Aggregation of a Heterogeneous Population of Solar Panels: Verification and Control

Andrea Peruffo

Linacre College
University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Trinity 2020

Abstract

The ever-growing presence of renewable energy sources has started a radical transformation of power grids worldwide. Their dynamical characteristics, the connections via power electronics, and their geographically distributed locations render their connection to the power grid substantially different from the traditional synchronous machines. Household solar panels are a striking example of this transition: their distribution is extremely sparse, and individual devices have a negligible contribution to the global electricity network; however, the presence of a large population of such devices influences the power grid in ways yet to be fully understood.

In the first part of this thesis we test the behaviour of a variety of solar devices, in order to develop models representing the aggregated, heterogeneous population that is connected to the power network. We then present a model of the power grid whose parameters depend on the amount of solar devices connected to the grid. In particular we focus on the network frequency, and we study how features of the PV population affect the overall frequency signal. We investigate the frequency response after a generation loss contingency at varying levels of renewable penetration. We aim at identifying scenarios that lead to significant frequency deviations which activate a load shedding procedure.

Simulations of critical circumstances are useful to highlight potential issues, yet are often not sufficient to guarantee the reliability of a complex, stochastic system as the power grid. In the second part of this dissertation we introduce formal methods, a suite of techniques to model complex systems as mathematical entities and to outline tight requirements about such systems. In particular we rely on the formal abstraction technique, which translates a stochastic system into a Markov model. By means of model checking, we assess the satisfaction of given requirements and we provide a formal guarantee of correctness.

Our formal tests must specify a significant number of parameters as we consider various populations of solar devices. The perfect tuning of such parameters is a cumbersome, handmade responsibility. Finally, we automate this task proposing a parameter synthesis framework which automatically returns the values of parameters that satisfy the predefined requirements.

Aggregation of a Heterogeneous Population of Solar Panels: Verification and Control

Andrea Peruffo

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2020

Acknowledgements

This thesis writes the final chapter to my four-year DPhil journey as part of the Oxford Control and Verification (OxCaV) group, within the Department of Computer Science. Whilst the last year has been marked by the pandemic, it has been a journey made by amazing companions and true relationships. In the following few lines, I hope to acknowledge their guidance, support and friendship.

First, I would like to express my deep gratitude to Alessandro, for his dedication, support and guidance. I would like to thank him for being a source of motivation, for the massive freedom. Thank you for the many thoughtful exchanges of ideas, opportunities, trust, advice and friendship along the way.

Second, I would like to express my gratitude to my examiners Dan Rogers and George Konstantopoulos for the constructive, precise, thorough feedback. Thanks also to my assessors Niky Trigoni, Malcom McCulloch, Kostas Margellos for their fruitful guidance and support. I will not forget my college advisor, Cezar Ionescu: for the delightful lunches, the precious dinners, the genuine conversations during my stay in Oxford.

I value the amazing research and personally extraordinary experience at the ERATO MMSD project lab in Tokyo. I would like to specifically thank Ichiro Hasuo, Satoshi Kura, Hiroshi Unno, Riccardo Treglia, Johan Ribom, Ivar Bengtsson, Agustin Martinez Suñé, Stefan Klikovits, Tsutomu Kobayashi, for their hospitality, work ethic and friendship. Thanks to you, I felt at home.

For the work related to the synthesis of Lyapunov and barrier functions, I would like to thank my co-authors. Namely, I would like to thank Daniele Ahmed, for your constant support, guidance and friendship; Alec Edwards, for your help and feedback, now the project rests on your shoulders; Mirco Giacobbe for the sharing of ideas and detailed discussions.

I had the privilege to share my time at Oxford with a group of smart, funny, witty people. A sincere thank you to the people of OxCaV (and beyond) Hosein

Hasanbeig, Francesco Cosentino, Gareth Molyneux, Muhammad Syifa'ul Mufid, Kyriakos Polymenakos, Timothy Seabrook, Viraj Wijesuriya, Mehran Hosseini. A special mention is due to Nathalie Cauchi and Elizabeth Polgreen, for the friendship and guidance, for the many hours spent sharing our thoughts. I had enjoyed these years so much and I will always be grateful: this is your fault.

I would like to thank all my friends back in Italy. Even from far away, I know their support is there. Many things change, we may be distant, yet home is home.

The concluding acknowledgement goes to my family, my mother Nadia, my father Claudio, my sister Giulia and to Sofia. You are my strength, my motivation, my support. It has not always been easy, and I can solemnly promise this will not change in the future.

Abstract

The ever-growing presence of renewable energy sources has started a radical transformation of power grids worldwide. Their dynamical characteristics, the connections via power electronics, and their geographically distributed locations render their connection to the power grid substantially different from the traditional synchronous machines. Household solar panels are a striking example of this transition: their distribution is extremely sparse, and individual devices have a negligible contribution to the global electricity network; however, the presence of a large population of such devices influences the power grid in ways yet to be fully understood.

In the first part of this thesis we test the behaviour of a variety of solar devices, in order to develop models representing the aggregated, heterogeneous population that is connected to the power network. We then present a model of the power grid whose parameters depend on the amount of solar devices connected to the grid. In particular we focus on the network frequency, and we study how features of the PV population affect the overall frequency signal. We investigate the frequency response after a generation loss contingency at varying levels of renewable penetration. We aim at identifying scenarios that lead to significant frequency deviations which activate a load shedding procedure.

Simulations of critical circumstances are useful to highlight potential issues, yet are often not sufficient to guarantee the reliability of a complex, stochastic system as the power grid. In the second part of this dissertation we introduce formal methods, a suite of techniques to model complex systems as mathematical entities and to outline tight requirements about such systems. In particular we rely on the formal abstraction technique, which translates a stochastic system into a Markov model. By means of model checking, we assess the satisfaction of given requirements and we provide a formal guarantee of correctness.

Our formal tests must specify a significant number of parameters as we consider various populations of solar devices. The perfect tuning of such parameters is a cumbersome, handmade responsibility. Finally, we automate this task proposing a parameter synthesis framework which automatically returns the values of parameters that satisfy the predefined requirements.

Contents

List of Figures

List of Abbreviations

cdf . . . . . Cumulative Distribution Function

CPS . . . . . Cyber-Physical System

DG . . . . . Distributed Generation

dtMC . . . . . Discrete Time Markov Chain

dtMP . . . . . Discrete Time Markov Process

ENTSO-E . . European Network of Transmission System Operators for Electricity

LFC . . . . . Load Frequency Control

LTI . . . . . Linear Time Invariant

MC . . . . . Markov Chain

MP . . . . . Markov Process

MPPT . . . . . Maximum Power Point Tracking

pdf . . . . . Probability Density Function

PV . . . . . Photovoltaic

RES . . . . . Renewable Energy Sources

SDE . . . . . Stochastic Difference Equation

SHS . . . . . Stochastic Hybrid System

SMC . . . . . Statistical Model Checking

TCL . . . . . Thermostatically Controlled Load

TSO . . . . . Transmission System Operator

1 Introduction

Contents

1.1 Motivations and Research Questions

Academia, industry and the general public alike have recently shown an increased interest on renewable energy sources, which are seen as a core component towards environmental conservation and mitigation of rapid climate changes. Bolstered by initiatives such as the 2017 Paris Agreement [1] on climate, the deployment of renewables, e.g. hydroelectric, wind, and solar energy, has considerably grown and is planned to additionally increase worldwide [2–5]. The efforts in the past decade have been remarkable, as the installed capacities of wind and solar have increased by a factor of 6 and 40, respectively [6], bringing the total worldwide installed capacity to 650 GW of wind power [7] and 635 GW of solar power [8]. In 2019, solar and wind accounted for 2.1% and 4.8% of the world electricity production, respectively [9]. Solar energy is the fourth power source among renewable energy sources – more than 100 TWh produced in 2019 – in the EU area [2] and the third most important

renewable energy source, after hydro and wind power, in terms of global installed capacity [4]. By 2030, the European Union aspires to a 40% reduction in greenhouse gas emissions, a 32.5 % energy efficiency increase and a binding target for renewable energy (32%) [5]. The ultimate goal is the reduction of pollutants - especially CO_2 emissions - by means of installing a significant amount of renewable resources.

Conventional power sources (e.g. coal, nuclear) entail sizeable production sites endowed with mechanical inertia due to the presence of rotating bodies [10]. Their models are well known, uncertainty on model parameters is limited and their estimation is largely reliable. On the other hand, renewable sources such as wind or solar production sites are usually spread over wide areas and entail a population of devices, unlike conventional energy sources which tend to be located at specific sites. This results in the so-called *distributed generation* and is one of the main drivers of this dissertation. Nowadays, populations of distributed units significantly contribute to the power generation of the electric grid and their use raises issues of network reliability, especially solar panels (we will interchangeably denote photovoltaic panels as PV or solar panels). Another element that differentiates conventional sources from renewable energy power plants is the power production, which does not follow usual demand patterns and is largely not controllable. As an example, the power production of a panel during a clear day follows the irradiance of the sun, with a maximum at around midday; on the contrary, power consumption usually peaks in the morning (around breakfast time) and in the evening (around dinner time). Monitoring the discrepancy between production and consumption requires a sophisticated control strategy and adequate storage facilities: both items are a (relative) novelty in power systems and subject of ongoing research.

Despite the remarkable attention on renewable energy sources, not many works address the connection between populations of PV systems (as aggregation of individual, household devices) and the dynamics of the electric grid. Large PV farms have been modelled and studied aggregatively in, e.g., [11–13]; however, there is no similar study of distributed solar power generation. At the level of aggregation of households or buildings, the underlying assumption considers a limited power

production and often consumed at the source, and hence the net contribution to the grid is negligible. This assumption is evidently not tenable in view of the growing importance of populations of PV: as an example, in Germany PV generated power – distributed over 1.6 million PV setups – has provided for approximately 7.2% of the electricity demand in 2017, with peaks of 60% during weekends and holidays [3]. The growing relevance of PV panels justifies the modelling, analysis and control of this energy source.

The French transmission operator RTE is naturally interested in finding how a cluster of solar devices affects the grid. In particular, they focus on a possible chain-disconnection of solar devices which may lead to larger network issues, as causing a load shedding. Motivated by this, we have started a collaboration in order to study a heterogeneous, large, distributed population of solar-powered generation devices with high levels of uncertainty and volatility. We focus on household systems, which are composed of a combination of devices – inverters, (smart) meters, batteries. In view of the coupling between a solar panel and its inverter (namely the device that converts the continuous output of a solar panel to alternate current that is fed to the electricity grid), whenever we mention PV systems, we implicitly consider the panel-inverter combination. Generally speaking, we will study heterogeneous *PV systems* (equivalently *solar* systems), where the heterogeneity is to be intended as *different disconnection and reconnection patterns*.

In this work, we discuss a new and general analysis of the dynamics of a population of PV systems that is connected to the grid, particularly as a function of their heterogeneity. Such population operates following a disconnection-reconnection mechanism depending on the periodic sampling of the network frequency signal. This mechanism complies with grid connection codes and national regulations that may vary over the years; further, different countries may have different regulations. A large population of devices also includes units built by several manufacturers, with different functionalities and performance degradation patterns, hence offering a rather inhomogeneous population of devices.

The first milestone of this dissertation presents a model for the heterogeneous population of PV systems, which encompasses the different rules for the connection to the grid and manufactures variance. We analyse data from real solar units and propose a generalised model which is both high-level and flexible to describe a whole population of devices. We then study the dynamics of the grid in presence of a large penetration of PV systems and connect the network's dynamics with the solar devices' population model. Drawing inspiration from the canonical power grid's control architecture, we devise a decentralised proportional control scheme using the panel's output to enhance the reliability of the network, especially under challenging scenarios for the grid. Our case studies concern the consequences of a power generation incident: we couple the population's heterogeneity with the network's response to a large frequency deviation.

The grid's behaviour after an incident can be encoded as a safety requirement: a (almost) flawless operation must be guaranteed under a variety of conditions. Such a certificate is hard (if at all possible) to obtain using simulation-based methods for two reasons: the number of simulations to play out is huge – think of the number of parameters to set at every iteration, as the power generation and power demand pattern, the weather conditions, the mixture of power sources – and the stochastic setting naturally arising from the power grid. Formal methods have proved to be valid analysis tools for the task at hand. They can easily handle very complex models and provide a mathematical proof of (virtually) any specification. We thus use tools and approaches from the formal verification community to provide certificates about the power grid reliability. The second milestone of the doctoral thesis is the formal abstraction of the network's model and consequent verification of safety properties, as the absence of load shedding after a contingency.

On these premises, we aim at answering the following questions:

- How do we design a general modelling framework which: (i) describes the connection patterns of a large, heterogeneous population of PV systems; (ii) captures the stochastic nature of the power grid to a satisfying degree; and (iii) is both theoretically sound and practical to use?

- How do we couple performance requirements with the modelling framework, in order to: (i) formally check the performance of the system against a certain specification; and (ii) design the maximum (minimum) amount of heterogeneity of a population to ensure a safe network's operation?

We propose the use of stochastic hybrid systems as the modelling framework which can combine both uncertainty and the elaborate dynamics of the solar population, while eagerly describing the complex system dynamics. Certificates of the safe operation are achieved applying formal methods to Markov models, which arise as a translation of stochastic hybrid systems. However, this introduces a well-known conundrum: general verification queries are undecidable, and their finite approximations are frequently doomed by the state-space explosion problem. This restricts the level of detail attainable for the verification of stochastic models and a trade-off must be achieved to conjugate resource consumption and modelling detail.

