthesised control input for Σ_2 and the refined control input for Σ_1 using RSF are combined to provide inputs which guarantees the satisfaction of specification.

For the compositional approach, one defines a global specification that when a disturbance $\mathbf{v}_1^3 \leq 1$ is present in Σ^3 , no area's frequency should fall below $\Delta f^i \leq -0.6$ Hz. The simulation relation error is computed as $\epsilon^3 = 0.1992$, where $\mathbf{u}_2^3 = 1$. If the response of EVs is not included in the system ($\mathbf{u}_2^3 = 0$), the maximum frequency deviation violates the specification, as can be seen in Fig. 53. Therefore, the contribution of EVs is essential to satisfy the specification on the frequency, as shown in Fig. 54.

8.11 Conclusion

This chapter studied two compositional control approaches for large-scale power systems while providing guarantees over the system's behaviour. Including:

- Providing a temporal logic specification for frequency regulation in the GB Power Network;

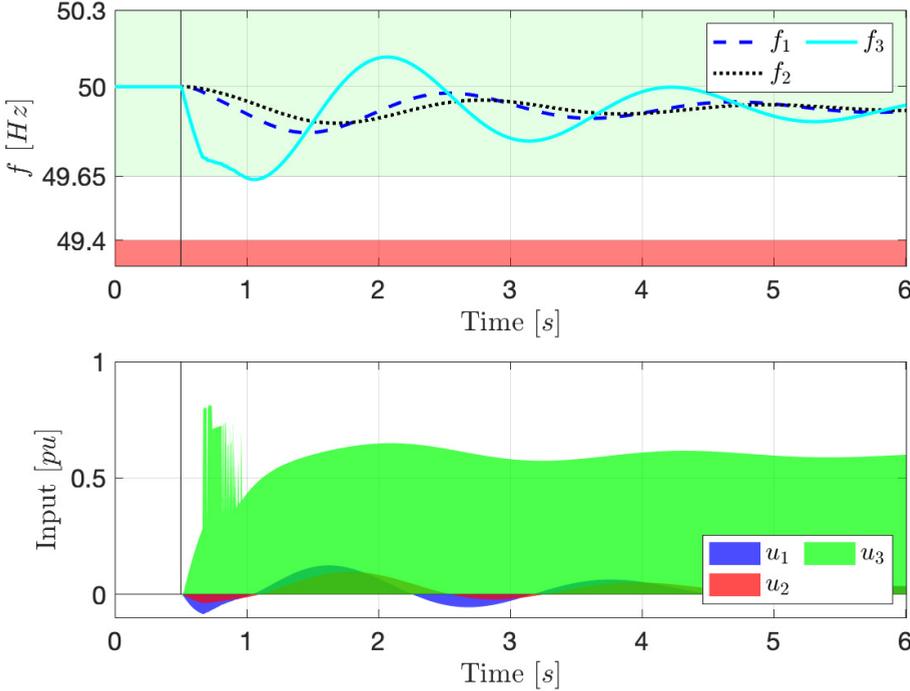


Figure 49: *Area 3 with formal control.* **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. The controller, designed using SCOTS, and the RSF with disturbance refinement successfully satisfy ψ . **Bottom.** The synthesised control input for Σ_2 and the refined control input for Σ_1 using RSF are combined to provide u_1^3 which guarantees the satisfaction of specification.

- Employing assume-guarantee contracts using robust simulation functions (RSF) with disturbance refinement to design decentralised controllers for distinct power system areas while providing guarantees for the fully interconnected power network;
- Demonstrating the compositional approach in two ways: i) using the principle of interconnected synchronous machines, ii) considering bounds on the power flows between neighbouring subsystems, relative to the area's frequency;
- Extending the notion of RSFs with disturbance refinement to a class of non-linear systems;
- Using the New England 39-Bus Test System as a challenging running case study to demonstrate the proposed approach.

The following chapter will conclude the thesis and provide recommended areas for future research.

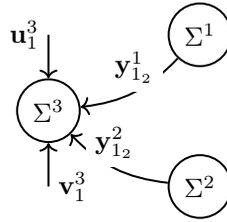


Figure 50: The NETS subsystem Σ^3 displayed as a vertex of a graph with subsystems as other vertices and interconnections with neighbours as edges. u_1^3 is an input, v_1^3 is an external disturbance, while $w_1^1 = y_{1_2}^1$ and $w_1^2 = y_{1_2}^2$ are internal disturbances for Σ^3 .

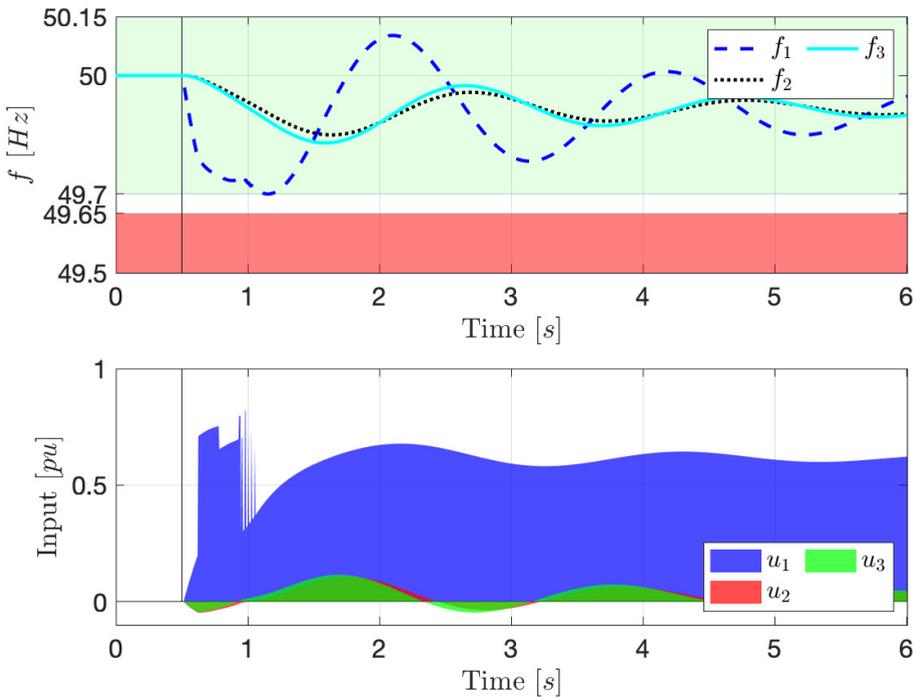


Figure 51: *Area 1 with formal control.* **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. The synthesised controller and the robust simulation function with disturbance refinement successfully satisfy ψ^1 . **Bottom.** The synthesised control input for Σ_2 and the refined control input for Σ_1 using RSF are combined to provide u^3 which guarantees the satisfaction of specification.

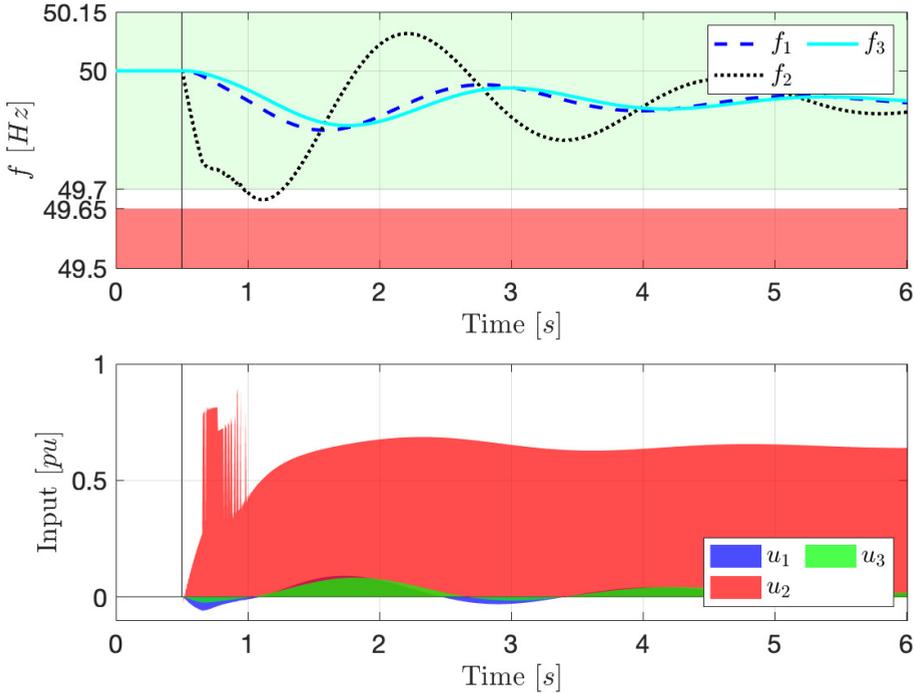


Figure 52: Area 2 with formal control. **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. The synthesised controller and the robust simulation function with disturbance refinement successfully satisfy ψ^2 . **Bottom.** The synthesised control input for Σ_2 and the refined control input for Σ_1 using RSF are combined to provide \mathbf{u}^2 which guarantees the satisfaction of specification.

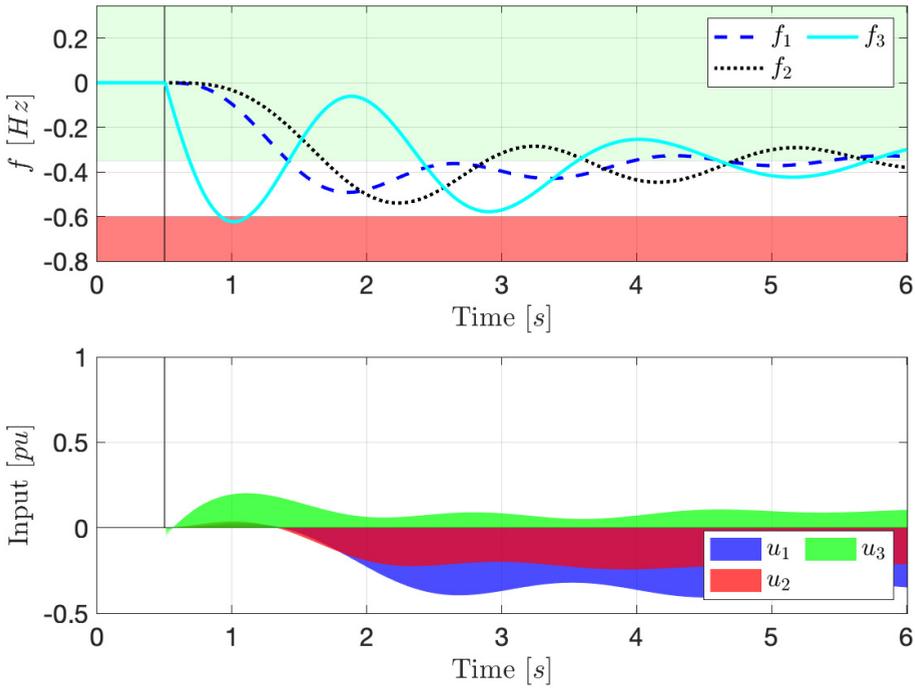


Figure 53: *Interconnected system without control.* **Top.** Target region \mathcal{T} and unsafe region \mathcal{A} are shown in green and red, respectively. For the baseline controller which keeps Σ_1 and Σ_2 ϵ -close, it can be seen that f^3 falls into the red unsafe region. **Bottom.** The inputs are a byproduct of the simulation relation interface keeping Σ_1 and Σ_2 ϵ -close. No additional controller is synthesised over Σ_2 .