1.2 Overview of the Thesis

This thesis sets up a formal modelling, verification and synthesis framework for large populations of devices connected to the electric network. The Markov modelling can be adapted and used for all power components that have frequency-dependent dynamics. This aims to address the research questions identified in Section 1.1. Figure 1.1 shows a pictorial representation of the relationship between the different chapters within this thesis. After the introductory and background chapters, we divide the thesis into three areas: (i) the modelling framework, composed of Chapter 3 and 4; (ii) the simulation-based tests, offered in Chapter 5; (iii) the embedding with formal methods in Chapter 6 and 7.

A breakdown of the structure of this manuscript is as follows:

Chapter 2 introduces background material about power systems and formal methods. We present the structure of an electric network and stress the importance of frequency measurements. This is followed by some formal

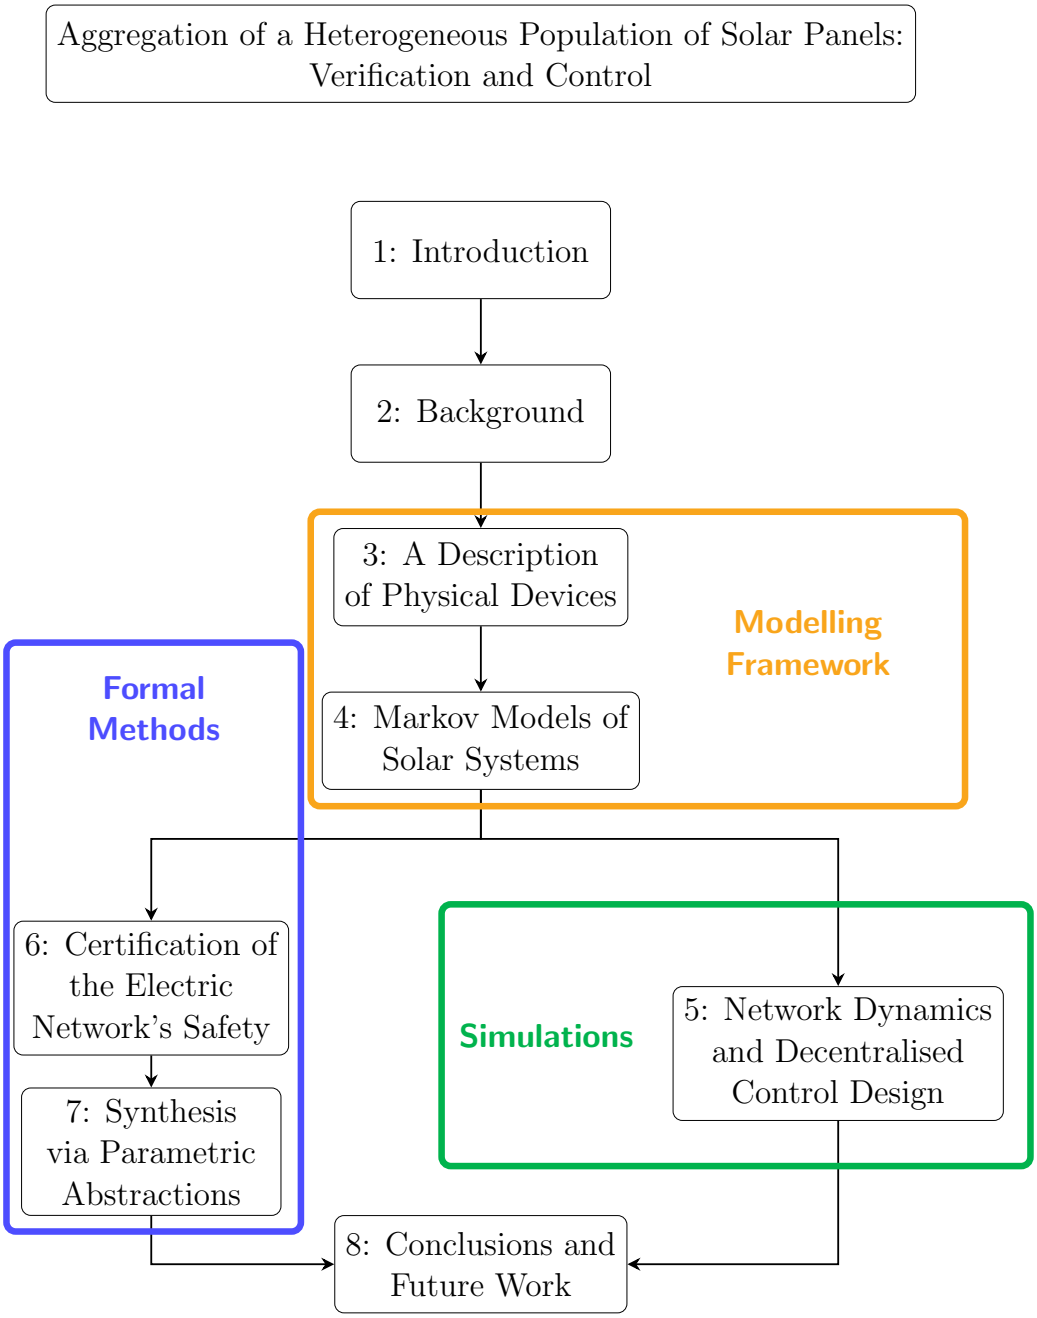

Figure 1.1: Structure of the dissertation.

methods and modelling preliminaries: stochastic systems, probabilistic models and formal abstractions.

Chapter 3 outlines the behaviour of a single solar device in response to frequency fluctuations. We further present data measured on real devices and generalise them, supported by European reports, to describe a large

population.

Chapter 4 builds up from the previous chapter and introduces a Markov model of a heterogeneous population of PV systems. The model of the aggregation exploits data-driven distributions of the working frequency intervals and delays structure.

Chapter 5 presents the model of the electric network and studies its sensitivity to solar penetration. Finally, we investigate how renewables influence the frequency response of the electricity grid in terms of load shedding risk via simulation-based scenarios. Further, we design a control action assuming PV systems are able to adjust their power output in response to a major frequency deviation. We study this hypothesis, designing a proportional control, in line with the network primary control.

Chapter 6 tests our methodology employing formal techniques. Methods as abstraction of non-linear systems are employed to certify the stability of the network, in terms of absence of black-out scenarios. The properties are asserted by varying distributions of PV systems' working intervals, population size, and network load.

Chapter 7 overturns the previous perspective by proposing a parameter synthesis framework. Whilst so far we have selected a numerical value for all the parameters at play and then have tested if the resulting model satisfied the safety specification, we now aim at directly asking the formal procedure to return the values of the parameters which make the model satisfy the requirements. We newly present a formal abstraction of parametric models that allows synthesising parameters to satisfy safety requirements.

Chapter 8 concludes the work by recapitulating the main contributions of this manuscript and outlines directions for future research.

1.3 Publications by the Author

The material presented in this doctoral thesis has appeared in top-tier international conference proceedings, or has been published in peer-reviewed journals. Each chapter builds upon different publications as follows:

Chapter 4 derives from the modelling framework originally published in [14]

Chapter 5 develops the incident scenarios and the tests under several conditions which are analysed in [15–17]

Chapter 6 introduces the formal abstraction approach for the electric grid as appears in [18, 19]

Chapter 7 extends the formal abstractions procedure to parametric models, as presented in [20]

Other publications by the author, listed below, are not included in this manuscript as they are not part of the same research path. In these publications, we devise a framework to synthesise Lyapunov functions and barrier certificates using a counterexample-based approach. This is outside the main focus of the present work, i.e. devising models and formally certify safety specifications of the electric grid. Nevertheless, these work provide formal certifications and they are concerned with the key aspect of the cyber-physical world: the stability of a dynamical system.

- Daniele Ahmed, Andrea Peruffo and Alessandro Abate, *Automated and Sound Synthesis of Lyapunov Functions with SMT Solvers*, in Tools and Algorithms for the Construction and Analysis of Systems (TACAS) - 26th International Conference, pages 97–114, Springer, 2020.

- Alessandro Abate, Daniele Ahmed, Mirco Giacobbe and Andrea Peruffo, *Formal Synthesis of Lyapunov Neural Networks*, in IEEE Control Systems Letters, pages 773-778, volume 5, number 3, 2020.

- Andrea Peruffo, Daniele Ahmed and Alessandro Abate, *Automated Formal Synthesis of Neural Barrier Certificates for Dynamical Models*, in Tools and Algorithms for the Construction and Analysis of Systems (TACAS) - 27th International Conference, Springer, 2021.

- Andrea Peruffo, Alec Edwards, Daniele Ahmed, Mirco Giacobbe, and Alessandro Abate, *FOSSIL: A Software Tool for the Formal Synthesis of Lyapunov Functions and Barrier Certificates using Neural Networks*, in 24th ACM International Conference on Hybrid Systems (HSCC) 2021.

2
Background

Contents

2.1 Power Systems

We concisely present basic concepts from the power systems literature which will be useful throughout this work. In particular we outline the crucial notions of network frequency, load shedding, and distributed generation.

2.1.1 Grid Frequency and Power Balance

Electrical power is an everyday commodity that is *just there*, whenever we need it, constantly. Although the electric connection is often taken for granted, it is anything but simple. Hundred of thousands of miles of high-voltage power lines and of overhead transmission lines in Europe connect electrical power plants to homes and businesses. Electricity travels from the source, i.e. a power plant, to our houses through a system called the *power grid* (interchangeably we will refer to it as electrical grid, electricity network or grid). To this day, large amounts of energy cannot be stored, hence electricity must be produced as it is used. Whilst battery storage solutions are strongly encouraged by governments' subsidies and have been steadily growing in the past five years, the global capacity is far from reaching 0.1% of today's consumption [21].

Swing Equation Broadly speaking, practically all power plants are synchronous machines: electro-mechanical transducers that convert mechanical power into AC electrical power. The moving parts of such machines, i.e. the rotors, are designed to rotate with constant angular velocity: this quantity defines the so-called *system frequency*, that we can consider homogeneous across the whole grid. The basic dynamic frequency model for a large power system is traditionally based on the swing equation [10] for a set of machines with rotating masses. The swing equation describes the electromechanical oscillations in a power system. The differential equation may be written in several equivalent forms and we will refer to the following

$$H\frac{d\,f}{dt} = P_m - P_e, \tag{2.1}$$

where f represents the frequency of the rotor; H is related to the total moment of inertia of the synchronous machine; P_m and P_e are the mechanical and electrical powers acting on the rotor, respectively.

Let us now discuss the various terms of Eq. (2.1), starting from the right-hand side to the left-hand side. The swing equation outlines how the difference between the mechanical power fed into and the electrical output power generates a motion

of the rotor with constant angular velocity ω_0, linked to the frequency by $\omega_0 = 2\pi f$. We may think of the mechanical power P_m originated by a hydro, steam or gas turbine. This power is determined by the gate opening of the turbine, and a variation of mechanical power takes several seconds in most cases. The mechanical power changes to regulate the network's frequency. Usually, each generator must comply with a hourly dispatch plan, which is determined by the expected load of the power system. Dispatch plans take into consideration the available resources at each hour-slot as well as the expected selling price and cost of electricity, and are the result of extensive optimisations: the mechanical input powers are thus planned in accordance with the hourly dispatches.

Even employing the best algorithms and the most powerful computational engines, exactly forecast the power consumption of the grid remains impossible; further, unexpected events or faults may always occur. The imbalance between power production and power demand causes a deviation of the system's frequency (cf. Section 2.1.3). This deviation may trigger few chosen machines to change their (electrical) output in order to level the discrepancy and restore the nominal value of frequency. However, whilst the frequency controller might be relatively fast, the machine itself takes some time to change its mechanical, hence electrical, power.

Network Frequency Throughout this work, we assume the grid frequency is described by only one differential equation, derived from the general form presented in Eq. (2.1). From a practical point of view, the notion of a unique network frequency is a mathematical abstraction denoting the homogeneous rate of change of phase. However, over a synchronous area, slightly different frequencies might appear in different points of the interconnected grid, within transient regimes. Bearing in mind the real-world difficulties of maintaining a uniform frequency over a large network, we place this work within an ideal framework: over a time frame of few seconds the frequency remains homogeneous over the whole grid. The network model defines the total inertia of the system H that, as per physical models, defines the resistance to a change in velocity. Large power networks traditionally have a

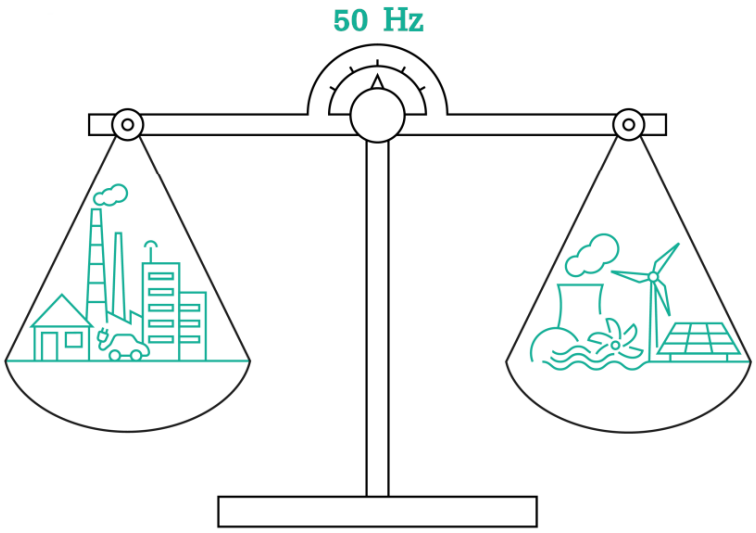

Figure 2.1: Frequency balance [22]

very large inertia coefficient that makes them slow to changes, but reliable and with a predictable output. However, the introduction of inertia-less, renewable power sources starts to challenge this modelling approach.