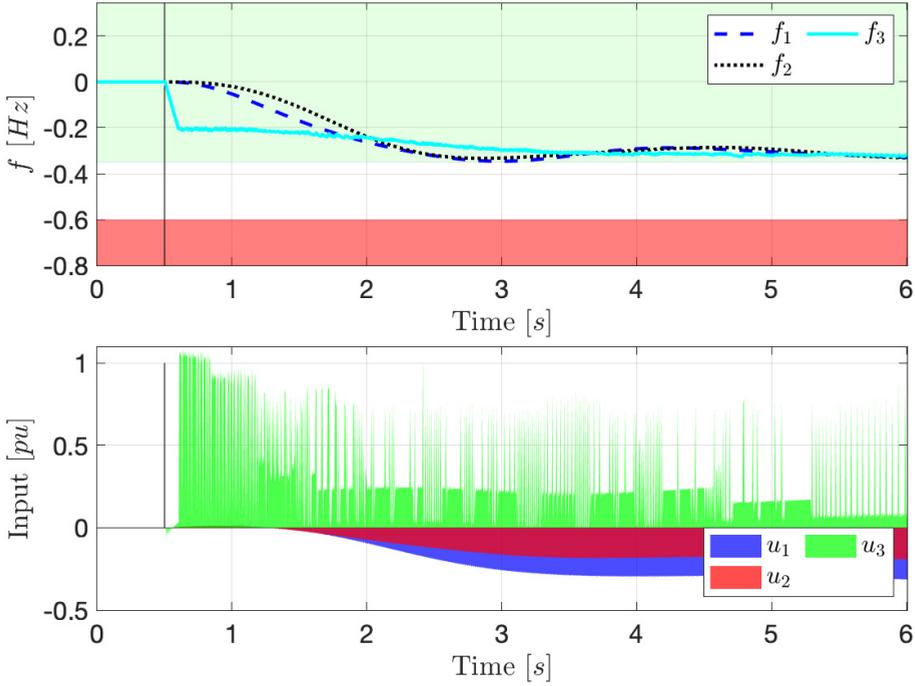


Figure 54: *Interconnected system with formal control.* **Top.** Target range \mathcal{T} is shown in green, unsafe region \mathcal{A} is shown in red. The controller designed using SCOTS and the RSF with disturbance refinement successfully satisfy ψ . **Bottom.** The control input designed using SCOTS for Σ_2 and the refined control input for Σ_1 using the RSF are combined to give the inputs which guarantees the systems meets the specification.

Conclusion and Future Directions

This chapter summarises the thesis and presents the main contributions of this manuscript. Some directions of future research are provided which the reader is recommended to pursue.

9.1 Conclusions

This thesis discussed symbolic control methods for smart grids with both model-based and data-driven methods. Example case studies were provided using *active buildings*, *electric vehicles* (EVs) and *energy storage systems* (ESSs) in demand-side primary frequency response control techniques. In the following, I summarise the main contributions.

- **Conversion of Great Britain power network specification from natural language to temporal logic specification.** A detailed description of the requirements on the behaviour of the Great Britain power network from the literature have been provided and encoded in *linear temporal logic* (LTL) specifications which can be verified using symbolic control methods, see Chapter 8.
- **Designing formal controllers for power systems with large disturbances.** Mathematical guarantees are provided for the correctness of controllers designed over power systems with bounded large disturbances. Giving confidence the system under the controller will always return to a target region and never fall into an unsafe frequency range that may lead to contingency events such as blackouts, see Chapter 5.
- **Formal control of primary frequency response using distributed energy resources.** It was demonstrated how formal control techniques can be applied to demand-side response in the emerging smart grids. It was demonstrated how *distributed energy resources* (DERs) can be used to provide a fast

control response to aid power system stability, particularly for primary frequency control, see Chapter 5.

- **Abstraction-based controller design for unknown systems with finite numbers of data samples.** A data-driven method to compute symbolic models (*a.k.a* finite abstractions) of systems with unknown dynamics has been presented. Robust convex programs were used to overapproximate reachable sets and solve a scenario convex program to find a feasible solution with given confidence. A lower bound on the number of trajectories required to achieve a certain confidence was also provided, see Chapter 6.
- **Designing symbolic controllers for reduced-order models.** Model-order reduction techniques were applied to power systems and then a relation was proven between the (original) concrete system and its (reduced-order) abstract system. By refining an approximate simulation function with an interface function for the disturbance the simulation relation error was reduced between these two systems. The simulation relation error was then used to reduce the regions defined as safe states or target states, and to increase the regions defined as unsafe states; making the symbolic controller robust to the model-reduction step of the controller design. Theoretical results for the approach were provided for both the class of linear systems and a class of nonlinear systems, see Chapter 7.
- **Formal control of interconnected power systems.** A formal design approach was provided for symbolic controllers of interconnected power systems using assume-guarantee contracts. In combination with the prior mentioned reduced-order model control approach, this enabled complex large interconnected power system models to be simplified to reduced-order subsystems that can be controlled independently. Assumptions of successful control of individual areas produced confidence that the interconnected system would behave as expected and satisfy the requirements, see Chapter 8.
- **Challenging power system case studies.** Throughout this thesis, challenging non-trivial power system case studies were presented that demonstrate the potential of the formal methods approaches. In particular, case studies of the Great Britain (GB) power network and the *New England 39-Bus Test System*(NETS) are used for the design of formal controllers.

9.2 Future Research Directions

In this section, I discuss some interesting topics and ideas for future investigation.

- **Fully distributed formal control of power systems.** The smart grid is evolving to be more distributed and electrified as more buildings become active buildings, more EVs are owned, and more renewable generating devices are installed. The increased load and uncertainty of renewable generation will add strain to the power network unless control schemes are developed to

correctly manage these devices while also factoring in user experience. A scalable strategy would be fully distributed multi-agent control schemes to anticipate this evolution in power networks. This future work comes from Chapter 8 and would be to the scale of a research project.

- **Extensions and applications to nonlinear systems.** As smart grids become more electrified with more network components they will become more complex to control. These complexities and nonlinearities need to be modelled and controlled. This future work comes from Chapter 8 and would be to the scale of a PhD project.
- **Evolving from model-based to data-driven control.** In this thesis, I showed a data-driven method to control a 3 area 3 machine power system. As the smart grid becomes more complex, developing accurate models of these systems will be more challenging and so data-driven control without underlying models will become necessary alternatives to rely on. This future work comes from Chapter 8 and would be to the scale of a PhD thesis or even a research project.
- **Considering noise and unknown disturbances.** For this thesis, I considered all case studies to be non-probabilistic, however in reality this is not often the case, e.g. due to uncertainties related to measurements, errors in the model, or external unpredictable factors (such as weather). Methods which consider these uncertainties in their controller design will be of great benefit. This future work comes from Chapter 5 and would be to the scale of a Master's project or PhD project.
- **Voltage control.** The control approaches of this thesis have focused particularly on primary frequency control. Voltage control is also important to smart grids; so formal methods could be developed for the transmission network, distribution network and the coordination between the two. This future work comes from Chapter 3 and would be to the scale of a research project.
- **Modelling smart grid vector components in the smart grid.** I discussed how controlling smart grids is a multi-vector approach, consisting of different smart grid component vectors i.e. water, heat, electricity, etc. Understanding and modelling these components accurately will enhance the control approaches that can be developed using them, e.g. EV charging strategies and availability. This future work comes from Chapter 3 and would be to the scale of a PhD project or research project.
- **Inertia considerations.** As more renewable generation is added to the power networks the less turbines will be required to provide energy. However, turbines provide inertia into power systems which greatly improve the stability of the power network. As the smart grid evolves, virtual synchronous generators with guarantees will need to be developed to solve this problem. This future work comes from Chapter 3 and would be to the scale of a PhD thesis or research project.

- **On-the-fly techniques for real-time power systems.** Due to the curse of dimensionality both computation and memory costs are high when using formal control methods. This means most controllers are designed offline through the use of look-up tables. Improvements in these algorithms so they can run on-the-fly and be updated should changes in the environment occur, would greatly improve the usability of such models. They would also prevent the designed controller from becoming obsolete due to changes in the environment invalidating a priori guarantees on the controller. Some work in this area has already begun [123]. This future work comes from Chapter 6 and would be to the scale of a challenging research project. Current technology and methods are not to the level this is feasible at the moment.
- **Parallelised data-driven tool implementation.** Model-based symbolic approaches have a large computational overhead. This is even more true for data-driven symbolic control as each state has a large number of samples taken from that state. Developing data-driven tools that run in parallel would dramatically speed up this process. This future work comes from Chapter 6 and would be to the scale of a large Master's project or a PhD project.
- **Extending robust simulation functions with disturbance refinement to a general system.** In this thesis I demonstrated the approach for linear systems and a class of nonlinear systems but it would be of interest to extend this idea to a general nonlinear system. This way, the simulation relation error can always be included inside the abstract system design from the model-order reduction step, to synthesise formal controllers for the original systems. This future work comes from Chapter 7 and would be to the scale of a PhD project.
- **ϵ -approximate digital twins.** Another interesting direction regards the robust simulation function approach as a way of forming digital twins. If the abstract model can be learned which is an ϵ -approximately correct model of a real-life system the a controller can be designed on this digital twin which is robust enough to be used on the real-life system and to provide behaviour guarantees. This future work comes from Chapter 7 and would be to the scale of a PhD thesis or research project.
- **Optimal interfaces for robust simulation functions.** I discussed how using two interface functions can reduce the simulation relation error in the model-reduction step when finding robust simulation functions. It would be of interest to find the optimal parameters of the interfaces to reduce the simulation relation error and provide guarantees on the upper bound of the input required to satisfy the specification. This future work comes from Chapter 7 and would be to the scale of a good Master's student project or a PhD project.