Frequency is the key physical quantity for a large network: as such, the control of its value is crucial. In particular, its value is a measure of mismatch between demand and generation: it will vary as load and generation change. As an example, the frequency will decrease after an increase of load or a generation reduction; the frequency will instead increase following a loss of load or an increase in generation. Figure 2.1 symbolically depicts the power balance between generation and demand, and Fig. 2.2 pictures the frequency values of the Continental Europe Synchronous Area during 24 hours. We rapidly notice the stochastic, noisy nature of the frequency signal. We thus define a range of operational values around to the nominal value ($f_0 = 50$ Hz): we depict with two red lines a 150 mHz interval around f_0, namely [49.925, 50.075] Hz. We account for a *significant* frequency deviation whenever the frequency is measured outside this interval. More detailed data analysis on the number of frequency deviations over the last 9 years is provided in the next chapter, in Section 3.4.

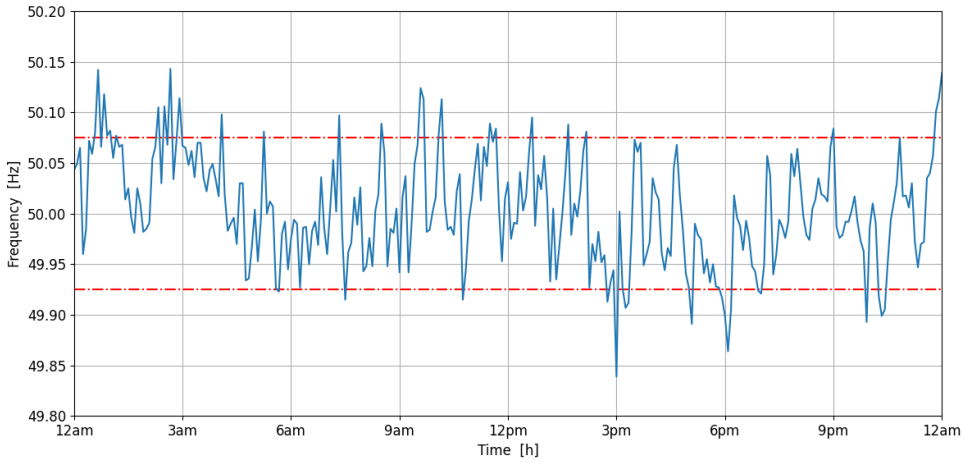

Figure 2.2: Frequency of the Continental Europe Synchronous Grid on 24/12/2015.

Synchronous Area We refer to a *synchronous area* as power grid that has international scale and operates at a synchronized system frequency being electrically intertwined during normal operations. In the following Chapters, we consider and model the synchronous grid of Continental Europe, which connects 26 countries and is the largest (by connected power) in the world. A wide synchronous area improves reliability and allows sharing resources, spreading both the energy generation and the load demand. On the other hand, problems are also shared. A recent example intertwines politics with energy provision: in 2018, a dispute between Kosovo and Serbia has led the whole Continental Europe grid to a months-long frequency drop. As a result, frequency-dependent clocks have run six minutes late [23].

2.1.2 Network Structure

The power grid is an interconnected network for delivering electricity from producers to consumers [24]. Its structure can be divided into three stages, as depicted in Fig. 2.3:

1. Power Generation: using either conventional sources (e.g. coal, natural gas) or renewable sources (e.g. wind, solar), electricity is usually generated in power plants located away from the populated areas.

2. Transmission: carrying electricity from the power plant to distribution substations. For the transmission of power over longer distances, the generated voltages are stepped up to reduce power losses – computed as $P_L = RI^2$. The transmitted electricity thus maintains the same output power by increasing voltage and reducing current[1].

3. Distribution: reduced-voltage electricity is finally delivered to consumers. Through a series of substations, the power is progressively stepped down to match consumers' needs, from large industrial consumers (4 kV up to 220 kV) to household users (110 V or 220 V). Household PV systems are located in this section of the grid, in contrast with large power sources.

The entity entrusted to monitor the transport of electricity from its generation to the local distribution is the transmission system operator (TSO). The TSO is required to maintain a continuous balance between electricity supply from power stations and demand from consumers, hence regulate the network frequency, and ensure sufficient reserves in case of sudden contingencies. Further, since the liberation of the energy market in 2009, the ENTSO-E (European Network of Transmission System Operators for Electricity [26]) gathers 42 TSOs from 35 countries to pursue the European energy and climate agenda, along with a further liberalisation of the energy market.

Control architectures are historically based on conventional power sources (e.g. coal, nuclear), that are endowed with mechanical inertia due to the presence of rotating bodies [10]. Hence, the control actions may take a few tens of seconds to be effective [27]. Solar energy and renewable sources in general can instead provide a much faster response, because they have no rotating mass [28]. We are witnessing a historical transition: from conventional, slow, predictable energy sources, the power grid is relying more and more on renewable, fast, less predictable

[1]With advancements in power electronics, HVDC (High Voltage DC) has proved efficient for longer distance transmission. HVDC cables are employed for very long distances or to connect networks with different frequency.

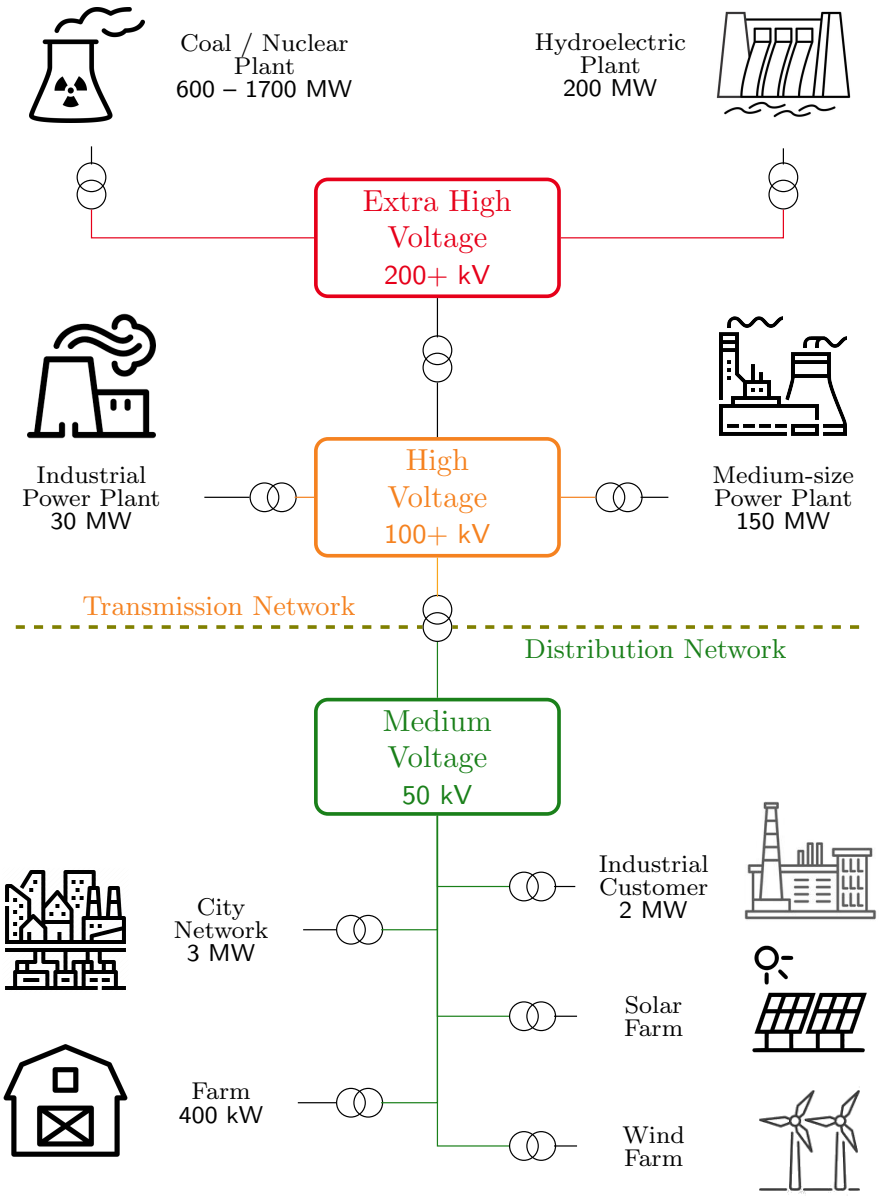

Figure 2.3: Electricity grid scheme, based on [25].

generation. The existing infrastructure and current control algorithm need to be updated and adapted to the inevitable generation shift.

2.1.3 Load Frequency Control

A power grid requires that generation and load demand closely balance at every moment; this balance can be judged by means of the system frequency. Frequency regulation can be compared to a cruise control of a car or a bike trying to travel at the same constant speed. On a level road it is easy to maintain speed. Approaching

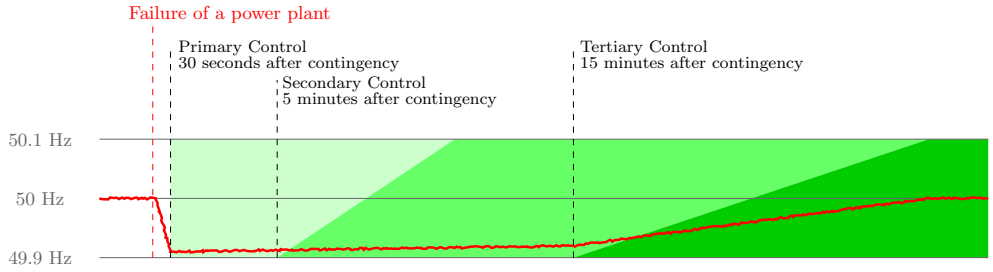

Figure 2.4: The three stages of control energy use [22].

a slope, however, the rider needs to make more effort to achieve the same speed; going downhill, the rider needs to apply the brakes to keep the same speed. Power consumption is indeed like the slope for the rider. If the consumption of electrical power is lower than production, the frequency is higher (the rider goes downhill); if consumption is higher than production, the frequency is lower (uphill). The electrical generators of an electricity grid rotate more readily and faster when consumption is low. Consequently, they rotate with a higher frequency. Conversely, the electrical generators rotate more laboriously and with a lower frequency when consumption is greater. Hence, when the frequency value is increasing, more power is being generated than used; if the system frequency is decreasing, the load demand is higher than the power generation.

In the entire European network the electrical generators are set up in such a way that they automatically and quickly respond to a change in grid frequency. They increase or lower their capacity depending on the level of consumption. To secure a high-quality operation of the electric network TSOs use a technique denoted Load-Frequency Control (LFC). LFC methods are classically based upon proportional-derivative-integral (PID) techniques [29, 30]; nevertheless, other approaches are available, as linear-quadratic regulators [31, 32] and recent developments include machine learning techniques, e.g. neural network [33, 34] or genetic algorithm [35] based techniques.

In the Continental Europe synchronous grid, control actions are performed in successive phases, each with different characteristics and goals. It entails a three-stage process, divided into *primary*, *secondary* and *tertiary* control, as depicted in

Fig. 2.4. First, primary control energy is activated. The turbines of the power plants throughout Europe respond to frequency fluctuations and increase or decrease their power accordingly. Primary control stabilises the network frequency at a stationary value after a disturbance or incident, and acts within a time-frame of seconds without restoring the frequency to its nominal value. After a few minutes, the secondary control replaces the primary control. The power is provided by power plants, which receive an automatic signal from the TSOs. Secondary Control maintains a balance between generation and consumption in a time-frame from 30 seconds up to typically 15 minutes after an incident. After 15 minutes, the operators manually switch to tertiary control energy. They instruct individual domestic or international power plants to feed more or less energy into the grid.

2.1.4 Islanding and Load Shedding

Islanding Islanding is the situation in which a distribution system becomes electrically isolated from the main power grid, yet continues to be energized by a distributed generator (DG) connected to it [36]. Traditionally, a distribution system does not include any active power generating source but, with the advent of DGs, this assumption no longer holds. In the islanded circuit, the balance between load and generation is likely to be violated in view of the absence of a strict frequency control, leading to frequency and voltage deviations. Hence, DGs must detect and rapidly disconnect from the island – a procedure called anti-islanding [37]. A natural example of islanding is a distribution network encompassing PV systems: such devices will continue to supply power as long as the solar irradiance is sufficient. For this reason, solar inverters are usually required to have an automatic anti-islanding strategy. Notably, microgrids may allow islanding operations [38]: in such scenario, the controller disconnects the local circuit from the grid and forces the DGs to provide power to match the local demand.

Load Shedding Electrical generation and transmission systems may not always meet the demand requirements. In extreme situations, when the frequency decreases rapidly and deviates significantly from the reference value, the demand must be lowered. Load shedding [39] is an intentionally engineered electrical power shutdown where electricity delivery is curtailed to selected network areas for a short time period, in order to avoid a widespread service disruption. Power systems are designed to withstand the effects of contingencies within given limits: the load shedding is a back-up network protection mechanism, a coarse tool for extreme situations. The ENTSO-E sets the limit value of 49.2 Hz [27]: if the network frequency trips below this value, a (large) load shedding procedure is automatically activated.

Nevertheless, little load shedding events are activated with smaller frequency deviations and thus are relatively common. As an example [40], on the 10^{th} January 2019, the Continental Europe Synchronous Area has faced an extraordinary frequency deviation in view of the coincidence of two events:

- A technical failure, caused by a frozen frequency measurement, has affected the German load frequency control and has resulted in a frequency deviation (in average) of -30 mHz;

- A deterministic frequency deviation has happened during the transition between one power source to another, as planned in the hourly schedule.