A

APPENDIX

Appendix

A.1 NETS Matrices from Chapter 7

The matrices of the NETS single area Σ_1 are given as:

$$A_1 = \begin{bmatrix} -12.5 & 0 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & -16.67 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & 0 & -14.29 & 0.05 & -0.61 & 0 & 0 & 0 & -0.05 \\ 0 & 0 & 0 & 0 & 0.93 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -6.28 & -0.09 & 2.5 & 2.78 & 2.38 & 0 \\ 12.5 & 0 & 0 & 0 & 0 & -2.5 & 0 & 0 & 0 \\ 0 & 16.67 & 0 & 0 & 0 & 0 & -2.78 & 0 & 0 \\ 0 & 0 & 14.29 & 0 & 0 & 0 & 0 & -2.38 & 0 \\ 0 & 0 & 0 & 6.28 & 2.08 & 0 & 0 & 0 & 0 \end{bmatrix}$$
$$B_1 = [0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0]^T$$
$$D_1 = [0 \ 0 \ 0 \ 0 \ -1 \ 0 \ 0 \ 0 \ 0]^T$$
$$C_1 = [0 \ 0 \ 0 \ 0 \ 2.05 \ 0 \ 0 \ 0 \ 0]$$

The reduced-order model Σ_2 is constructed as:

$$A_2 = \begin{bmatrix} -0.6333 & 3.0028 & 0.4428 \\ -3.0028 & -0.0026 & -0.0263 \\ -0.4428 & -0.0263 & -1.5159 \end{bmatrix}$$
$$B_2 = [-0.8580 \ 0.5378 \ 0.6956]^T$$
$$D_2 = [0.8580 \ -0.5378 \ -0.6956]^T$$
$$C_2 = [-1.7990 \ 0.1141 \ 0.5998]$$

Note that we have $D_2 = 0$ for the method without the disturbance refinement. The

matrices obtained for establishing our robust simulation relation are as follows:

$$M = \begin{bmatrix} 0.22 & 0 & 0 & 0 & 0.01 & -0.01 & 0 & 0 & 0 \\ 0 & 0.26 & 0 & 0 & 0.01 & 0 & -0.01 & 0 & 0 \\ 0 & 0 & 0.26 & 0 & 0.01 & 0 & 0 & -0.01 & 0 \\ 0 & 0 & 0 & 82.14 & 20.22 & 0 & 0 & 0 & 16.80 \\ 0.01 & 0.01 & 0.01 & 20.22 & 11.62 & 0 & 0 & 0 & 11.68 \\ -0.01 & 0 & 0 & 0.01 & 0 & 0.02 & 0 & 0 & 0 \\ 0 & -0.01 & 0 & 0.01 & 0 & 0 & 0.02 & 0 & 0 \\ 0 & 0 & -0.01 & 0.01 & 0 & 0 & 0 & 0.02 & 0 \\ 0 & 0 & 0 & 16.80 & 11.68 & 0 & 0 & 0 & 29.44 \end{bmatrix}$$

$$P = \begin{bmatrix} 0.04 & 0.03 & 0.03 & 0.02 & -0.88 & 0.025 & 0.044 & 0.03 & 0.66 \\ -0.01 & -0.01 & -0.01 & 0.29 & 0.06 & -0.10 & -0.98 & -0.99 & 0.36 \\ -0.03 & -0.18 & -0.018 & -0.18 & 0.29 & -0.33 & -0.25 & -0.31 & 0.52 \end{bmatrix}^T$$

$$Q_1 = K_1 = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$$

$$K_2 = [-0.2 \ -0.2 \ -0.2 \ -482.5 \ -278.9 \ -2.5 \ -2.8 \ -2.4 \ -279.9]$$

$$Q_2 = [0.0238 \ -0.0407 \ 0.3401]$$

$$R_1 = R_2 = 1.$$

A.2 NETS Matrices from Chapter 8

Proof of Concept - Nonlinear Area 1

$$A_1 = \begin{bmatrix} -12.5 & 0 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & -16.67 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & 0 & -14.29 & 0.05 & -0.61 & 0 & 0 & 0 & -0.05 \\ 0 & 0 & 0 & 0 & 0.93 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -6.28 & -0.09 & 2.5 & 2.78 & 2.38 & 0 \\ 12.5 & 0 & 0 & 0 & 0 & -2.5 & 0 & 0 & 0 \\ 0 & 16.67 & 0 & 0 & 0 & 0 & -2.78 & 0 & 0 \\ 0 & 0 & 14.29 & 0 & 0 & 0 & 0 & -2.38 & 0 \\ 0 & 0 & 0 & 6.28 & 2.08 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$B_1 = [0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0]^T$$

$$D_1 = [0 \ 0 \ 0 \ 0 \ -1 \ 0 \ 0 \ 0 \ 0]^T$$

$$C_1 = [0 \ 0 \ 0 \ 0 \ 2.05 \ 0 \ 0 \ 0 \ 0]$$

$$E_1 = [0 \ 0 \ 0 \ 0 \ 0.0285 \ 0 \ 0 \ 0 \ 0]^T$$

$$F_1 = [0 \ 0 \ 0 \ 0 \ 2 \ 0 \ 0 \ 0 \ 0]$$

$$A_2 = \begin{bmatrix} -0.6333 & 3.0028 & 0.4428 \\ -3.0028 & -0.0026 & -0.0263 \\ -0.4428 & -0.0263 & -1.5159 \end{bmatrix}$$

$$B_2 = [-1.0204 \ 0.6395 \ 0.8273]^T$$

$$D_2 = [1.0204 \ -0.6395 \ -0.8273]^T$$

$$C_2 = [-1.5128 \ 0.096 \ 0.5044]$$

Linear Area 1 - Isolated

$$A_1 = \begin{bmatrix} -12.5 & 0 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & -16.67 & 0 & 0.09 & -0.65 & 0 & 0 & 0 & -0.09 \\ 0 & 0 & -14.29 & 0.05 & -0.61 & 0 & 0 & 0 & -0.05 \\ 0 & 0 & 0 & 0 & 0.93 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -6.28 & -0.09 & 2.5 & 2.78 & 2.38 & 0 \\ 12.5 & 0 & 0 & 0 & 0 & -2.5 & 0 & 0 & 0 \\ 0 & 16.67 & 0 & 0 & 0 & 0 & -2.78 & 0 & 0 \\ 0 & 0 & 14.29 & 0 & 0 & 0 & 0 & -2.38 & 0 \\ 0 & 0 & 0 & 6.28 & 2.08 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$B_1 = [0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0]^T$$

$$D_1 = [0 \ 0 \ 0 \ 0 \ -1 \ 0 \ 0 \ 0 \ 0]^T$$

$$C_1 = [0 \ 0 \ 0 \ 0 \ 2.05 \ 0 \ 0 \ 0 \ 0]$$

$$A_2 = \begin{bmatrix} -0.6333 & 3.0028 & 0.4428 \\ -3.0028 & -0.0026 & -0.0263 \\ -0.4428 & -0.0263 & -1.5159 \end{bmatrix}$$

$$B_2 = [-0.8580 \ 0.5378 \ 0.6956]^T$$

$$D_2 = [0.8580 \ -0.5378 \ -0.6956]^T$$

$$C_2 = [-1.7990 \ 0.1141 \ 0.5998]$$

$$M = \begin{bmatrix} 0.22 & 0 & 0 & 0 & 0.01 & -0.01 & 0 & 0 & 0 \\ 0 & 0.26 & 0 & 0 & 0.01 & 0 & -0.01 & 0 & 0 \\ 0 & 0 & 0.26 & 0 & 0.01 & 0 & 0 & -0.01 & 0 \\ 0 & 0 & 0 & 82.14 & 20.22 & 0 & 0 & 0 & 16.80 \\ 0.01 & 0.01 & 0.01 & 20.22 & 11.62 & 0 & 0 & 0 & 11.68 \\ -0.01 & 0 & 0 & 0.01 & 0 & 0.02 & 0 & 0 & 0 \\ 0 & -0.01 & 0 & 0.01 & 0 & 0 & 0.02 & 0 & 0 \\ 0 & 0 & -0.01 & 0.01 & 0 & 0 & 0 & 0.02 & 0 \\ 0 & 0 & 0 & 16.80 & 11.68 & 0 & 0 & 0 & 29.44 \end{bmatrix}$$

$$P = \begin{bmatrix} 0.04 & 0.03 & 0.03 & 0.02 & -0.88 & 0.025 & 0.044 & 0.03 & 0.66 \\ -0.01 & -0.01 & -0.01 & 0.29 & 0.06 & -0.10 & -0.98 & -0.99 & 0.36 \\ -0.03 & -0.18 & -0.018 & -0.18 & 0.29 & -0.33 & -0.25 & -0.31 & 0.52 \end{bmatrix}^T$$

$$Q_1 = K_1 = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$$

$$K_2 = [-0.2 \ -0.2 \ -0.2 \ -482.5 \ -278.9 \ -2.5 \ -2.8 \ -2.4 \ -279.9]$$

$$Q_2 = [0.0238 \ -0.0407 \ 0.3401]$$

$$R_1 = R_2 = 1.$$

Area 1 is formed from rows and columns 1 – 8 and 25 from A in the fully interconnected system.

Linear Area 1 - with Internal Disturbance

The internal disturbances can be derived from column 13 and column 21 of A returning the following disturbance matrix.

$$D_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.37 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.52 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}^T$$

Linear Area 2 and Area 3

Area 2 and Area 3 can be derived in the same way that Area 1 was derived using the fully interconnected New England 39-Bus System. The matrices for these areas are not provided.

Bibliography

- [1] Alessandro Abate, Henk Blom, Nathalie Cauchi, Joanna Delicarlis, Sofie Haesaert, Birgit van Huijgevoort, Abolfazl Lavaei, Anne Remke, Oliver Schön, Stefan Schupp, Fedor Shmarov, Sadegh Soudjani, and Lisa and Paolo Zuliani Willemsen. Arch-comp23 category report: Stochastic models. In *10th International Workshop on Applied Verification of Continuous and Hybrid Systems, ARCH 2023*, pages 126–150. EasyChair, 2023.
- [2] Alessandro Abate, Henk Blom, Nathalie Cauchi, Joanna Delicarlis, Arnd Hartmanns, Mahmoud Khaled, Abolfazl Lavaei, Carina Pilch, Anne Remke, Stefan Schupp, Fedor Shmarov, Sadegh Soudjani, Abraham Vinod, Ben Wooding, Majid Zamani, and Paolo Zuliani. ARCH-COMP20 Category Report: Stochastic Models. In *ARCH20. 7th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH20)*. EasyChair, 2020.
- [3] Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.
- [4] Y. P. Agalgaonkar, B. C. Pal, and R. A. Jabr. Distribution voltage control considering the impact of pv generation on tap changers and autonomous regulators. *IEEE Transactions on Power Systems*, 29(1):182–192, 2014.
- [5] Ahmad Ahmad, Calin Belta, and Roberto Tron. Adaptive sampling-based motion planning with control barrier functions. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 4513–4518. IEEE, 2022.
- [6] Waseem Akram and Muaz A. Niazi. A formal specification framework for smart grid components. *Complex Adaptive Systems Modeling*, 6(1):5, Sep 2018.
- [7] Shravan Kumar Akula and Hossein Salehfar. Frequency control in microgrid communities using neural networks. In *2019 North American Power Symposium (NAPS)*, pages 1–6. IEEE, 2019.
- [8] Hassan S Haes Alhelou, MEH Golshan, and Masoud Hajiakbari Fini. Multi agent electric vehicle control based primary frequency support for future smart micro-grid. In *2015 Smart Grid Conference (SGC)*, pages 22–27. IEEE, 2015.

- [9] Matthias Althoff. Formal and compositional analysis of power systems using reachable sets. *IEEE Transactions on Power Systems*, 29(5):2270–2280, 2014.
- [10] Matthias Althoff and Bruce H. Krogh. Reachability analysis of nonlinear differential-algebraic systems. *IEEE Transactions on Automatic Control*, 59(2):371–383, 2014.
- [11] R. Alur, T. A. Henzinger, G. Lafferriere, and G. J. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(7):971–984, 2000.
- [12] Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. *2019 18th European Control Conference (ECC)*, Jun 2019.
- [13] Amir Anees and Yi Ping Phoebe Chen. True real time pricing and combined power scheduling of electric appliances in residential energy management system. *Applied Energy*, 165:592–600, 2016.
- [14] Anuradha M Annaswamy, Ahmad R Malekpour, and Stefanos Baros. Emerging research topics in control for smart infrastructures. *Annual Reviews in Control*, 2016.
- [15] Luca Arnaboldi, Ricardo M Czekster, Charles Morisset, and Roberto Metere. Modelling Load-Changing Attacks in Cyber-Physical Systems. *Electronic Notes in Theoretical Computer Science*, 2020.
- [16] Pouya Babahajiani, Qobad Shafiee, and Hassan Bevrani. Intelligent demand response contribution in frequency control of multi-area power systems. *IEEE Transactions on Smart Grid*, 9(2):1282–1291, 2018.
- [17] Thom Badings, Licio Romao, Alessandro Abate, David Parker, Hasan A Poonawala, Marielle Stoelinga, and Nils Jansen. Robust control for dynamical systems with non-gaussian noise via formal abstractions. *Journal of Artificial Intelligence Research*, 76:341–391, 2023.
- [18] Thom S. Badings, Alessandro Abate, Nils Jansen, David Parker, Hasan A. Poonawala, and Marielle Stoelinga. Sampling-based robust control of autonomous systems with non-Gaussian noise. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9):9669–9678, June 2022.
- [19] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- [20] Andrea Bajcsy, Somil Bansal, Eli Bronstein, Varun Tolani, and Claire J Tomlin. An efficient reachability-based framework for provably safe autonomous navigation in unknown environments. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1758–1765. IEEE, 2019.
- [21] Tamajit Banerjee, Rupak Majumdar, Kaushik Mallik, Anne-Kathrin Schmuck, and Sadegh Soudjani. A direct symbolic algorithm for solving stochastic rabin games. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 81–98. Springer, 2022.

- [22] Tamajit Banerjee, Rupak Majumdar, Kaushik Mallik, Anne-Kathrin Schmuck, and Sadegh Soudjani. Fast symbolic algorithms for omega-regular games under strong transition fairness. *TheoretCS*, 2, 2023.
- [23] K. Bao, S. Li, and H. Zheng. Battery charge and discharge control for energy management in EV and utility integration. In *2012 IEEE Power and Energy Society General Meeting*, pages 1–8, July 2012.
- [24] Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. SIAM, 1998.
- [25] Calin Belta, Boyan Yordanov, and Ebru Aydin Gol. *Formal methods for discrete-time dynamical systems*, volume 15. Springer, 2017.
- [26] Albert Benveniste, Benoît Caillaud, Dejan Nickovic, Roberto Passerone, Jean-Baptiste Raclet, Philipp Reinkemeier, Alberto Sangiovanni-Vincentelli, Werner Damm, Thomas A Henzinger, Kim G Larsen, et al. Contracts for system design. *Foundations and Trends® in Electronic Design Automation*, 12(2-3):124–400, 2018.
- [27] Guillaume O Berger, Raphaël M Jungers, and Zheming Wang. Data-driven invariant subspace identification for black-box switched linear systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 32–37. IEEE, 2022.
- [28] Hassan Bevrani. *Robust power system frequency control*. Springer, 2009.
- [29] Hassan Bevrani, Toshifumi Ise, and Yushi Miura. Virtual synchronous generators: A survey and new perspectives. *International Journal of Electrical Power & Energy Systems*, 54:244–254, 2014.
- [30] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [31] S. L. Brunton and J. N. Kutz. *Data Driven Science & Engineering - Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2017.
- [32] Giuseppe Carlo Calafiore and Marco C Campi. The scenario approach to robust control design. *IEEE Transactions on automatic control*, 51(5):742–753, 2006.
- [33] D. S. Callaway and I. A. Hiskens. Achieving controllability of electric loads. *Proceedings of the IEEE*, 99(1):184–199, Jan 2011.
- [34] Eduardo F Camacho and Carlos Bordons. Introduction to model predictive control. In *Model Predictive Control*, pages 1–11. Springer, 2007.
- [35] Nathalie Cauchi and Alessandro Abate. StocHy: automated verification and synthesis of stochastic processes. In *25th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2019.

- [36] J.H. Chow and K.W. Cheung. A toolbox for power system dynamics and control engineering education and research. *IEEE Transactions on Power Systems*, 7(4):1559–1564, 1992.
- [37] Matthew Cleaveland, Lars Lindemann, Radoslav Ivanov, and George J Pappas. Risk verification of stochastic systems with neural network controllers. *Artificial Intelligence*, 313:103782, 2022.
- [38] Max H. Cohen and Calin Belta. Model-based reinforcement learning for approximate optimal control with temporal logic specifications. In *HSCC '21: 24th ACM International Conference on Hybrid Systems: Computation and Control, Nashville, Tennessee, May 19-21, 2021*, pages 12:1–12:11. ACM, 2021.
- [39] Max H Cohen, Zachary Serlin, Kevin Leahy, and Calin Belta. Temporal logic guided safe model-based reinforcement learning: A hybrid systems approach. *Nonlinear Analysis: Hybrid Systems*, 47:101295, 2023.
- [40] Wenqi Cui, Yan Jiang, and Baosen Zhang. Reinforcement learning for optimal primary frequency control: A lyapunov approach. *IEEE Transactions on Power Systems*, 38(2):1676–1688, 2022.
- [41] Imma Curiel. *Cooperative game theory and applications: cooperative games arising from combinatorial optimization problems*, volume 16. Springer Science & Business Media, 2013.
- [42] Fatemeh Daneshfar and Hassan Bevrani. Multiobjective design of load frequency control using genetic algorithms. *International Journal of Electrical Power & Energy Systems*, 42(1):257–263, 2012.
- [43] Howard Demuth. *Neural Network Toolbox Documentation. NN with Matlab*, 2004.
- [44] Paul Devaux and Mohammed Mehdi Farid. Benefits of PCM underfloor heating with PCM wallboards for space heating in winter. *Applied Energy*, 191:593–602, 2017.
- [45] Alex Devonport, Adnane Saoud, and Murat Arcak. Symbolic abstractions from data: A PAC learning approach. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 599–604. IEEE, 2021.
- [46] Franck Djeumou, Abraham P. Vinod, Eric Goubault, Sylvie Putot, and Ufuk Topcu. On-the-fly control of unknown systems: From side information to performance guarantees through reachability. *IEEE Transactions on Automatic Control*, pages 1–16, 2022.
- [47] Mohammad Dreidy, H. Mokhlis, and Saad Mekhilef. Inertia response and frequency control techniques for renewable energy sources: A review. *Renewable and Sustainable Energy Reviews*, 69:144–155, 2017.
- [48] Tom Elliott, Joachim Geske, and Richard Green. White paper: Can you make money from active buildings? challenges facing business models. Technical report, Active Buildings Centre Research Programme, 2020.

- [49] Energy Emergencies Executive Committee. GB power system disruption – 09 AUGUST 2019. Technical Report January, Department for Business, Energy and Industrial Strategy, 2020.
- [50] Anne Mai Ersdal, Lars Imsland, and Kjetil Uhlen. Model predictive load-frequency control. *IEEE Transactions on Power Systems*, 31(1):777–785, 2015.
- [51] Peyman Mohajerin Esfahani, Tobias Sutter, and John Lygeros. Performance bounds for the scenario approach and an extension to a class of non-convex programs. *IEEE Transactions on Automatic Control*, 60(1):46–58, 2014.
- [52] Chuchu Fan, Bolun Qi, Sayan Mitra, and Mahesh Viswanathan. DryVR: Data-Driven Verification and Compositional Reasoning for Automotive Systems. In *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24–28, 2017, Proceedings, Part I*, volume 10426 of *Lecture Notes in Computer Science*, pages 441–461. Springer, 2017.
- [53] Mohsen Farahani, Soheil Ganjefar, and Mojtaba Alizadeh. Pid controller adjustment using chaotic optimisation algorithm for multi-area load frequency control. *IET Control Theory & Applications*, 6(13):1984–1992, 2012.
- [54] João Figueiredo and José Sá da Costa. A SCADA system for energy management in intelligent buildings. *Energy and Buildings*, 2012.
- [55] Gene F Franklin, J David Powell, Abbas Emami-Naeini, and J David Powell. *Feedback control of dynamic systems*, volume 4. Prentice hall Upper Saddle River, 2002.
- [56] Goran Frehse, Colas Le Guernic, Alexandre Donzé, Scott Cotton, Rajarshi Ray, Olivier Lebeltel, Rodolfo Ripado, Antoine Girard, Thao Dang, and Oded Maler. SpaceEx: Scalable Verification of Hybrid Systems. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification*, pages 379–395, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [57] Oleg Gasparyan. *Linear and nonlinear multivariable feedback control: a classical approach*. John Wiley & Sons, 2008.
- [58] Mukesh Gautam, Narayan Bhusal, and Mohammed Benidris. A cooperative game theory-based approach to under-frequency load shedding control. In *2021 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2021.
- [59] Mukesh Gautam, Narayan Bhusal, Mohammed Benidris, and Hanif Livani. A cooperative game theory-based secondary frequency regulation in distribution systems. In *2021 North American Power Symposium (NAPS)*, pages 1–6. IEEE, 2021.
- [60] Vahid Gholamrezaie, Mehdi Ghazavi Dozein, Hassan Monsef, and Bin Wu. An optimal frequency control method through a dynamic load frequency control (lfc) model incorporating wind farm. *IEEE Systems Journal*, 12(1):392–401, 2017.