The cumulated effect has resulted in a frequency reaching the value of 49.808 Hz, precisely at 21:02. This low frequency value automatically triggers the activation of the RTE Industrial Interruptible Service: this disconnects 1250 MW within 5 seconds and, if the frequency value remains low, up to 1700 MW within 30 seconds. The load shedding procedure has helped the frequency to return within a normal operational interval.

2.1.5 Distributed Generation

In the last decade, pushed by a new environmental awareness and the opening of the energy market, distributed energy resources have risen as a clean and viable energy generation. Small generators and storage units, usually privately owned, can help lower the overall power demand or supply electricity to the local grid. Common examples are home-owners with excess power from their PV systems or a smart building connected to a wind turbine. As the 21st century progresses, the energy industry is migrating towards novel solutions, devices and infrastructures to deal with the ever-growing energy demand. The simple, yet effective, network topology outlined in Section 2.1.2 has started to evolve to accommodate for distributed generators. Whilst the power grid was originally built considering a unidirectional power flow, the introduction of DGs makes the power flow two ways. In the next years, the once clear distinction between transmission grids and distribution grids will continue to blur.

Low Inertia Environmental and sustainability awareness leads to the turnover of an important portion of conventional power plants with renewable energy-based generation. The replacement of synchronous machines, which carry their well-understood dynamics, hides the great challenge of connecting power electronics-based generating units to the electricity network, as their regulation and interactions are not fully understood yet.

Beyond the control-related concerns illustrated in Section 2.1.3, several other issues arise in a renewable-rich grid [6]. First, the modelling of the components and their relations require new efforts. In a low-inertia system, we need an accurate representation of the strengths and limitations of power electronic devices, especially on the (very) short time scales which were not a concern with conventional power sources. Another issue arises from inertia variability, as the amount of inertia is now heavily influenced on the generation combination. Several conventional power sources may be disconnected under weather conditions congenial to renewables, with peaks reaching 50% [3].

Notably, low inertia has a significant impact on system stability and frequency response. A reduced inertia leads to steeper responses reaching lower frequency values in the instants following an incident, as witnessed in [41], where the ENTSO-E reports the risks related to a massive adoption of low inertia devices in terms of lowest frequency values reached after a simulated incident.

Further, the global inertia constant H becomes heterogeneous. Studies [28] suggest that rather than a single constant we are in presence of several, different, local constants H_i for each individual sub-area, as functions of the proportion renewable generation versus conventional power sources. As such, the inertia (once) constants become time-variant, i.e. $H_i(t)$, in view of the variability of the power generation: frequency responses hence become differently fast within different areas.

Models and control designs for low-inertia grids are matter of ongoing investigation: a recent effort in [42] presents a model that captures the relevant physical properties and associated dynamic of a mixed-generation network, integrating both synchronous and renewable power sources. Control techniques for network with high penetration of solar and wind power may be split into two groups [43]: either via the integration of power reserves to supply or absorb additional power during times of low or high generation, or via the so-called inertia emulation and droop techniques, which use frequency deviations to release purposely stored power.

2.2 Formal Methods

2.2.1 General Notation

We denote the set of real numbers by $\mathbb{R}$ and the natural numbers (including 0) by $\mathbb{N}$. We use $[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\} \subset \mathbb{R}$ to denote the interval of reals between a and b including the boundaries. Given two homogeneous sets A, B, their difference is defined as $A \setminus B = \{x : x \in A \wedge x \notin B\}$. We represent their Cartesian product by $A \times B = \{(a, b) \mid a \in A \wedge b \in B\}$. A probability measure P for a sample space $\mathbb{X}$ and σ-algebra C defined over $\mathbb{X}$ is a non-negative map $\mathsf{P} : C \to [0, 1]$ such that $\mathsf{P}(\mathbb{X}) = 1$ and such that for all countable collections of $\{Z_i\}_{i=1}^{\infty}$ of pairwise disjoint sets in C, it holds that $\mathsf{P}(\bigcup_i Zi) = \sum_i \mathsf{P}(Z_i)$. The triple $(\mathbb{X}, C, \mathsf{P})$ defines

the probability space and has realisations $x \sim \mathcal{P}$. Further note, a special instance of C is the Borel σ-algebra. We denote the Borel σ-algebra on the set S as $\mathcal{B}(S)$ and the Borel measurable space is $(S, \mathcal{B}(S))$. We define a $\mathcal{N}(\mu, \sigma^2)$ as a normal distribution with mean μ and variance σ^2. If μ and σ are an n-dimensional vector and an $n \times n$ symmetric, positive semi definite matrix, respectively, $\mathcal{N}(\mu, \sigma^2)$ is a multi-dimensional distribution. We highlight general concepts using yellow boxes.

2.2.2 Stochastic Hybrid Systems

In literature, a wide portfolio of modelling frameworks are employed to capture the dynamics of solar panels with the surrounding components, such as the inverter, the storage battery, and so on [44–46].

Modelling Frameworks for PV Systems Based on the desired application, one can choose different modelling frameworks, from the microscopic details to the macroscopic effects: quantum models [47–49], equivalent electrical circuit models [47, 50–52], household PV systems [47], PV farms [11, 53], smart grids and demand response schemes [54–57].

Quantum models are physical models of the semiconductor materials which compose a PV cell. To understand the inner behaviour of a cell, we leverage quantum theory and the interactions between electrons and holes[2]. Physical models have been applied towards the improvement of the cell material and towards the development of new materials. Notably, in the last few years, the concept of quantum dot solar cells (QDSCs)[48, 58] has rapidly drawn attention. A QDSC is a solar cell that uses a tiny layer (few nanometers in width) of semiconductor material as the absorbing photovoltaic material.

Equivalent electrical circuits [47] are theoretical circuits that represent all of the electrical characteristics of a solar cell. The equivalent circuit simplifies calculation, helps the understanding the of basic behaviour of a device while aiding the analysis. The simplest and most common circuit evaluates an ideal solar cell as a current

[2]A *hole* is a positive charge originated from movement of an electron.

source in parallel with a diode; in practice, a shunt resistance and a series resistance component are added to the model to represent natural power losses. More complex models encompass additional components, e.g. adding another diode in parallel, taking into account the temperature dependency of the resistors [50–52].

Moving towards a macroscopic description of solar installations, we focus on the devices' power output. Models of a PV system (see Chapter 3 for a detailed dissertation) comprise solar panels, an inverter for the connection with the main grid, an optional battery storage system, and a meter to measure the input and output power flow. Models of PV farms maintain the same blocks, at a larger scale. We may define a solar farm as a large-scale PV system designed for the supply of power into the electricity network. As opposed to household systems, they supply power at the distribution or transmission level rather than to local users. This business has proved profitable and the construction of solar farms is on the rise: as an example, in 2020 a new 103 MW PV farm has been installed in Southern Italy [53].

A thriving research direction involves smart grids. A smart grid is, without doubt, the next-generation power grid, which uses a two-way flow of electricity and information to create a distributed, automated energy network [54]. It typically includes a variety of operation and energy measures as smart meters, smart appliances, and renewable energy sources. The study of smart grids is entangled with demand-response schemes [55] [56], and leaps towards the integration under the wide umbrella of cyber-physical systems (CPSs) [57].

The large variety of solar model frameworks possesses one common element: the models are merely an approximation of reality and inevitably contain uncertainty and inaccuracies, which arise from the chosen model or the sensor measurements used to identify and make predictions. Intuitively, uncertainty plays a crucial role in a model's ability to capture the behaviour of the PV system. To this end, we use models based on stochastic processes.

Solar devices may be represented as an amalgam of continuous (the network frequency) and discrete elements (the on/off power supply). A well-known and rich modelling framework that interleaves stochastic processes and draws upon

the hybrid theory stemming from continuous and discrete elements is known as Stochastic Hybrid Systems (SHSs) [59]. Recently SHSs have been deployed in a wide spectrum of fields with great success, as control systems [60, 61], formal methods [62, 63] and cyber-physical systems [64].

Stochastic Hybrid Systems Stochastic hybrid systems are a mathematical framework that captures the interactions between discrete features, continuous dynamics and probabilistic uncertainty [64]. Thanks to their universal characteristics, SHSs have been used to represent a variety of complex systems. An SHS is composed of discrete variables, which define probabilistic jumps among discrete locations, and continuous variables, which evolve according to stochastic differential equations within each location. Over the last few years, we have witnessed significant progress in this challenging research area that bridges notions from computer science, control engineering and stochastic analysis. A survey on the current and potential application of SHSs is reported in [64].

In this work, we focus on discrete-time SHS (dtSHS) according to the definitions in [59, 65].

Definition 1 (dtSHS [65]) *A discrete time stochastic hybrid system (or automata) is a collection* $\mathcal{H} = (Q, n, \mathscr{I}, T_x, T_q, R)$*, where*

- $Q := \{q_1, q_2, \ldots, q_m\}$ *with* $m \in \mathbb{N}$ *represents a discrete state space;*
- $n : Q \to \mathbb{N}$ *assigns to each discrete state* $q \in Q$ *the dimension of the continuous state space* $\mathbb{R}^{n(q)}$*;*
- $\mathscr{I} : \mathcal{B}(\mathcal{S}) \to [0, 1]$ *is a probability measure on* $\mathcal{S}$ *for the initialisation of the solution process;*
- $T_x : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{S} \to [0, 1]$ *is a conditional stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $\mathcal{S}$*. It assigns to each* $s = (q, x) \in \mathcal{S}$ *a probability measure,* $T_x(\cdot \mid s)$*, on the Borel space* $(\mathbb{R}^{n(q)}, \mathcal{B}(\mathbb{R}^{n(q)}))$*. The function* $T_x(A|(q, \cdot))$ *is assumed to be Borel measurable, for all* $q \in Q$ *and all* $A \in \mathcal{B}(R^{n(q)})$*;*

- $T_q : Q \times \mathcal{S} \to [0,1]$ *is a conditional discrete stochastic kernel on* Q *given* $\mathcal{S}$*, which assigns to each* $s \in \mathcal{S}$ *a probability distribution,* $T_q(\cdot|s)$*, over* Q*;*

- $R : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{S} \times Q \to [0,1]$ *is a conditional stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $\mathcal{S} \times Q$*, that assigns to each* $s \in S$ *and* $q' \in Q$*, a probability measure,* $R(\cdot|s,q')$*,on the Borel space* $(\mathbb{R}^{n(q')}, \mathcal{B}(\mathbb{R}^{n(q')})$*. The function* $R(A|(q,\cdot),q')$ *is assumed to be Borel measurable for all* $q, q' \in Q$ *and all* $A \in \mathcal{B}(\mathbb{R}^{n(q')})$*.*

Practically speaking, the discrete component takes values in a finite set Q of modes; each mode is equipped with a continuous domain (typically $\mathcal{S} = \mathbb{R}^n$). A point d of the hybrid space is thus composed of the pair $d = (q, s)$, where $q \in Q$ and $s \in \mathcal{S}$. Point d belongs to the hybrid space $\mathbf{D} = Q \times \mathcal{S}$. Discrete-time transitions among states abide by the following conventions. Given the initial point $d_0 = (q_0, s_0)$, the evolution of $s(k)$ follows the stochastic dynamics corresponding to the mode q_0. The continuous variable $s(k)$ maintains the same dynamics governed by T_x. At each time instant, the state d may jump to a new mode q' according to the stochastic kernel T_q. Whenever the jump occurs, the state d is "transported" into a new state $d' = (q', s')$ following the reset map $R(\cdot|q, s, q')$. Variable s will then evolve according to the dynamics of q' until a new jump occurs.

Formally, we may define $B(k)$ as a standard stochastic motion in $\mathbb{R}$ at time k. A typical stochastic execution starts from (q_0, s_0) and the continuous state $s(k)$ evolves according to the stochastic difference equation (SDE)

$$s(k+1) = f(q_0, s(k)) + g(q_0, s(k))B(k), \qquad s(0) = s_0, \tag{2.2}$$

until s jumps into a new mode q'. We will use this formalism to describe the stochastic dynamics of the network frequency and of the PV systems' power output in Chapters 5, 6, 7.

2.2.3 Probabilistic Models

Real-life systems often present complex, hybrid and stochastic behaviours. Photovoltaic systems, as an example, activate or deactivate their power output in response

to the grid frequency value (see Chapter 4): this type of switching behaviour denotes a hybrid system. Further, the noisy frequency measurements can be modelled as a stochastic signal; thus, the interconnection between the power grid and PV systems represents a stochastic hybrid system.

To reason over complex systems, oftentimes we need an *abstraction* of the system that would yield a simpler model, over which analysis and computation can be performed. Different abstraction models can be used to comprehensively represent a system: a common, yet effective, modelling framework is represented by Markov models. Markov models are stochastic state models that exhibit the Markov property: given any current state, the probability distribution describing the next state is independent of previous states. In this thesis, we use Markov models to refer to discrete-time Markov chains (dtMCs) [66–68]. We use the notation from the verification community – in particular, from [69].