- [61] A. Girard and G. J. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.
- [62] Antoine Girard. Reachability of uncertain linear systems using zonotopes. In *Proceedings of the 8th International Conference on Hybrid Systems: Computation and Control*, HSCC’05, page 291–305, Berlin, Heidelberg, 2005. Springer-Verlag.
- [63] Antoine Girard and George J Pappas. Approximate bisimulation relations for constrained linear systems. *Automatica*, 43(8):1307–1317, 2007.
- [64] Antoine Girard and George J Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.
- [65] Antoine Girard and George J Pappas. Approximate bisimulation: A bridge between computer science and control theory. *European Journal of Control*, 17(5-6):568–578, 2011.
- [66] Antoine Girard, Giordano Pola, and Paulo Tabuada. Approximately bisimilar symbolic models for incrementally stable switched systems. *IEEE Transactions on Automatic Control*, 55(1):116–126, 2009.
- [67] Girard, Antoine and Pappas, George. J . Hierarchical control system design using approximate simulation. *Automatica*, 45(2):566–571, 2009.
- [68] Lazaros Gkatzikis, Iordanis Koutsopoulos, and Theodoros Salonidis. The role of aggregators in smart grid demand response markets. *IEEE Journal on Selected Areas in Communications*, 31(7):1247–1257, 2013.
- [69] DM Greenwood, Khim Yan Lim, C Patsios, PF Lyons, Yun Seng Lim, and PC Taylor. Frequency response services designed for energy storage. *Applied Energy*, 203:115–127, 2017.
- [70] Kush Grover, Fernando dos Santos Barbosa, Jana Tumova, and Jan Kretinsky. Semantic Abstraction-Guided Motion Planning for sLTL Missions in Unknown Environments. In *Robotics: Science and Systems*, 2021.
- [71] Laurens Haan and Ana Ferreira. *Extreme value theory: an introduction*, volume 3. Springer, 2006.
- [72] G. Hackenberg, M. Irlbeck, V. Koutsoumpas, and D. Bytschkow. Applying formal software engineering techniques to smart grids. In *2012 First International Workshop on Software Engineering Challenges for the Smart Grid (SE-SmartGrids)*, pages 50–56, June 2012.
- [73] Sofie Haesaert and Sadegh Soudjani. Robust dynamic programming for temporal logic control of stochastic systems. *IEEE Transactions on Automatic Control*, 66(6):2496–2511, 2020.
- [74] Akash Harapanahalli, Saber Jafarpour, and Samuel Coogan. Forward invariance in neural network controlled systems. *IEEE Control Systems Letters*, pages 1–1, 2023.

- [75] Yogesh V Hote and Shivam Jain. Pid controller design for load frequency control: Past, present and future challenges. *IFAC-PapersOnLine*, 51(4):604–609, 2018.
- [76] Kai-Chieh Hsu, Vicenç Rubies-Royo, Claire J Tomlin, and Jaime F Fisac. Safety and liveness guarantees through reach-avoid reinforcement learning. In *Robotics: Science and Systems*, 2021.
- [77] Kyle Hsu, Rupak Majumdar, Kaushik Mallik, and Anne-Kathrin Schmuck. Multi-layered abstraction-based controller synthesis for continuous-time systems. In *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (part of CPS Week)*, pages 120–129, 2018.
- [78] Junjie Hu, Guangya Yang, Koen Kok, Yusheng Xue, and Henrik W Bindner. Transactive control: a framework for operating power systems characterized by high penetration of distributed energy resources. *Journal of Modern Power Systems and Clean Energy*, 2017.
- [79] Bing Huang, Alvaro A. Cardenas, and Ross Baldick. Not everything is dark and gloomy: Power grid protections against IoT demand attacks. *Proceedings of the 28th USENIX Security Symposium*, pages 1115–1132, 2019.
- [80] Md. Monirul Islam, Xiao Zhong, Zeyi Sun, Haoyi Xiong, and Wenqing Hu. Real-time frequency regulation using aggregated electric vehicles in smart grid. *Computers & Industrial Engineering*, 134:11 – 26, 2019.
- [81] S. Izadkhast, P. Garcia-Gonzalez, and P. Frías. An aggregate model of plug-in electric vehicles for primary frequency control. *IEEE Transactions on Power Systems*, 30(3):1475–1482, May 2015.
- [82] Saber Jafarpour, Akash Harapanahalli, and Samuel Coogan. Interval reachability of nonlinear dynamical systems with neural network controllers. In Nikolai Matni, Manfred Morari, and George J. Pappas, editors, *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pages 12–25. PMLR, 15–16 Jun 2023.
- [83] Krishna C. Kalagarla, Rahul Jain, and Pierluigi Nuzzo. Model-Free Reinforcement Learning for Optimal Control of Markov Decision Processes Under Signal Temporal Logic Specifications. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 2252–2257, 2021.
- [84] M. Kamgarpour, C. Ellen, S. Soudjani, S. Gerwinn, J.L. Mathieu, N. Mullner, A. Abate, D.S. Callaway, M. Fränzle, and J. Lygeros. Modeling options for demand side participation of thermostatically controlled loads. In *International Conference on Bulk Power System Dynamics and Control (IREP)*, pages 1–15, August 2013.
- [85] M Karimi, H Mokhlis, K Naidu, S Uddin, and A H A Bakar. Photovoltaic penetration issues and impacts in distribution network – A review. *Renewable and Sustainable Energy Reviews*, 53:594–605, 2016.

- [86] Milad Kazemi, Rupak Majumdar, Mahmoud Salamati, Sadegh Soudjani, and Ben Wooding. Data-driven abstraction-based control synthesis. *arXiv preprint arXiv:2206.08069*, 2022.
- [87] Willett Kempton and Jasna Tomić. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1):268–279, 2005.
- [88] Florian Kerber and Arjan van der Schaft. Compositional analysis for linear control systems. In *Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control, HSCC '10*, page 21–30, New York, NY, USA, 2010. Association for Computing Machinery.
- [89] Nassim Khaled and Bibin Pattel. Chapter 2 - Theoretical Foundation of MPC. In Nassim Khaled and Bibin Pattel, editors, *Practical Design and Application of Model Predictive Control*. Butterworth-Heinemann, 2018.
- [90] Muhammad Khalid and Andrey V Savkin. Model predictive control based efficient operation of battery energy storage system for primary frequency control. In *2010 11th International Conference on Control Automation Robotics & Vision*, pages 2248–2252. IEEE, 2010.
- [91] G. J. Kish, M. Ranjram, and P. W. Lehn. A Modular Multilevel DC/DC Converter With Fault Blocking Capability for HVDC Interconnects. *IEEE Transactions on Power Electronics*, 30(1):148–162, 2015.
- [92] Niklas Kochdumper, Bastian Schürmann, and Matthias Althoff. Utilizing dependencies to obtain subsets of reachable sets. In *Hybrid Systems: Computation and Control, HSCC '20*, New York, NY, USA, 2020. Association for Computing Machinery.
- [93] Nour EL Yakine Kouba, Mohamed Mena, Mourad Hasni, and Mohamed Boudour. A new optimal load frequency control based on hybrid genetic algorithm and particle swarm optimization. *International journal on electrical engineering and informatics*, 9(3):418–440, 2017.
- [94] V. Koutsoumpas and P. K. Gupta. Towards a constraint based approach for self-healing smart grids. In *2013 2nd International Workshop on Software Engineering Challenges for the Smart Grid (SE4SG)*, pages 17–24, May 2013.
- [95] P. Kundur, N.J. Balu, and M.G. Lauby. *Power System Stability and Control*. EPRI power system engineering series. McGraw-Hill Education, 1994.
- [96] Vince Kurtz, Patrick M Wensing, and Hai Lin. Robust approximate simulation for hierarchical control of linear systems under disturbances. In *2020 American Control Conference (ACC)*, pages 5352–5357. IEEE, 2020.
- [97] Marta Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G. Gopalakrishnan and S. Qadeer, editors, *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011.

- [98] A. Lavaei and E. Frazzoli. Data-driven synthesis of symbolic abstractions with guaranteed confidence. *IEEE Control Systems Letters*, 7:253–258, 2022.
- [99] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani. Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica*, 146, 2022.
- [100] A. Lavaei, S. Soudjani, and M. Zamani. Compositional construction of infinite abstractions for networks of stochastic control systems. *Automatica*, 107:125–137, 2019.
- [101] A. Lavaei and M. Zamani. From dissipativity theory to compositional synthesis of large-scale stochastic switched systems. *IEEE Transactions on Automatic Control*, 67(9):4422–4437, 2022.
- [102] Abolfazl Lavaei, Mahmoud Khaled, Sadegh Soudjani, and Majid Zamani. AMYTISS: Parallelized automated controller synthesis for large-scale stochastic systems. In *Proc. 32nd International Conference on Computer Aided Verification (CAV)*, LNCS. Springer, 2020.
- [103] Abolfazl Lavaei, Ameneh Nejati, Pushpak Jagtap, and Majid Zamani. Formal safety verification of unknown continuous-time systems: A data-driven approach. In *Proceedings of the 24th International Conference on Hybrid Systems: Computation and Control, HSCC '21*, New York, NY, USA, 2021. Association for Computing Machinery.
- [104] Abolfazl Lavaei, Sadegh Soudjani, Emilio Frazzoli, and Majid Zamani. Constructing MDP abstractions using data with formal guarantees. *IEEE Control Systems Letters*, 7:460–465, 2022.
- [105] Maciej Ławryńczuk. Nonlinear predictive control of a boiler-turbine unit: A state-space approach with successive on-line model linearisation and quadratic optimisation. *ISA Transactions*, 67, 2017.
- [106] Edward A Lee. Cyber physical systems: Design challenges. In *2008 11th IEEE international symposium on object and component-oriented real-time distributed computing (ISORC)*, pages 363–369. IEEE, 2008.
- [107] Benoît Legat, Raphaël M Jungers, and Jean Bouchat. Abstraction-based branch and bound approach to q-learning for hybrid optimal control. In *Learning for Dynamics and Control*, pages 263–274. PMLR, 2021.
- [108] Karen Leung and Marco Pavone. Semi-supervised trajectory-feedback controller synthesis for signal temporal logic specifications. In *2022 American Control Conference (ACC)*, pages 178–185. IEEE, 2022.
- [109] Thomas Lew, Lucas Janson, Riccardo Bonalli, and Marco Pavone. A simple and efficient sampling-based algorithm for general reachability analysis. In *Learning for Dynamics and Control Conference*, pages 1086–1099. PMLR, 2022.
- [110] Yan Li, Peng Zhang, Matthias Althoff, and Meng Yue. Distributed formal analysis for power networks with deep integration of distributed energy resources. *IEEE Transactions on Power Systems*, 34(6):5147–5156, 2019.