Definition 2 (Discrete Time Markov chain (dtMC)) *A discrete time Markov chain is a tuple defined by* $\mathscr{M} = (\mathcal{S}, P, s_0)$ *where*

- $\mathcal{S}$ *is a finite set of states,*
- $P : \mathcal{S} \times \mathcal{S} \to [0, 1]$ *is a discrete transition probability function that assigns, to each* $s \in \mathcal{S}$*, a probability distribution over* $\mathcal{S} : P(s, \cdot)$*; further, for all states* s*,*

$$\sum_{s' \in \mathcal{S}} P(s, s') = 1,$$

- s_0 *is an initial state.*

Labels can be introduced to describe states: as an example, a *target* state may represent a desirable model condition as drawing a 6 out of a die, or an *unsafe* state may encode a bad, undesirable condition. To this end, two additional elements can be added to a dtMC to form a *labelled* dtMC:

Definition 3 (Labelled dtMC) *A labelled dtMC is a Markov chain* $\mathscr{M} = (\mathcal{S}, P, s_0, AP, L)$ *where*

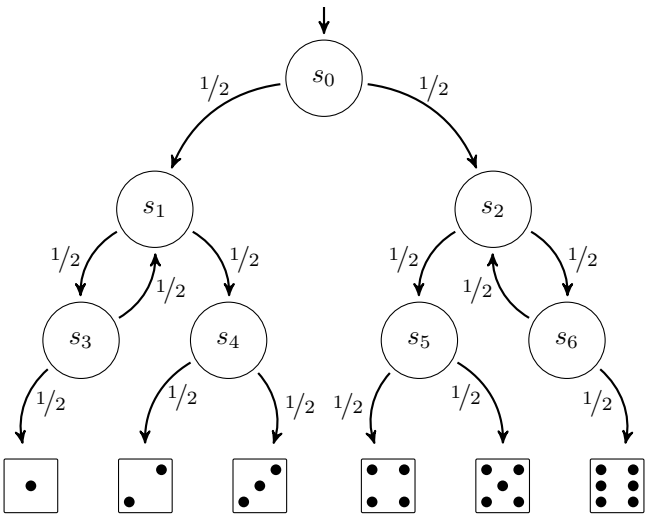

Figure 2.5: Knuth-Yao die.

- *AP is finite set of atomic propositions;*
- *$L : \mathcal{S} \to 2^{AP}$ is a labelling function that maps each state with a label;*

so that (possibly) all states have a label drawn from the alphabet 2^{AP}.

Figure 2.5 depicts the famous Knuth-Yao algorithm [70] for simulating a six-sided die by repeatedly tossing a fair coin. The MC consists of 13 states; the set AP is composed by the natural numbers from 1 to 6. The lower six states are marked by their semantics: the outcome of a die. States $s_0, \ldots, s_6$ are transitory, thus have no practical meaning: they have an empty label. The algorithm works as follows. We start in the initial state s_0 and flip (represented by the circular states) fair coins. If we flip heads, we take the leftmost transition, otherwise, we take the rightmost transition, until we reach one of the states ⚀ – ⚅.

Labels are useful from an intuition perspective, to give a practical meaning to abstract states, as per the final 6 states of the Knuth-Yao algorithm. Furthermore, labels help us ask quantitative and qualitative questions about a specific model, and more generally to develop model checking techniques.

Parametric Models

In practice, precise models that exactly describe a system, possibly under several different environments, are hard (if at all possible) to develop. The parametric

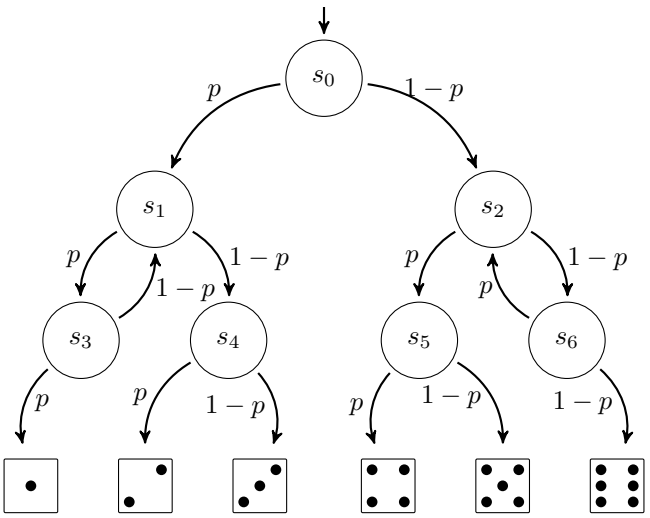

Figure 2.6: Knuth-Yao parametric die.

framework rather succinctly describes uncountably many models: one is often not only interested in verifying a single (known) model, but rather in deriving (any or all) model parameters ensuring the validity of a property under study [71–73]. In particular, [71] presents one of the first explicit approaches for parametric model checking of discrete-time Markov Chains (dtMCs): the probability of a path is computed symbolically by translating the dtMC into a finite automaton and obtaining a regular expression from it. Conditions for the satisfaction of a given property for a parametric dtMC, with applications to fair communication protocols, are outlined in [74]. [72] provides a more efficient computation of reachability properties for parametric dtMCs. The complexity of reachability properties for parametric Markov Decision Processes (MDPs) is discussed in [75]. Finally, an extensive illustration of methods for parametric Markov models is found in [76]. Typically, parameters are allowed to vary within predefined ranges: as said above, we do not treat them as inherent (adversarial) uncertainties in the model, but rather as viable values that can be selected. The goal is to find a subset (either a single instance, or the maximal subset) of the allowed parameter space, such that the satisfaction of a specification can be guaranteed.

Figure 2.6 shows a parametric representation of the Knuth-Yao die of Fig. 2.5. All the leftmost and righmost transitions now have probability p and $1-p$, respectively,

where $p \in [0, 1]$ is the probability of the outcome *heads* and $1 - p$ the probability of the outcome *tails*. The power of parametric model lies in their expressiveness: rather than reporting that the probability of reaching a certain state is exactly $1/6$, we are allowed to assert properties over a set of models. In practice, we hunt for instantiations of the model that satisfy a given property. Two canonical questions are

- Feasibility: does it exist an evaluation of the parameters (possibly within a region) such that the instantiated model satisfies the property?
- Verification: do all evaluations of the parameters (within a region) instantiate a model that satisfies the property?

Applications of parameter synthesis techniques span over several fields, especially chemical networks [77] and biological systems [78]. In particular, [79] presents an approach for tuning gene network parameters to verify a given specification. Beyond the life sciences and within engineering, [80] studies a parameter synthesis problem for optimal energy consumption with a combination of discrete and continuous parameters; [14] embeds the heterogeneity (in terms of reconnection and disconnection patterns) of a population of renewable energy sources as parameters in the model of an energy network, which is then investigated under safety requirements (lack of blackouts) [19] – the very same practical spirit drives this dissertation.

2.2.4 Model Checking

Assuming that our developed model is an appropriate representation of a system of interest, we may use the model to check whether the system satisfies a given requirement. Notably, we would like to automate it, as reasoning about probabilistic systems is all but trivial, as famously shown by the brain teasers: the Monty Hall problem[3], the Three Prisoner's problem[4], the Bertrand's box paradox[5]. The automatic check is schematically depicted in Fig. 2.7, and known as model checking

[3] `https://en.wikipedia.org/wiki/Monty_Hall_problem`
[4] `https://en.wikipedia.org/wiki/Three_Prisoners_problem`
[5] `https://en.wikipedia.org/wiki/Bertrand's_box_paradox`

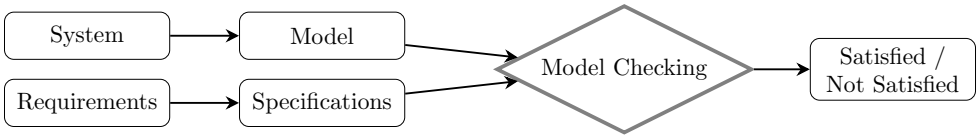

Figure 2.7: Model checking scheme.

[69, 81, 82]: "Model checking is an automated technique that, given a finite-state model of a system and a formal property, systematically checks whether this property holds for that model" [69]; when the underlying models are Markov models, we talk about *probabilistic* model checking. Thanks to its reliability and mathematical guarantees, over the last decade (probabilistic) model checking has proved to be an incredible success for the diffusion of formal methods in several scientific communities, from engineering to automated coding.

Model checking comprises a specification of interest, which is usually encoded in some modal logic (e.g. LTL, CTL, PCTL), and the system under study expressed as a formal model, commonly either a transition model [69], a probabilistic model [83], a Petri net [84], a timed [85] or hybrid automaton [86], or via a process algebra [87]. Stochastic models in particular describe systems involving random phenomena and probabilistic behaviours.

Probabilistic model checking [88–90] is a formal technique for Markov models which combines correctness guarantees with a quantitative performance evaluation. The success of probabilistic model checking comes from the fact that it is based on a systematic exploration of all possible states (hence the correctness guarantee) and it can be used to analyse a wide range of quantitative properties of the original system, relating for example to its performance or reliability. In contrast to simulation-based techniques, which generate approximate results by averaging simulation outcomes from a large number of random samples, probabilistic model checking applies numerical computation to yield exact results.

Formal verification of stochastic models is an active area of investigation, which has witnessed the development of quite a few software tools. Prism [83] and

Storm [91] are probabilistic model checkers covering a large variety of finite-state models, supporting a variety of numerical and symbolic solvers, as well as multi-objective model checking and modern parameter synthesis techniques. For an extensive summary of both features and performance of state-of-the-art probabilistic model checkers, the interested reader may refer to [92]. Towards models with uncountably many states, FAUST^2 [93] provides formal abstractions, whereas Modest [94] and StocHy [95] combine nondeterministic choices, continuous system dynamics, and stochastic decisions.

Probabilistic model checking is a formal analysis technique that provides an exact, quantitative performance evaluation of a Markov model.

Let us now fix a labelled dtMC $\mathscr{M} = (\mathcal{S}, s_0, P, AP, L)$ with *target* states $\mathcal{T} \subseteq \mathcal{S}$, and assume we are interested in asking questions like: "What is the probability to reach a specified set of target states?" or "What is the probability to reach the target states within h steps?".

Remark 1 *We are never interested in what happens once we have reached a target state, hence we render target states $t \in \mathcal{T}$ absorbing: with a unique transition $P(t,t) = 1$.* □

Paths and Probability Measures We now define the notions of paths and probability measure for reachability probabilities. A path of a (labelled) dtMC $\mathscr{M}$ is an (in)finite sequence $\pi : s_0 \to s_1 \to \dots$, where $s_i \in \mathcal{S}$ and $\mathsf{P}(s_i, s_{i+1}) \neq 0$ $\forall i \in \mathbb{N}$. A probability measure $\mathsf{P}^{\mathscr{M}}$ for finite paths $\pi = s_0 s_1 \dots s_n$ is given by the product of transition probabilities:

$$\mathsf{P}^{\mathscr{M}}(\pi) = \prod_{i=0}^{n-1} \mathsf{P}(s_i, s_{i+1}). \tag{2.3}$$

A reachability probability, as the name suggests, denotes the probability of reaching a given (target) set $\mathcal{T}$. We might be interested in a *bounded* reachability, i.e. being able of reaching the target within a finite number of steps, or unbounded reachability, namely being able to reach the target at all. Within the formal methods

community, the reachability of $\mathcal{T}$ is denoted as $\Diamond\mathcal{T}$ and the bounded reachability as $\Diamond^{\leq h}\mathcal{T}$, where h is the number of steps.

Definition 4 (Reachability probability) *Let $\mathscr{M}$ be a dtMC with target states $\mathcal{T}$. We define the reachability probability $\mathsf{P}_{\mathscr{M}}(s \models \Diamond\mathcal{T})$ for reaching $\mathcal{T}$ from state s*

$$\mathsf{P}_{\mathscr{M}}(s \models \Diamond\mathcal{T}) = \sum_{\pi \in P(s,\Diamond T)} \mathsf{P}^{\mathscr{M}}(\pi),$$

and the bounded reachability probability $\mathsf{P}_{\mathscr{M}}(s \models \Diamond^{\leq h}\mathcal{T})$ for reaching $\mathcal{T}$ from state s within $h \in \mathbb{N}$ steps

$$\mathsf{P}_{\mathscr{M}}(s \models \Diamond\mathcal{T}) = \sum_{\pi \in P(s,\Diamond^{\leq h} T)} \mathsf{P}^{\mathscr{M}}(\pi).$$

We refer to h as the horizon. We define the (bounded) reachability probability $\mathsf{P}_{\mathscr{M}}(s \models \Diamond\mathcal{T})$ (resp. $\mathsf{P}_{\mathscr{M}}(s \models \Diamond^{\leq h}\mathcal{T})$) for reaching $\mathcal{T}$ in $\mathscr{M}$ (in h steps) as reachability probability from the initial state, i.e.,

$$\mathsf{P}_{\mathscr{M}}(\Diamond\mathcal{T}) := \mathsf{P}_{\mathscr{M}}(s_0 \models \Diamond\mathcal{T}) \text{ and } \mathsf{P}_{\mathscr{M}}(\Diamond^{\leq h}\mathcal{T}) := \mathsf{P}_{\mathscr{M}}(s_0 \models \Diamond^{\leq h}\mathcal{T}).$$

Let $s \in \mathcal{S}$ with $p := \mathsf{P}_{\mathscr{M}}(s \models \Diamond\mathcal{T})$: we say that a state s has probability p to reach the target $\mathcal{T}$.

When the target state represents particularly beneficial conditions for the system, e.g. stability for a dynamical system, network frequency close to the reference value for a power grid, we might call $\mathcal{T}$ a *safe* set. In this case, we talk about a safety probability.