- [111] Lars Lindemann and Dimos V. Dimarogonas. Control barrier functions for signal temporal logic tasks. *IEEE Control Systems Letters*, 3(1):96–101, 2019.
- [112] Lars Lindemann, Haimin Hu, Alexander Robey, Hanwen Zhang, Dimos Dimarogonas, Stephen Tu, and Nikolai Matni. Learning hybrid control barrier functions from data. In *Conference on Robot Learning*, pages 1351–1370. PMLR, 2021.
- [113] Lars Lindemann, George J. Pappas, and Dimos V. Dimarogonas. Control barrier functions for nonholonomic systems under risk signal temporal logic specifications, 2020.
- [114] Wei Liu, Wei Gu, Wanxing Sheng, Xiaoli Meng, Zaijun Wu, and Wu Chen. Decentralized multi-agent system-based cooperative frequency control for autonomous microgrids with communication constraints. *IEEE Transactions on Sustainable Energy*, 5(2):446–456, 2014.
- [115] Lennart Ljung. Perspectives on System Identification. *Annual Reviews in Control*, 34, 2010.
- [116] J. Löfberg. YALMIP : A Toolbox for Modeling and Optimization in MATLAB. In *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [117] Minyue Ma and Lingling Fan. Implementing consensus based distributed control in power system toolbox. In *2016 North American Power Symposium (NAPS)*, pages 1–6, 2016.
- [118] Nasif Mahmud and A. Zahedi. Review of control strategies for voltage regulation of the smart distribution network with high penetration of renewable distributed generation. *Renewable and Sustainable Energy Reviews*, 64:582–595, 2016.
- [119] Rupak Majumdar, Kaushik Mallik, Mateusz Rychlicki, Anne-Kathrin Schmuck, and Sadegh Soudjani. A flexible toolchain for symbolic rabin games under fair and stochastic uncertainties. In *International Conference on Computer Aided Verification*, pages 3–15. Springer, 2023.
- [120] Rupak Majumdar, Kaushik Mallik, Anne-Kathrin Schmuck, and Sadegh Soudjani. Symbolic control for stochastic systems via finite parity games. *Nonlinear Analysis: Hybrid Systems*, 51:101430, 2024.
- [121] Rupak Majumdar, Kaushik Mallik, and Sadegh Soudjani. Symbolic controller synthesis for Büchi specifications on stochastic systems. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control, HSCC’20*, New York, NY, USA, 2020. Association for Computing Machinery.
- [122] Rupak Majumdar, Necmiye Ozay, and Anne-Kathrin Schmuck. On abstraction-based controller design with output feedback. In *HSCC ’20: 23rd ACM International Conference on Hybrid Systems: Computation and Control, Sydney, New South Wales, Australia, April 21-24, 2020*, pages 15:1–15:11. ACM, 2020.

- [123] Rupak Majumdar, Mahmoud Salamati, and Sadegh Soudjani. Neural abstraction-based controller synthesis and deployment, 2023.
- [124] Anas Makdesi, Antoine Girard, and Laurent Fribourg. Efficient data-driven abstraction of monotone systems with disturbances. In *7th IFAC Conference on Analysis and Design of Hybrid Systems, ADHS 2021, Brussels, Belgium, July 7-9, 2021*, number 5 in IFAC-PapersOnLine, pages 49–54. Elsevier, 2021.
- [125] Anas Makdesi, Antoine Girard, and Laurent Fribourg. Data-driven models of monotone systems. *IEEE Transactions on Automatic Control*, 2023.
- [126] K. Mallik, A. Schmuck, S. Soudjani, and R. Majumdar. Compositional synthesis of finite-state abstractions. *IEEE Transactions on Automatic Control*, 64(6):2629–2636, 2019.
- [127] M. M. C. Merlin, T. C. Green, P. D. Mitcheson, D. R. Trainer, R. Critchley, W. Crookes, and F. Hassan. The alternate arm converter: A new hybrid multilevel converter with dc-fault blocking capability. *IEEE Transactions on Power Delivery*, 29(1):310–317, 2014.
- [128] L Miao, G Wei, X Fang, and J Risheng. The strategy of the voltage control in smart grid based on modern control method and FPGA. In *2015 34th Chinese Control Conference (CCC)*, 2015.
- [129] Ioanna Mitsioni, Pouria Tajvar, Danica Kragic, Jana Tumova, and Christian Pek. Safe data-driven contact-rich manipulation. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 120–127. IEEE, 2021.
- [130] Alexander McFarlane Mood. *Introduction to the Theory of Statistics*. McGraw-hill, 1950.
- [131] Bruce Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE transactions on automatic control*, 26(1):17–32, 1981.
- [132] MOSEK ApS. *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019.
- [133] C. Mu, W. Liu, and W. Xu. Hierarchically adaptive frequency control for an ev-integrated smart grid with renewable energy. *IEEE Trans. on Industrial Informatics*, 14(9):4254–4263, Sep. 2018.
- [134] Y. Mu, J. Wu, J. Ekanayake, N. Jenkins, and H. Jia. Primary frequency response from electric vehicles in the Great Britain power system. *IEEE Transactions on Smart Grid*, 4(2):1142–1150, June 2013.
- [135] Soumyadeep Nag and Namitha Philip. Application of neural networks to automatic load frequency control. In *Proceedings of The 2014 International Conference on Control, Instrumentation, Energy and Communication (CIEC)*, pages 345–350. IEEE, 2014.

- [136] National Grid. The electricity safety, quality and continuity regulations. Technical report, legislation.gov.uk, 2002.
- [137] National Grid. Future Requirements for Balancing Services. Technical report, National Grid, 2016.
- [138] National Grid. The Grid Code. Technical Report 5, National Grid, 2020.
- [139] National Grid. Enhanced Frequency Response. Technical Report V5.0, National Grid, 29 March 2016.
- [140] National Grid ESO. Firm Frequency Response Balancing Service. Technical Report V14.0, National Grid ESO, July 2019.
- [141] Sajid Nazir, Shushma Patel, and Dilip Patel. Assessing and augmenting SCADA cyber security: A survey of techniques. *Computers & Security*, 2017.
- [142] NETS. National Electricity Transmission System Security and Quality of Supply Standard Version 2.4. Technical report, National Grid, 1 April 2019.
- [143] CCC (Committee on Climate Change). Net Zero: The UK's Contribution to Stopping Global Warming, 2019.
- [144] Soroush Oshnoei, Arman Oshnoei, Ali Mosallanejad, and Farhad Haghjoo. Novel load frequency control scheme for an interconnected two-area power system including wind turbine generation and redox flow battery. *International Journal of Electrical Power & Energy Systems*, 130:107033, 2021.
- [145] M. Paramasivam, A. Salloum, V. Ajarapu, V. Vittal, N. B. Bhatt, and S. Liu. Dynamic optimization based reactive power planning to mitigate slow voltage recovery and short term voltage instability. *IEEE Transactions on Power Systems*, 28(4):3865–3873, 2013.
- [146] Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (SFCS 1977)*, pages 46–57, 1977.
- [147] Giordano Pola, Antoine Girard, and Paulo Tabuada. Approximately bisimilar symbolic models for nonlinear control systems. *Automatica*, 44(10):2508–2516, 2008.
- [148] Giordano Pola, Pierdomenico Pepe, and Maria Domenica Di Benedetto. Symbolic models for networks of control systems. *IEEE Transactions on Automatic Control*, 61(11):3663–3668, 2016.
- [149] Iraj Rahimi Pordanjani, Hooman Erfanian Mazin, and Wilsun Xu. A novel genetic programming approach for frequency-dependent modeling. *IEEE transactions on evolutionary computation*, 17(3):353–367, 2012.
- [150] Theis Bo Harild Rasmussen, Qiuwei Wu, Jakob Glarbo Møller, and Menglin Zhang. Mpc coordinated primary frequency support of small-and large-scale heat pumps. *IEEE Transactions on Smart Grid*, 13(3):2000–2010, 2022.

- [151] M. Mazhar Rathore, Awais Ahmad, Anand Paul, and Seungmin Rho. Urban planning and building smart cities based on the Internet of Things using Big Data analytics. *Computer Networks*, 2016.
- [152] Robert Reed, Luca Laurenti, and Morteza Lahijanian. Promises of deep kernel learning for control synthesis. *IEEE Control Systems Letters*, 2023.
- [153] G. Reissig, A. Weber, and M. Rungger. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE TAC*, 62(4):1781–1796, 2016.
- [154] Jonathan Reynolds. *Real-time and semantic energy management across buildings in a district configuration*. PhD thesis, Cardiff University, 2019.
- [155] Jonathan Reynolds, Muhammad Waseem Ahmad, and Yacine Rezgui. Holistic modelling techniques for the operational optimisation of multi-vector energy systems. *Energy and Buildings*, 169:397–416, 2018.
- [156] Seyed Hamid Reza Hosseini, Adib Allahham, Vahid Vahidinasab, Sara Louise Walker, and Phil Taylor. Techno-economic-environmental evaluation framework for integrated gas and electricity distribution networks considering impact of different storage configurations. *International Journal of Electrical Power & Energy Systems*, 125:1–13, 2021.
- [157] Robin Roche, Fabrice Lauri, Benjamin Blunier, Abdellatif Miraoui, and Abderrafiâa Koukam. Multi-agent technology for power system control. In *Power Electronics for Renewable and Distributed Energy Systems*, pages 567–609. Springer, 2013.
- [158] Katherine M. Rogers, Ray Klump, Himanshu Khurana, Angel A. Aquino-Lugo, and Thomas J. Overbye. An authenticated control framework for distributed voltage support on the smart grid. *IEEE Transactions on Smart Grid*, 1(1):40–47, 2010.
- [159] Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.
- [160] Pritam Roy, Paulo Tabuada, and Rupak Majumdar. Pessoa 2.0: A controller synthesis tool for cyber-physical systems. In *HSCC'11*, page 315–316. ACM, 2011.
- [161] Sergio Rozada, Dimitra Apostolopoulou, and Eduardo Alonso. Load frequency control: A deep multi-agent reinforcement learning approach. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.
- [162] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 1986.
- [163] Matthias Rungger and Majid Zamani. SCOTS: A tool for the synthesis of symbolic controllers. In *HSCC*, page 99–104. ACM, 2016.

- [164] Walid Saad, Zhu Han, H. Vincent Poor, and Tamer Başar. Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications. *IEEE Signal Processing Magazine*, 2012.
- [165] Sadra Sadraddini and Calin Belta. Formal guarantees in data-driven model identification and control synthesis. In *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (part of CPS Week), HSCC 2018, Porto, Portugal, April 11-13, 2018*, pages 147–156. ACM, 2018.
- [166] Ali Salamati, Abolfazl Lavaei, Sadegh Soudjani, and Majid Zamani. Data-driven verification and synthesis of stochastic systems via barrier certificates. *Automatica*, 159:111323, 2024.
- [167] Stanly Samuel, Kaushik Mallik, Anne-Kathrin Schmuck, and Daniel Neider. Resilient abstraction-based controller design. In *HSCC '20: 23rd ACM International Conference on Hybrid Systems: Computation and Control, Sydney, New South Wales, Australia, April 21-24, 2020*, pages 33:1–33:2. ACM, 2020.
- [168] Mohammad Javad Sanjari and Gevorg B Gharehpetian. Game-theoretic approach to cooperative control of distributed energy resources in islanded microgrid considering voltage and frequency stability. *Neural Computing and Applications*, 25:343–351, 2014.
- [169] Adnane Saoud, Antoine Girard, and Laurent Fribourg. Assume-guarantee contracts for continuous-time systems. *Automatica*, 134:109910, 2021.
- [170] Pierluigi Siano and Debora Sarno. Assessing the benefits of residential demand response in a real time distribution energy market. *Applied Energy*, 161:533–551, 2016.
- [171] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [172] SMA Solar Technology AG. Statement by SMA Solar Technology AG on the Cyber Security of PV Inverters (Horus Scenario). *White Paper*, 2017.
- [173] Richard Smith. System Operability Framework 2016. Technical Report November, National Grid, 2016.
- [174] S. Soudjani and A. Abate. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *IEEE Trans. on Control Systems Technology*, 23(3):975–990, 2015.
- [175] S. Soudjani, C. Gevaerts, and A. Abate. FAUST²: Formal abstractions of uncountable-state stochastic processes. In *TACAS*, volume 9035 of *LNCS*, pages 272–286. Springer, 2015.

- [176] Sadegh Soudjani and Alessandro Abate. Higher-order approximations for verification of stochastic hybrid systems. In S. Chakraborty and M. Mukund, editors, *Automated Technology for Verification and Analysis*, volume 7561 of *Lecture Notes in Computer Science*, pages 416–434. Springer Verlag, Berlin Heidelberg, 2012.
- [177] Sadegh Soudjani and Alessandro 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.
- [178] Sadegh Soudjani, Sebastian Gerwinn, Christian Ellen, Martin Fränzle, and Alessandro Abate. Formal synthesis and validation of inhomogeneous thermostatically controlled loads. In *International Conference on Quantitative Evaluation of Systems*, pages 57–73. Springer, 2014.
- [179] Sadegh Soudjani and Rupak Majumdar. Concentration of measure for chance-constrained optimization. *IFAC-PapersOnLine*, 51(16):277–282, 2018.
- [180] Soudjani, Sadegh and Abate, A. Aggregation of thermostatically controlled loads by formal abstractions. In *European Control Conference*, pages 4232–4237, Zurich, Switzerland, July 2013.
- [181] Aleksandar M. Stanković, Savo D. Đukić, and Andrija T. Sarić. Approximate bisimulation-based reduction of power system dynamic models. *IEEE Transactions on Power Systems*, 30(3):1252–1260, 2015.
- [182] Goran Strbac, Matt Woolf, Danny Pudjianto, Xi Zhang, Sara Walker, and Vahid Vahidinasab. White Paper: The Role of Active Buildings in the Transition to a Net Zero Energy System. Technical report, Active Buildings Centre Research Programme, 2020.
- [183] Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *Conference on Robot Learning*, pages 1519–1539. PMLR, 2021.
- [184] Hongbin Sun, Qinglai Guo, Junjian Qi, Venkataramana Ajarapu, Richard Bravo, Joe Chow, Zhengshuo Li, Rohit Moghe, Ehsan Nasr-Azadani, Ujjwol Tamrakar, Glauco N. Taranto, Reinaldo Tonkoski, Gustavo Valverde, Qiuwei Wu, and Guangya Yang. Review of Challenges and Research Opportunities for Voltage Control in Smart Grids. *IEEE Transactions on Power Systems*, 34:2790–2801, 2019.
- [185] Zexin Sun, Zhenyi Yuan, Changhong Zhao, and Jorge Cortés. Learning decentralized frequency controllers for energy storage systems. *IEEE Control Systems Letters*, 7:3459–3464, 2023.
- [186] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [187] T. H. Szymanski. Security and privacy for a green internet of things. *IT Professional*, 19(5):34–41, 2017.

- [188] Paulo Tabuada. *Verification and control of hybrid systems: A symbolic approach*. Springer US, 2009.
- [189] Z. Tan, P. Yang, and A. Nehorai. An optimal and distributed demand response strategy with electric vehicles in the smart grid. *IEEE Transactions on Smart Grid*, 5(2):861–869, March 2014.
- [190] Yuichi Tazaki and Jun-ichi Imura. Bisimilar finite abstractions of interconnected systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 514–527. Springer, 2008.
- [191] Kotub Uddin, Tim Jackson, Widanalage D. Widanage, Gael Chouchelamane, Paul A. Jennings, and James Marco. On the possibility of extending the lifetime of lithium-ion batteries through optimal V2G facilitated by an integrated vehicle and smart-grid system. *Energy*, 133:710–722, 2017.
- [192] Riaz Uddin, Syed Atif Naseem, and Zafar Iqbal. Formal reliability analyses of power line communication network-based control in smart grid. *IJCAS*, 17(12):3047–3057, Dec 2019.
- [193] Birgit Van Huijgevoort, Oliver Schön, Sadegh Soudjani, and Sofie Haesaert. Syscore: Synthesis via stochastic coupling relations. In *Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control*, pages 1–11, 2023.
- [194] Cees Ferdinand Verdier, Niklas Kochdumper, Matthias Althoff, and Manuel Mazo Jr. Formal synthesis of closed-form sampled-data controllers for nonlinear continuous-time systems under STL specifications. *Automatica*, 139:110184, 2022.
- [195] Carlo Vezzoli, Fabrizio Ceschin, Lilac Osanjo, Mugendi M’Rithaa, Richie Moalosi, Venny Nakazibwe, and Jan Carel Diehl. Designing Sustainable Energy for All. Sustainable Product-Service System Design Applied to Distributed Renewable Energy. *Designing Sustainable Energy for All*, 2018.
- [196] Abraham P. Vinod, Joseph D. Gleason, and Meeko M. K. Oishi. SReachTools: A MATLAB Stochastic Reachability Toolbox. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, HSCC ’19, page 33–38, New York, NY, USA, 2019. Association for Computing Machinery.
- [197] Petr Vorobev, David M. Greenwood, John H. Bell, Janusz W. Bialek, Philip C. Taylor, and Konstantin Turitsyn. Deadbands, Droop, and Inertia Impact on Power System Frequency Distribution. *IEEE Transactions on Power Systems*, 34(4):3098–3108, 2019.
- [198] Xiao Wang, Saasha Nair, and Matthias Althoff. Falsification-based robust adversarial reinforcement learning. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 205–212. IEEE, 2020.

- [199] Ye Wang, Vera Silva, and Miguel Lopez-Botet-Zulueta. Impact of high penetration of variable renewable generation on frequency dynamics in the continental europe interconnected system. *IET Renewable Power Generation*, 10(1):10–16, 2016.
- [200] Zheming Wang and Raphaël M Jungers. Probabilistic guarantees on the objective value for the scenario approach via sensitivity analysis. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5668–5673. IEEE, 2022.
- [201] Kandai Watanabe, Nicholas Renninger, Sriram Sankaranarayanan, and Morteza Lahijanian. Probabilistic specification learning for planning with safety constraints. In *Intelligent Robots and Systems (IROS)*, page TBA. IEEE, 2021.
- [202] Tsui-Wei Weng, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, Dong Su, Yupeng Gao, Cho-Jui Hsieh, and Luca Daniel. Evaluating the robustness of neural networks: An extreme value theory approach. In *International Conference on Learning Representations*, 2018.
- [203] Matthew Wicker, Luca Laurenti, Andrea Patane, and Marta Kwiatkowska. Probabilistic safety for bayesian neural networks. In Jonas Peters and David Sontag, editors, *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124 of *Proceedings of Machine Learning Research*, pages 1198–1207. PMLR, 03–06 Aug 2020.
- [204] GR Wood and BP Zhang. Estimation of the lipschitz constant of a function. *Journal of Global Optimization*, 8(1):91–103, 1996.
- [205] Ben Wooding. Using Formal Methods and Proof to Verify a CANDO Epilepsy Medical Device. Master’s thesis, Newcastle University, 2019.
- [206] Ben Wooding, Abolfazl Lavaei, and Sadegh Soudjani. Formal Control of New England 39-Bus Test System: An Assume-Guarantee Approach. *arXiv preprint arXiv:2307.03467*, 2023.
- [207] Ben Wooding, Abolfazl Lavaei, Vahid Vahidinasab, and Sadegh Soudjani. Robust Simulation Functions with Disturbance Refinement. In *European Control Conference (ECC) 2023*, 2023.
- [208] Ben Wooding, Vahid Vahidinasab, Milad Kazemi, and Sadegh Soudjani. Control and management of active buildings. In *Active Building Energy Systems*, pages 161–192. Springer, 2022.
- [209] Ben Wooding, Vahid Vahidinasab, and Sadegh Soudjani. Formal controller synthesis for frequency regulation utilising electric vehicles. In *2020 International Conference on Smart Energy Systems and Technologies (SEST)*, pages 1–6. IEEE, 2020.
- [210] Ben Wooding, Vahid Vahidinasab, and Sadegh Soudjani. Operation and control of a population of active buildings at network level. In *Active Building Energy Systems*, pages 193–218. Springer, 2022.

- [211] Di Wu, Nikitha Radhakrishnan, and Sen Huang. A hierarchical charging control of plug-in electric vehicles with simple flexibility model. *Applied Energy*, 2019.
- [212] Bai Xue, Miaomiao Zhang, Arvind Easwaran, and Qin Li. PAC model checking of black-box continuous-time dynamical systems. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39(11):3944–3955, 2020.
- [213] Shahram Yadollahi, Hamed Kebriaei, and Sadegh Soudjani. Generalized stochastic dynamic aggregative game for demand-side management in microgrids with shared battery. *IEEE Control Systems Letters*, 2023.
- [214] M. Yilmaz and P. T. Krein. Review of battery charger topologies, charging power levels, and infrastructure for plug-in electric and hybrid vehicles. *IEEE Transactions on Power Electronics*, 28(5):2151–2169, May 2013.
- [215] P P Zarina, S Mishra, and P C Sekhar. Exploring frequency control capability of a PV system in a hybrid PV-rotating machine-without storage system. *International Journal of Electrical Power & Energy Systems*, 60:258–267, 2014.
- [216] H. Zhang, J. Zhou, Q. Sun, J. M. Guerrero, and D. Ma. Data-Driven Control for Interlinked AC/DC Microgrids Via Model-Free Adaptive Control and Dual-Droop Control. *IEEE Transactions on Smart Grid*, 2017.
- [217] Bingzhuo Zhong, Majid Zamani, and Marco Caccamo. Synthesizing safety controllers for uncertain linear systems: A direct data-driven approach. In *2022 IEEE Conference on Control Technology and Applications (CCTA)*, pages 1278–1284, 2022.
- [218] D. Zonetti, A. Saoud, A. Girard, and L. Fribourg. A symbolic approach to voltage stability and power sharing in time-varying dc microgrids. In *2019 18th European Control Conference (ECC)*, pages 903–909, 2019.