Properties Properties evaluate MC by comparing some measure to a threshold $\kappa \in \mathbb{Q}_{\geq 0}$ using a threshold comparison relation $\bowtie \in \{<, \leq, >, \geq\}$. We thus turn from the quantitative view to a Boolean view. For each property ϕ, we define the satisfaction relation $\models$:

Definition 5 ((Un)Bounded Reachability Property) *Let $\mathscr{M}$ be a dtMC with target states $\mathcal{T}$, κ a threshold and $\bowtie$ a comparison relation. We define the following properties:*

- *unbounded reachability property* $\mathsf{P}_{\bowtie\kappa}(\Diamond\mathcal{T})$*:*
  $\mathscr{M} \models \mathsf{P}_{\bowtie\kappa}(\Diamond\mathcal{T}) \Leftrightarrow \mathsf{P}_{\mathscr{M}}(\Diamond\mathcal{T}) \bowtie \kappa$.

- *bounded reachability property* $\mathsf{P}_{\bowtie\kappa}(\Diamond^{\leq h}\mathcal{T})$*:*
  $\mathscr{M} \models \mathsf{P}_{\bowtie\kappa}(\Diamond^{\leq h}\mathcal{T}) \Leftrightarrow \mathsf{P}_{\mathscr{M}}(\Diamond^{\leq h}\mathcal{T}) \bowtie \kappa$.

When a property is upper bounded, i.e. $\bowtie \in \{<, \leq\}$ is called *safety* property; on the contrary, when a property is lower bounded, i.e. $\bowtie \in \{>, \geq\}$, is called *liveness* property.

Whilst in this work we limit ourselves to the study of safety/reachability properties, evidently in formal verification much more complex specifications are allowed. These notions may be extended to richer logics like arbitrary PCTL [96], PCTL* [97], or ω-regular properties [98]. For these logics, reachability is considered an elementary property.

2.2.5 Formal Abstractions of Markov Processes

The investigation of heterogeneous, stochastic, continuous models represents a challenging task for scientists, both analytically and computationally – think of the design of optimal controllers [99] or the reachability analysis [100], as an example. Notions of language equivalence and bisimulation [101] present a correspondence between a complex model and a simpler, smaller one. However, the requirements to prove equivalence between models are quite restrictive: generally we expect perfect correspondence between the two models' trajectories. Unlike bisimilarity, approximate abstraction is a rather novel and computationally-friendly approach which generates a model with a smaller, generally finite, state space that is approximately equivalent to the original, complex system [102]. These approximate abstractions are equipped with a metric that quantifies the distance between the trajectories of the abstract and original systems [102–104].

The purpose of approximate or exact abstractions of dynamical systems is twofold. We first need to prove the existence and well-foundedness of a finite abstraction [101], to further develop finite-time, tunable abstraction algorithms [103].

Abstraction techniques have been applied to probabilistic models: for instance to discrete-space, continuous-time models [105]. Arguably, the largest accomplishment of bisimulation – and of the corresponding abstract models – regards discrete Markov processes (MPs) [106–108], and jump-linear stochastic systems in [104]. Further, [109] introduces a technique to approximately abstract continuous-time probabilistic models as locally-consistent Markov chains; this approach is then developed and adapted to hybrid models [110, 111]. However, these approaches do not derive any certificate on the reliability of the abstraction: in other words, they do not provide any explicit approximation bound.

In recent years, the method of formal abstraction [65, 112, 113] has arisen as a reliable procedure that generates tunable abstractions of continuous, stochastic models as MPs and SHSs, while providing an explicit bound on the offered abstracted model. It is aimed at reducing a discrete-time, uncountable state-space Markov process to a discrete-time finite-state Markov model, for the purpose of probabilistic model checking. This method, exploiting notions of distance between probability distributions, represents a general procedure to approximate probabilistic models [106–108], .

As outlined in Section 2.2.2, we recall that SHSs evolve over uncountable spaces with discrete jumps. Directly and blindly applying formal verification to such complex models is in general not decidable [65, 114]. Quantitative finite abstractions as formal abstractions may be applied to SHSs in order to overcome this issue. These accurate approximations come in two main different flavours: abstractions into dtMC or MDP [65], [93] and, recently, into interval MDP (IMDP) [95, 115, 116]. Notice that both methods require the definition of a compact set over which analysis is performed, as a consequence of partitioning the state-space. Once the finite, abstract models are provided, we may formally verify them attaching an approximation error. The formal abstractions methodology has the following features:

- it provides a procedure to construct an approximate abstraction of a stochastic model via the partition of the state-space and the approximation of the transition stochastic kernel, with an explicit computation of the error;

- the abstraction results in a Markov model, which is a useful and commonly used framework to closely represent probabilistic systems;
- under rather general continuity assumptions, the procedure offers an explicit and tunable bound on the error between the probability distribution of the abstracted model and that of the original model, for each time instant. Notably, refining the abstraction guarantees a tighter error;
- finally, it provides an algorithm to compute the abstraction with desired precision, based on the derived error bound.

Further, this technique enables the verification of probabilistic properties of the original system on a finite-dimensional Markov abstracted model with arbitrary precision.

In the following, we briefly outline formal abstraction procedures for SDE [93], which we newly extend to classes of parametric models in Chapter 7.

A Formal Abstraction Procedure Let us consider a discrete-time stochastic differential equation as follows:

$$x(k+1) = f(x(k)) + \omega(k), \tag{2.4}$$

where $k \in \mathbb{Z}$ indexes the discrete time, $x \in \mathbb{R}^n$ represents the state space variable, f is a continuous vector field $f : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$, and ω denotes an additive stochastic term.

The model in Eq. (2.4) can be conceived as a discrete-time Markov Process (dtMP), defined over a pair $(\mathcal{Q}, T_w)$, where $\mathcal{Q}$ is a properly measurable continuous state space [112], and T_w is a stochastic kernel that assigns to each point $q \in \mathcal{Q}$ a probability measure $T_w(\cdot \,|\, q)$, so that for any properly measurable set $\mathcal{A}$

$$\mathsf{P}_q(\mathcal{A}) = \int_{\mathcal{A}} T_w(dq \,|\, q), \tag{2.5}$$

where P_q denotes the conditional probability $\mathsf{P}(\cdot \,|\, q)$. Notice that T_w represents a *continuous* kernel, i.e. takes value over uncountably many states q.

The basic procedure to abstract a dtMP $\mathscr{F} = (\mathcal{Q}, T_w)$ as a finite-state dtMC $\mathscr{M} = (\mathcal{Z}, P)$ is as follows [65, 93]. $\mathcal{Z} = \{z_1, z_2, ..., z_p\}$ is a finite set of abstract states, and $P : \mathcal{Z} \times \mathcal{Z} \to [0, 1]$ is a transition probability matrix over $\mathcal{Z}$. We denote $P(z, z')$ to be the probability of transitioning from state z to state z'. The finite state space $\mathcal{Z}$ is constructed by partitioning the state space of $\mathscr{F}$, and selecting representative points r_i in each partition to make up the states in $\mathcal{Z}$. The probability of transitioning from the (abstract) state z to state z', namely $P(z, z')$, is computed by marginalising the stochastic kernel T_w, conditional on z, over the partition corresponding to z'.

Assume now that the kernel T_w can be expressed [112] via its density function t_w

$$t_w : \mathcal{Q} \times \mathcal{Q} \to \mathbb{R} \geq 0,$$

namely $T_w(dq' \,|\, q) = t_w(q' \,|\, q)\, dq'$ for any $q, q' \in \mathcal{Q}$. The next-step error [65] incurred by the abstraction is shown to depend on the regularity of function t_w: assuming that t_w is Lipschitz continuous, namely there exists l_w such that

$$|t_w(\bar{q} \,|\, q) - t_w(\bar{q} \,|\, q')| \leq l_w ||q - q'||, \quad \forall q,\, q',\, \bar{q} \in \mathcal{Q},$$

where l_w is a positive constant, then the next-step error is bounded from above as

$$\epsilon_{abs} \leq l_w\, \delta_q\, \mathscr{L}(\mathcal{Q}), \tag{2.6}$$

where δ_q is the max diameter of the introduced partitions and $\mathscr{L}(\mathcal{Q})$ depends on the volume of the partitions [65]. Furthermore, the h-step error is upper bounded by

$$\epsilon_{abs} \leq h\, l_w\, \delta_q\, \mathscr{L}(\mathcal{Q}). \tag{2.7}$$

This formulation allows tuning the partition size δ_q to ensure that the abstraction error is smaller than a desired, user-provided threshold ϵ^{max}, as follows:

$$\delta_q \leq \frac{\epsilon^{max}}{h l_w \mathscr{L}(\mathcal{Q})}. \tag{2.8}$$

A smaller δ_q can thus be selected to ensure higher accuracy, but can result in a larger finite state space $\mathcal{Z}$. In the additive noise of Eq. (2.4), the constant l_w depends on the variance of the noise w, and more generally it is related to the Lipschitz constant of the kernel density t_w [65, 93].

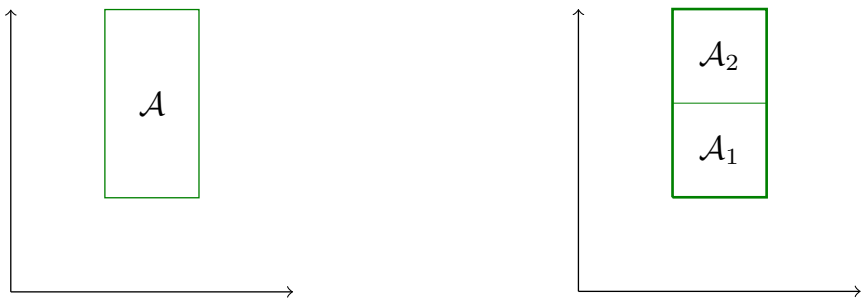

(a) **(b)**

Figure 2.8: Depiction of a safe set $\mathcal{A}$ (left) and state space partitioning (right).

Formal abstractions allow the translation of a Markov process into a Markov chain, attaching a tunable precision error. The procedure entails a state space partitioning procedure and a marginalisation of the stochastic kernel to compute the transition probabilities among the partitions.

Example 1 *Consider a discrete time system of two variables*

$$x(k+1) = Ax(k) + \Sigma\omega(k), \tag{2.9}$$

where $A = I_2$ *and* $\Sigma = I_2$*, where* I_2 *is the two-dimensional identity matrix, and* ω *represents a Gaussian process* $\mathcal{N}(0,1)$*. We first notice that* $x(k+1) \simeq \mathcal{N}(Ax(k), \Sigma)$*, hence the value of* x *at the next time step belongs to a Gaussian distribution with average* $Ax(k)$ *and variance* Σ*. Let us assume a rectangular safe set* $\mathcal{A}$ *defined as* $[1,2] \times [1,3]$*, i.e.* $1 \leq x \leq 2$ *and* $1 \leq y \leq 3$*, shown in Fig. 2.8 (right). Assume also the initial state of the system is within* $\mathcal{A}$*, located at* $x_0 = (1.5, 1.5)$*.*

In this example we will compute the probability of remaining within $\mathcal{A}$ *for* $h = 2$ *time steps.*

Let us divide $\mathcal{A}$ *into two partitions,* $\mathcal{A}_1 = [1,2] \times [1,2)$ *and* $\mathcal{A}_2 = [1,2] \times [2,3]$ *– see Fig. 2.8. Each partition is abstracted into a state of a Markov chain:* z_1 *and* z_2 *respectively. The remaining state space,* $\mathbb{R}^2 \setminus \mathcal{A}$*, is abstracted as the unsafe state* z_3*.*

To compute transitions between states, we select a reference point for each of the safe partitions, $r_1 = (1.5, 1.5)$ *and* $r_2 = (1.5, 2.5)$ *respectively, shown in Fig. 2.10 (left). The reference points are used as the center, i.e. the average value, of the*

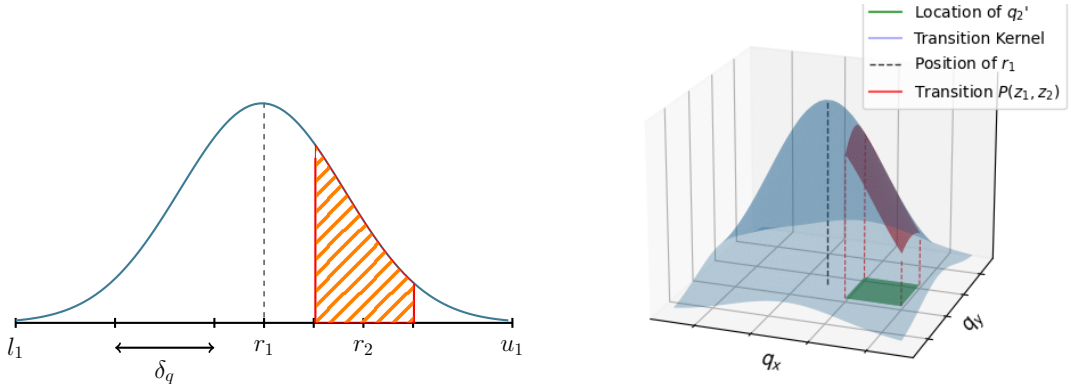

(a) depiction of $P(z_1, z_2)$ in a one-dimensional system.

(b) Depiction of $P(z_1, z_2)$ in a two-dimensional system.

Figure 2.9: Depiction of the transition probability $P(z_1, z_2)$.

Gaussian distribution. Let us calculate $\mathsf{P}(z_1, z_2)$*, the transition probability from state* z_1 *to the next state* z_2*, corresponding to the transition from region* $\mathcal{A}_1$ *to region* $\mathcal{A}_2$*, as*

$$\mathsf{P}(z_1, z_2) = \int_{\mathcal{A}_2} p_{\mathcal{N}}(u \mid r_1) du, \tag{2.10}$$

where $p_{\mathcal{N}}$ *represents the Gaussian probability distribution function. Explicitly, Eq.* (2.10) *becomes*

$$\mathsf{P}(z_1, z_2) = \frac{1}{2\pi} \int_1^2 \int_2^3 \exp\left(-\frac{1}{2}(u - r_1)^T \Sigma^{-1} (u - r_1)\right) du_1 du_2. \tag{2.11}$$

Similarly we compute $\mathsf{P}(z_1, z_1)$*, the self loop transition probability, corresponding to the transition from region* $\mathcal{A}_1$ *to region* $\mathcal{A}_1$*, as*

$$\mathsf{P}(z_1, z_1) = \frac{1}{2\pi} \int_1^2 \int_1^2 \exp\left(-\frac{1}{2}(u - r_1)^T \Sigma^{-1} (u - r_1)\right) du_1 du_2. \tag{2.12}$$

Finally, to compute transition $\mathsf{P}(z_1, z_3)$ *we exploit the fact that the sum of all transitions starting from a given state is equal to 1. Hence,*

$$\mathsf{P}(z_1, z_3) = 1 - \mathsf{P}(z_1, z_1) - \mathsf{P}(z_1, z_2). \tag{2.13}$$

We compute transitions from state z_2 *by swapping* r_2 *into the integral computation. Figure 2.10 (right) schematically depicts the transitions. Notice that in Fig. 2.10 we mark the state space outside* $\mathcal{A}$ *as unsafe and symbolically denote it with a red*

Figure 2.10: Construction of the Markov chain: selection of the reference points (left), computation of transition probabilities (right).

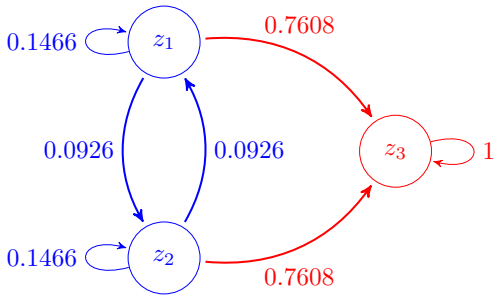

Figure 2.11: The final Markov chain: state representing $\mathcal{A}$ are depicted in blue, the unsafe state in red.

cross. A pictorial representation of the transitions' computation is reported in Fig. 2.9, both for a one-dimensional system (left) and a two-dimensional system (right).

The unsafe state z_3 is made absorbing, i.e. with a probability-one self loop. This represents a standard procedure, indicating that the system's safety is violated once it gets to an unsafe state.

The resulting Markov chain, represented in Fig. 2.11, offers the transition probability matrix

$$P = \begin{bmatrix} 0.1466 & 0.0926 & 0.7608 \\ 0.0926 & 0.1466 & 0.7608 \\ 0 & 0 & 1 \end{bmatrix}. \tag{2.14}$$

Recall that we aim at computing the probability of remaining within $\mathcal{A}$ over a time horizon of $h = 2$ steps starting from $x_0 = (1.5, 1.5)$. We rewrite the initial condition as a probability vector: x_0 belongs to $\mathcal{A}_1$, thus we have probability 1 of being in $\mathcal{A}_1$ at time 0, and probability 0 of being in $\mathcal{A}_2$ or outside $\mathcal{A}$. We write the initial probability vector as

$$v_0 = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}.$$

We compute the probability vector after two steps by multiplying the transition matrix twice, as

$$v_2 = v_0 \cdot P^2 = \begin{bmatrix} 0.0301 & 0.0272 & 0.9427 \end{bmatrix}, \tag{2.15}$$

thus the probability of remaining within $\mathcal{A}$ *is* $0.0301 + 0.0272 = 0.0573$*. Note that the probability of reaching the unsafe state within two steps is* $1 - 0.0573 = 0.9427$*.*

The abstraction procedure introduces an error in view of the partitioning procedure and the reference point selection. The abstraction error ϵ_{abs} *is upper bounded [65] as in Eq. (2.7). This illustrative example presents* $\mathscr{L}(\mathcal{A}) = 2$*, the partitions size* δ_A *is the diagonal of the partitions, hence* $\delta_A = \sqrt{2}$*, whereas the time horizon is* $h = 2$*. The Lipschitz constant can be shown [117] to be* $l_w = \dfrac{1}{2\pi} \simeq 0.15$*. The overall abstraction error is upper bounded by*

$$\epsilon_{abs} \leq 0.6\sqrt{2} \simeq 0.84. \tag{2.16}$$

The abstraction error should be limited to values below one, the smaller the better. In order to reduce ϵ_{abs} *the formal abstraction procedure can be refined. We might increase the number of partitions, i.e. of states of the Markov chain, by reducing the partition diameter* δ_A*. Let us consider a desired maximum abstraction error* $\epsilon^*_{abs} = 0.01$*. Given the same system characteristics, we set the equality*

$$\epsilon^*_{abs} = h\, l_w\, \delta_A\, \mathscr{L}(\mathcal{A}) = 0.01 \Longrightarrow \delta_A = 0.167. \tag{2.17}$$

The value of δ_A *is defined as the diameter of the partition, hence we need squares with sides* $^{0.167}/_{\sqrt{2}} \simeq 0.1$ *to fulfil the requirement of* $\epsilon_{abs} \leq 0.01$*. The corresponding Markov chain is composed of (at least) 200 states.* □

We offer an insight on the computation of the transition probabilities in the following. We choose an arbitrary reference point r_i within each (hyper-)rectangle $\mathcal{A}_i$, e.g. its centre of mass. Consider, as an example, the transition from region $\mathcal{A}_1$ to region $\mathcal{A}_2$, that is, from state z_1 to state z_2. Imagine to centre the stochastic kernel over point r_1, as per Fig. 2.9. The transition probability between z_1 and z_2 is obtained as the integral of the density of the noise kernel, centred at r_1, over region $\mathcal{A}_2$ (the area corresponding to z_2).

An extension of this methodology is represented by adaptive and sequential partitioning [93, 118]. The procedure first constructs a coarse grid, which provides an abstraction error greater than the desired value, usually corresponding to $10\epsilon^*_{abs}$. Then, a *local* abstraction error ϵ_l is computed for each partition. If a partition provides a local error $\epsilon_l > \epsilon^*_{abs}$, the procedure splits it into smaller cells. This procedure is repeated until each cell achieves a local abstraction error $\epsilon_l \leq \epsilon^*_{abs}$ across all cells.

Uniform abstractions and adaptive abstractions represent an intuitive trade-off: the first method requires more memory, in view of the generally larger state-space, but less time, whereas the second method generates abstractions with a smaller state-space that require more time to be computed.

2.2.6 Formal Methods in Power Systems

Renewable power has significantly increased its exposure over the last decade and their presence in larger and smaller power systems is becoming preponderant. On the economic side, the energy market accommodates more and more solar, wind, thermal power sources [2] while the cost of producing electricity from renewables is lower than the cheapest fossil fuel-based option [119] – tables have turned! The future will see the absolute majority of electricity sourced from renewables: the ENTSO-E foresees that within the 28 countries of the EU, electricity from renewable sources meets up to 63% of power demand in 2030 and 83% in 2040 [5].

From the electrical connection point of view, renewable energy producers are starting to inject power into the grid and to participate in the regulatory action [120]. An extraordinary example is Ireland, where wind generation reaches peaks of 84% of the demand, with a total of more than 30% electricity supplied in 2019 [121, 122]. However, the intrinsic unpredictability of RESs, connected mainly to weather conditions or seasonal effects, remains an obstacle to the full deployment of this technology. Just to name a few, the integration of renewables into the power grid has impact on optimum power flow, power quality, voltage and frequency control and system economics. An overview of the issues and main challenges

on frequency regulation with the integration of renewables can be found in [43, 123]. Further, renewables modify the electric grid characteristics in view of their reduced inertia. As a consequence, the response to oscillations or incidents differs considerably (see Chapter 5 for a full analysis). Formal methods provide tools to model and certify safety requirements over complex systems, as the electricity network. Formal verification techniques are indispensable and highly recommended for the development of safe and reliable power systems. Conventional validation methods, i.e. testing and simulations, have limitations in terms of assessing the reliability of complex systems due to both stochastic and non-linear dynamical features. Whilst the literature on formal analysis techniques that support verification of safety and security constraints is rich and vast, the application to identify the potential impacts of security related threats on safety properties of power systems is much less developed.

In the following, we review several works tackling models for PV systems or the strictly related literature on TCLs. We schematically report the literary efforts in Table 2.1. In particular, we highlight the analysis approach (whether by machine learning (ML) techniques, non-linear models (NL), Markov chains (MC), Markov decision proccesses (MDP), mean field analysis (MF), abstractions (abstr.), model-predictive control (MPC) or linear time invariant (LTI) models. The column 'System' denoted whether the work considers PV systems (PV), energy storage systems (ESS), thermostatically controlled loads (TCL) or common smart loads (L); when the population is heterogeneous, we use the prefix h- before the system's acronym. We further note the size of the population and the time frame used for each simulation. Finally, we report whether the submission tackles frequency deviations or jumps after a contingency, as done in this work.

Models for the Aggregation of PV Systems Models of a population of PV systems are widespread across several scientific domains, used for economic, optimisation, game-theoretic research. A dedicated modelling framework is presented in [124], where the authors offer models for (thermal) solar plants, boilers, turbines

and power generators. A novel control design is proposed to reduce generation costs and maintenance. In view of the different time constants, the work simulates various time scales, from milliseconds to minutes. The nonlinear model of PV system is used to represent a solar plant composed of 1500 devices.

High-level techniques as machine learning-based estimation may also be used to derive models of solar applications. As an example, [125] uses machine learning techniques (random forests regression) to estimate in real time the grid inertia. Assuming a proportional control (as per Chapter 5) the algorithm computes the minimum headroom needed to support the grid in case of contingency. The time frame is few seconds, and the approach is model-free as the power output of a single PV system is (assumed) known at each time instant. Model-free approaches are further used to maximise the power output of grid-connected solar devices, as in [126], exploiting identification techniques.

Similarly, [127] proposes an adaptive power point tracking method to provide frequency regulation on an AC micro-grid. The authors use the equivalent circuit model of 6 PV panels embedded in a micro-grid to simulate the response after a frequency jump. A population-based optimisation method may be instead used to overcome generation sub-optimal results in presence of partial shading: [128] simulates the results of several algorithms in a population of 100 devices to achieve the yearly maximum power output. Equivalent circuit models for 15 PV systems are used also in [129], along with 15 battery systems, where a fuzzy control logic is proposed to reduce frequency oscillations in a smart grid.

Models for a population of PV systems may also be found within the voltage regulation literature. Game-theoretic approaches are frequently used for this purpose: [130] uses a multi-objective optimisation algorithm to minimise phase and voltage imbalances, simulating the IEEE 123 test feeder with 75 PV systems. Closer to the scopes of this dissertation, a simulation of under- and over-frequency contingencies of the Hawaiian grid is presented in [131], where the software PSS/E is used to model the grid and the PV systems, which account for 40% of the power demand. In particular, the authors compare the frequency response with and without a

power droop control to counterbalance incidents. Several tools exist for research on power grids, providing models of renewable energy sources and numerical analysis, as HOMER [132], PSS/E [133], Matlab [134].

A part of literature is dedicated to the estimation of the presence of PV systems from power consumption data: e.g. via fuzzy logic approach [135] or via machine learning [136, 137]. This interesting research may be used in conjunction with the work in this dissertation: once the population of PV systems is finely estimated, we may perform more accurate tests on the safety of the grid, exploiting the real-world measurements.

Markov models of Smart Loads The demand-response literature is one of the first research paths to have introduced Markov models for loads [55, 56, 117, 138–140]. Commonly, this literature considers a population of loads, often thermostatically controlled loads (TCLs): such devices are fridge/freezers, coolers or air conditioning systems. We recall several examples of control of a TCLs population for flexible demand-response. The work in [138] presents a stochastic control solution to manage the power consumption of domestic refrigerators in response to frequency deviations and sudden power plants outages, highlighting the problem of devices' synchronisation. In particular, it considers a homogeneous population ($N = 10^3$) of loads using an LTI model framework with a ctMC-based scheduling. Building on top of this work, [55, 56] propose a power consumption scheme based on the frequency deviation, testing the overall response against ENTSO-E requirements. The authors model a *heterogeneous* population ($N = 10^3$) of devices, in terms of initial state, temperature setpoints, power consumption.

The works in [141–143] present an analysis of power grids with a substantial portion of photovoltaic micro-generation, while testing a proportional control technique. The authors study the proposed regulations, comparing them with alternative approaches inspired by communication protocol designs, using the verification tool Modest [94]. They, however, do not offer an aggregation framework, but rather a 600-second simulation considering 32 devices. The present thesis builds

Reference	Approach	System	Size	Time Frame	Contingency
[125]	ML	PV	1	seconds	✓
[127]	NL	PV	6	seconds	✓
[129]	NL	PV, ESS	15	hours	×
[130]	NL	PV	75	hours	×
[131]	NL	PV	40%†	seconds	✓
[124]	NL	PV	10^3	mins (ms)	×
[128]	NL	PV	10^3	hours	×
[141–143]	LTI	PV, TCL	32	minutes	[142]
[138]	LTI, ctMC	TCL	10^3	hours	✓
[55, 56]	LTI, ctMC	h-TCL	10^3	hours	[55]
[144, 145]	LTI	h-TCL	10^4	hours	×
[117, 139, 140]	MC, abstr.	h-TCL	10^3	hours	×
[146]	LTI	h-TCL	10^3	hours (5 mins)	×
[147, 148]	MC, MPC	h-TCL	10^3	hours (10 s)	×
[149]	MF	h-TCL	30	200 steps	×
[150, 151]	MC	TCL	10^4	hours (1 min)	×
[152]	LTI	TCL	10^2	hours (1 min)	×
[153]	MDP, MF	L	10^6	30 min (150 s)	×
this work	MC, abstr.	h-PV	var.	seconds (ms)	✓

Table 2.1: Literature review categorised on the used modelling approach, the application (including heterogeneity), the size of the population, the simulation time frame and the corresponding time step (where applicable). The column 'Contingency' indicates whether the work considers a sudden generation or load loss. This work considers a heterogeneous population of PV systems (h-PV) at various (10%, 20%, 40%) penetrations of network load in a time frame of few seconds.
†: the work considers a 1 GW load demand Hawaiian grid.

upon this study: we however propose different, data-driven models (see Chapter 4), a more complex and more realistic network model (Chapter 5) and finally we analyse the importance of disconnection and reconnection patterns.