List of symbols

\mathbb{N}	set of natural numbers
$\mathbb{N}_{\geq 0}$	set of non-negative natural numbers
\mathbb{R}	set of real numbers
$\mathbb{R}_{>0}$	set of positive real numbers
$\mathbb{R}_{\geq 0}$	set of non-negative real numbers
\mathbb{R}^n	Cartesian product of n copies of sets in \mathbb{R}
\emptyset	empty set
2^A	power set of a set A
t	time
τ	sampling time
X	set of states
X_0	set of initial states
Y	set of outputs
U	set of control inputs
V	set of external disturbances
W	set of internal disturbances
$X \rightarrow Y$	mapping from set X to set Y
$X_1 \times X_2$	Cartesian product of set X_1 with X_2
$\mathcal{B} = [\underline{\mathcal{B}}, \overline{\mathcal{B}}]$	interval on set of reals
$\overline{\mathcal{B}}$	upper bound of \mathcal{B}
$\underline{\mathcal{B}}$	lower bound of \mathcal{B}
\mathcal{B}	interval of safe set
\mathcal{A}	interval of avoid set
\mathcal{T}	interval of target set
\mathcal{S}	interval of statutory limits

\mathcal{L}	normal infeed loss
\mathcal{Z}	containment zone
$\mathbf{x}, \mathbf{x}(t)$	vector of system state (with respect to time)
\mathbf{y}	vector of system output
\mathbf{u}	vector system control input
\mathbf{v}	vector of external system disturbance
\mathbf{w}	vector of internal system disturbance
\mathbf{x}_i	i^{th} component of vector \mathbf{x}
$\dot{\mathbf{x}}, \frac{\partial \mathbf{x}}{\partial t}$	derivative of \mathbf{x} with respect to time
\mathbf{x}'	successor state from predecessor state \mathbf{x}
\mathbf{x}_{in}	initial state
$Post_{\mathbf{u}}(\mathbf{x})$	set of successor states of \mathbf{x} under control input \mathbf{u}
$U(\mathbf{x})$	set of control inputs of state \mathbf{x} that have successor states \mathbf{x}'
$ \cdot $	element-wise absolute value
$\ \cdot\ $	vector norm
\mathbb{I}^n	identity matrix of size $\mathbb{R}^{n \times n}$
A^T	transpose of a matrix A
$a \ll b$	a much less than b
f	power system frequency
f_0	nominal power system frequency (50 Hz or 60 Hz)
Δf	change in f from nominal value
Σ	transition system
Σ_1	original system model
Σ_2	(possibly) reduced-order abstraction
Σ^i	subsystem i of transition system Σ
$\hat{\Sigma}$	symbolic model or finite abstraction
$\mathfrak{B}(\Sigma)$	system behaviour
$\mathfrak{B}_{\mathbf{x}}(\Sigma)$	system behaviour initialised at \mathbf{x}
$\preceq_{\mathfrak{B}}$	behavioural inclusion
$\preceq_{\mathfrak{G}}$	simulation relation
$\preceq_{\mathfrak{G}}^{\epsilon}$	approximate simulation relation
$\cong_{\mathfrak{B}}$	behavioural equivalence
$\cong_{\mathfrak{G}}$	bisimulation
$\cong_{\mathfrak{G}}^{\epsilon}$	approximate bisimulation

$Reach(\Sigma)$	reachable set of Σ
$\mathcal{I}(\Sigma^1, \dots, \Sigma^N)$	interconnection of subsystems i through N
$\prod_{i=1}^N X^i$	Cartesian product of N sets of subsystem state spaces X
η_x	discretisation of state space
η_u	discretisation of input space
ϵ	simulation relation error
$\rho_u \in U^\omega$	word over set U
U^ω	set of infinite words over set U
ψ	LTL formula
\mathcal{A}	assumption
\mathcal{G}	guarantee
\mathcal{C}	contract
\preceq	refinement relation
\oplus	contract composition
L_φ	upper bound of Lipschitz constant of φ
$\zeta_{\mathbf{x}_0, \mathbf{u}, \mathbf{v}}$	continuous time trajectory
$1 - \beta$	confidence
$\Omega_\epsilon(c)$	ball with centre c and radius ϵ
$(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$	probability space
κ^e	modified growth bound
$N(\epsilon, \beta)$	number of samples
ϕ	nonlinear term of slope restriction
$\varrho(\cdot)$	class- κ functions
$\text{sat}(\cdot)$	saturation function

List of Abbreviations

3A3M	3 Area 3 Machine Power System
ABCD	Abstraction-Based Controller Design
ANN	Artificial Neural Network
BMS	Building Management System
CHP	Combined Heat and Power Unit
CPS	Cyber-Physical System
DAE	Differentiable Algebraic Equation
DDC	Data-driven Control
DER	Distributed Energy Resource
DSM	Demand-side Management
DSO	Distributed System Operator
DSR	Demand-side Response
EMS	Energy Management System
EFR	Enhanced Frequency Response
ESS	Energy Storage System
EV	Plug-in Electric Vehicle
FFR	Firm Frequency Response
GA	Genetic Algorithm
GP	Genetic Programming
GB	Great Britain
IoT	Internet of Things
LFDD	Low Frequency Demand Disconnection
LMI	Linear Matrix Inequalities
LTL	Linear Temporal Logic
MAS	Multi-agent System Control

MIMO	Multi-input Multi-output
ML	Machine Learning
MPC	Model Predictive Control
NETS	New England 39-Bus Test System
ODE	Ordinary Differential Equation
P2G	Power-to-Gas
PAC	Probably Approximately Correct
PMU	Phasor Measurement Unit
PV	Photovoltaic Panel
RL	Reinforcement Learning
RoCoF	Rate of Change of Frequency
ROM	Reduced-order Model
RSA	Robust Scenario Approach
RSF	Robust Simulation Function
RCP	Robust Convex Optimisation Program
SCADA	Supervisory Control and Data Acquisition
SCP	Scenario Convex Program
SISO	Single Input Single Output
SoC	State of Charge
TCL	Thermostatically Controlled Load
TSO	Transmission System Operator
VSG	Virtual Synchronous Generator

Curriculum Vitae

Benjamin James Wooding was born on 16 November 1996 in Liverpool, England. His early years were spent living on the compound of Kiwoko Hospital, in rural Uganda, with his parents and younger sister. Growing up, he went to school at the International School of Uganda (ISU) in Kampala, Uganda, and Cheney School in Oxford, England.

It was at ISU that he developed his interest in mathematics, taking IGCSE Mathematics a year early and IGCSE Additional Mathematics. He was also a talented athlete; participating in football, basketball, tennis, track and field, and volleyball tournaments. In 2013, he returned to Oxford to study A-levels. He took A-level Mathematics one year early, as well as studying Physics, Chemistry and Further Mathematics. He also continued to play volleyball, including representing the South East regional team.

In 2015, Ben was accepted to Newcastle University with a Sport Scholarship for Volleyball. He studied an integrated Masters degree in Computer Science, focusing on Security and Resilience. At this time he became interested in formal methods and the need for mathematical guarantees for safety-critical systems. His Masters dissertation investigated formal methods to verify an epilepsy medical device [205]. He received 1st Class Honours for his Masters degree, with a score of 92% on the dissertation.

In 2019, he started an EPSRC PhD Studentship at the School of Computing, Newcastle University, supervised by Dr Sadegh Soudjani. His research focus was the intersection of formal methods and control theory applied to power system frequency regulation. He has contributed to the international academic community with published works, program committee memberships, conference and journal paper reviews, and conference presentations.

At Newcastle University, he has been the chair of the AMBER research group, he has given multiple internal research presentations, and assisted widely by teaching as a demonstrator, and marking. He has experience with supervising BSc, MSc and PhD student projects. Alongside these, he has been a responsible person for research communication and dissemination from the HyCoDeV Lab.

In 2023, Ben was awarded the EPSRC Doctoral Prize Fellowship to continue high-quality research at Newcastle University. His research topic is *Reliable AI-enabled Design of Cyber-Physical Systems* working with Dr Abolfazl Lavaei.

List of Publications

Publications: Journal

1. **B. Wooding**, A. Lavaei, S. Soudjani, "Formal Control for the New England 39-Bus Test System: An Assume-Guarantee Contract Approach", *under review*, 2023
2. A.S. Laino, **B. Wooding**, S. Soudjani, R. Davenport, "Logic-Based Robustness for Resilience of Water Resource Recovery Facilities (WRRFs)", *under submission*, 2023
3. M. Kazemi*, R. Majumdar, M. Salamati*, S. Soudjani, **B. Wooding***, "Data-Driven Abstraction-Based Controller Synthesis", *under review*, 2022 (**contributed equally*)

Publications: Conference

1. **B. Wooding**, A. Lavaei, V. Vahidinisab, S. Soudjani, "Robust Simulation Functions with Disturbance Refinement", *2023 European Control Conference (ECC)*, 2023
2. S. Bogomolov, J. Fitzgerald, FF. Foldager, C. Gamble, PG. Larsen, K. Pierce, P. Stankaitis, **B. Wooding**, "Tuning Robotti: the machine-assisted exploration of parameter spaces in multi-models of a cyber-physical system", *18th International Overture Workshop*, 2021 (*authors in alphabetical order*)
3. **B. Wooding**, V. Vahidinisab, S. Soudjani, "Formal Controller Synthesis for Frequency Regulation Utilising Electric Vehicles", *Smart Energy Systems and Technologies (SEST)*, 2020

4. A. Abate, H. Blom, N. Cauchi, J. Delicaris, A. Hartmanns, M. Khaled, A. Lavaei, C. Pilch, A. Remke, S. Schupp, F. Shmarov, S. Soudjani, A. Vinod, **B. Wooding**, M. Zamani, and P. Zuliani, "ARCH-COMP20 Category Report: Stochastic Models", *"7th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH20)"*, 2020 (*authors in alphabetical order*)

Publications: Book Chapter

1. **B. Wooding**, V. Vahidinisab, M. Kazemi, S. Soudjani, "Cyber-Physical Smart Homes/Buildings", *accepted book chapter*, 2023
2. **B. Wooding**, V. Vahidinisab, M. Kazemi, S. Soudjani, "Control and management of active buildings", *Active Building Energy Systems: Operation and Control*, 2021
3. **B. Wooding**, V. Vahidinisab, S. Soudjani, "Operation and control of a population of active buildings at network level", *Active Building Energy Systems: Operation and Control*, 2021