The typical technique for controlling TCLs exploits a negligible shift in temperature setpoints. This however may lead to oscillations in power consumption of a heterogeneous population of devices, as highlighted in [144, 145]. The authors thus propose a *safe* protocol to actuate the shifting of the temperature setpoints implementing a relaxed hysteresis cycle. The population ($N = 10^4$) heterogeneity, in terms of room thermal capacitance C, power consumption P, and thermal resistance R, is encompassed using a lognormal distribution.

Several works [117, 139, 140, 147, 148] propose a Markov Chain (MC) approach

to tackle the *aggregation* of a homogeneous and heterogeneous (in terms of C, P, R) population ($N = 10^3$) of TCL. The temperature range is divided into multiple partitions (usually 10 to 80), each represented by a state in the MC. The transitions among states are determined by the TCL dynamical model, and the simulations include thousands of devices in a time frame of few hours. Whilst the authors in [117, 139, 140] choose an approach based on a formal abstraction, the works in [147, 148] exploit model-predictive control (MPC) techniques to control the power consumption of the population. In particular, the works [117, 139, 140] introduce formal methods, formal abstractions and verification to synthesise controllers and provide guarantees on the safety of a population of TCLs. Further, [139] proposes extensions to battery models and optimisation techniques for power consumption. The work on this thesis follows this path, providing Markov models, and formal guarantees to control a population of PV systems. Altogether, this work integrates into a vision of frequency regulation at a household level: TCLs on the power consumption side, PV systems on the power generation side.

Battery models, coupled with a population ($N = 100$) of homogeneous TCLs are presented in [152] to achieve load reduction following a priority-based control design. The overall analytical model encompasses the thermodynamical model of a room, the electrical model of a battery and of the TCL, along with a *virtual* energy storage system to prioritise the control actions.

The work [146] devise an LTI model to encompass the aggregation of heterogeneous devices ($N = 1000$) in a time frame of hours with time steps of 5 minutes. The authors devise an MPC approach to control the population energy need and, anticipating real-world applications, study the parameters estimations from energy consumption data.

Building upon the MC representation, the work in [153] proposes an MDP framework: practically speaking, general loads are equipped with actions (i.e. a discrete, quantised control). The authors propose an aggregated population ($N = 10^6$) model based on a MF approach, to evaluate the energy performance of a population of smart and deferrable loads (pools, in particular).

The aggregation of homogeneous and heterogeneous devices may exploit a mean-field (MF) approach. The work [149] proposes a mean field technique to model population of heterogeneous (C, P, R, initial conditions, temperature operation limits) TCLs, whose individual dynamics is represented by a linear model. A MF control is devised using a quadratic cost function to minimise the overall power consumption. In particular, this work focusses on 30 clusters of homogeneous devices.

Linear stochastic models are also used in [150, 151] to model a homogeneous population of TCLs. Each individual device is equipped with a MC-based randomised control to avoid synchronised switching. A heterogeneous (C, P, R, initial conditions, uniformly and Gaussian distributed) population ($N = 10^4$) is simulated against a homogeneous, aggregated analytical model showing the relation between the degree of heterogeneity and the aggregative model framework.

An alternative, simulation-based, verification technique to verify complex system is *statistical model checking* (SMC). By and large, SMC deduces whether a system satisfies a property by observing some executions with a monitoring procedure [154–156]. In contrast to a formal approach, SMC cannot guarantee a correct result, yet it is more scalable than the formal counterparts and it offers the possibility to bound the probability of making an error. These techniques have been extended to handle several kinds of properties and languages and a variety of tools have been developed in the recent years. A detailed outline of the SMC techniques is out of the scopes of this thesis and it can be found in [154, 157].

Statistical model checking approaches have been applied to grid-related safety guarantees, especially to smart grids and smart buildings, as an engineer-friendly bridge to verification. The work in [158] presents a framework for the UPPAAL tool [159, 160] for smart buildings to evaluate the performance of given control strategies varying environmental settings. [161] proposes an approach to reduce electricity demand peaks while enforcing distribution-network constraints: an SMC technique verifies that the probability of violating those constraints in non-nominal cases remains low. Hybrid systems are used to model smart grids in [162, 163],

along with a tailored control design and safety verification; [164] proposes a parallel SMC technique to speed-up the verification procedure of smart grids.

The literature review has shown that, whilst models of heterogeneous TCL are widely used, models for a inhomogeneous population PV systems are less considered. Furthermore, the use of model checking techniques for TCLs, PV systems, or more in general for power grid is rather limited: besides SMC approaches, the worlds of power systems and formal methods have rather little intersection. We borrow approaches from the TCL-related literature and apply them to a solar-based analysis. This work thus fills an important gap: it offers a novel modelling framework for heterogeneous photovoltaic systems along with simulating a contingency over a short (few seconds) time frame. Using techniques of formal abstractions, this thesis offers one of the first efforts of formal verification of the power grid for safety applications. The results in this dissertation may be thought as the first effort towards the construction of a formal tool to verify the safety of extreme situations for the electricity grid.

3

Description of Solar Devices

Contents

3.1 Photovoltaic System Definition

Photovoltaic solar systems exist in many different configurations with regard to their relationship to inverter systems, external grids, battery banks, or other electrical loads. Generally speaking, we can define a (household) PV system as the connection of solar panels, a DC/AC inverter, the house electrical load and a meter bridging the house with the grid. Optionally, the load can include a storage system. Figure 3.1 illustrates a scheme of a basic PV system that is attached to the grid.

We define a photovoltaic system as the connection of a PV panel, a power inverter and a frequency meter.

A PV panel is a modular object, composed of PV arrays which are in turn built on a series connection of PV cells. A PV cell thus represents the basic power generation unit, and is essentially a semiconductor that absorbs light and converts it into electricity [47]. Their intrinsic behaviour can be understood in terms of quantum states of the electrons in the material: either an electron is at rest within the so-called *valence band* or is travelling in the *conduction band*. The electrons need energy to jump into the conduction band and photons have the fundamental role of being the carriers of this energy. When photons hit the PV cell, they are absorbed in the semiconductor material and clash against electrons making them jump. Solar cells are engineered so that the electrons are only allowed to move in a single direction: this process creates the DC current as the output of the PV cell. In terms of mechanical construction, the PV panel is made up of Silicon cells, which are protected by a number of components that are engineered to operate under a variety of weather conditions – some manufactures claim their devices can withstand winds up to 160 km/h and up to 1.5 meters of snow. The literature on electric circuits representing a PV cell or a PV panel is rich, and offers models with increasing levels of complexity – e.g. [50–52].

In the following of the dissertation, we abandon the circuit subject to focus instead on the power output of solar panels. The current output of a PV panel is a DC current, while the power grid uses AC currents. The injection of current from a DC source is possible only via an inverter, that commonly uses a PWM [10] technique to obtain the desired AC output. Its task includes also the so called maximum power point tracking (see Section 3.2.1 for more details) to get the maximum possible power output for the PV panel for each working condition. The complex relationship between solar irradiance, cell temperature and load resistance provides a non trivial problem to solve (possibly in realtime). Within this PV system, the inverter is a critical component to a PV system's efficiency and reliability.

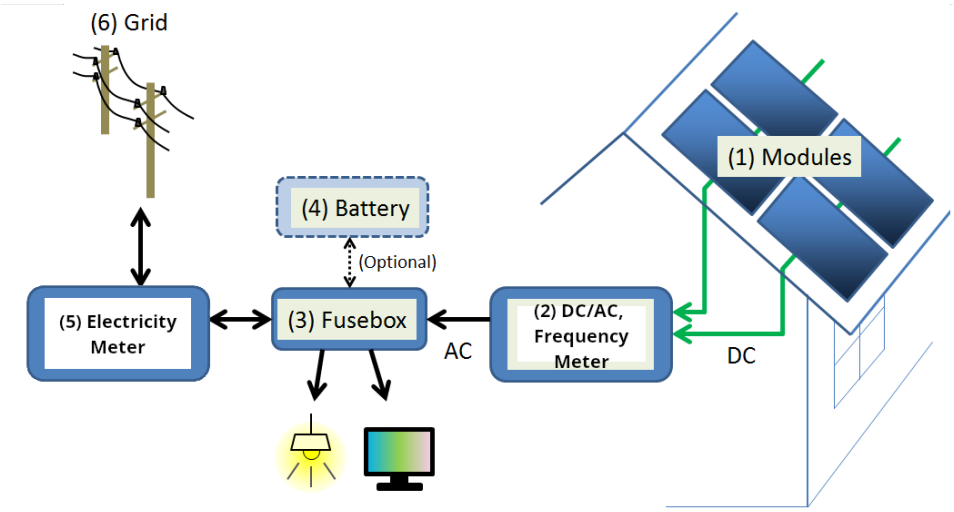

Figure 3.1: Household PV system scheme, modified from [47].

The inverter is also one of the most expensive components, after the PV modules themselves, and is responsible for up to 37% of system failures [47].

The literature about efficiency of household solar and solar is rich and varied: as an example, [165] compares the relative sizing of panels and inverters to improve overall performance; [12, 13] present in-field studies of household panels and solar farms assisted with computer simulations, under several weather conditions.

3.2 Photovoltaic System Operation

We illustrate the functioning of a single solar-inverter device, which will clarify our later modelling choice for the aggregation of a heterogeneous and large population of such devices. In the following we outline the *ideal* functioning of a solar unit; in the next Section 3.3 we instead report the *real* functioning of several devices, as per tests offered by our industrial collaborator at RTE France. We focus on how an inverter-panel system reacts as a function of the local electric grid. In this work we consider household devices, namely PV panels that are installed on the roof of private houses, as opposed to larger, less diverse industrial setups. Different manufacturers, weather conditions, ages, regulations, render this population naturally prone to a high level of heterogeneity.

A PV system can be either active or disconnected: accordingly, we define two working states, ON and OFF. At time k, the switch amongst these two states depends on two values: the local network frequency $f(k)$ (whose nominal value is $f_0 = 50$ Hz in Europe) and an internal time delay τ_r. Whilst in practice τ_r is a performance requirement given in seconds, we introduce a discretisation step h, and thus consider τ_r to be evaluated as a number of steps. Table 3.1 represents the switching behaviour of a PV system, considering the contributions of the frequency signal and of the internal time delay. The device is connected to the electric grid and samples its frequency discretely in time. Safety regulations demand the PV system to be active, i.e. in the ON mode, only when the network frequency lies within a given local frequency interval $\mathcal{I}_f$. If $f(k) \notin \mathcal{I}_f$ the device must disconnect from the grid, i.e. switch to the OFF mode, as soon as possible (ideally within one time step). We assume that the ON-to-OFF transition happens within a negligible time interval (see the following Section 3.3 and [166, 167]), so at time $(k+1)$ the device is OFF. On the other hand, the OFF-to-ON switch cannot happen before an internal delay τ_r has passed, during which the frequency f must continuously dwell within $\mathcal{I}_f$: this rule ensures the reconnection to a stable electric network, and avoids possible reconnection/disconnection scenarios (especially phenomena like chattering, i.e. the rapid disconnection and reconnection of devices) that may contribute to overall network instability.

A photovoltaic system's behaviour is frequency-dependent: it can be active only when the frequency remains close to the nominal value for a sufficient amount of time.

To guarantee that the frequency dynamics is stable, the inverter samples the grid frequency within the τ_r-long time interval: if the frequency is measured outside $\mathcal{I}_f$, the internal counter τ is reset. Notice that the network frequency sampled by different devices will differ, depending on several factors, as natural frequency fluctuations, measurement noise, high or low quality components: this motivates the use of (probabilistic) Markov model to describe such systems.

Current state $q(k)$	Frequency	Delay	Next state $q(k+1)$
OFF	$f(k) \in \mathcal{I}_f$	$\tau(k) \geq \tau_r$	ON
ON	$f(k) \in \mathcal{I}_f$	$\tau(k) = 0$	ON
ON	$f(k) \notin \mathcal{I}_f$	$\tau(k) = 0$	OFF
OFF	$f(k) \in \mathcal{I}_f$	$\tau(k) < \tau_r$	OFF

Table 3.1: Behaviour of a single photovoltaic system connected to the power network at time k.

The power output of an aggregation of PV systems adds uncertainty to the intrinsic diversity of the devices. We focus on a constant power production: this is a reasonable assumption over small time scales (seconds) and during a clear sunny day. However, this assumption might not hold during a cloudy day: as observed in [168], individual solar plants can significantly vary their outputs over the course of seconds. Weather should be encompassed in the model to have a better description of a real-time system: this is matter of future development.

3.2.1 Maximum Power Point Tracking

The maximum power point tracking (MPPT) is a technique used commonly in PV systems to maximize the power output under (possibly) all weather conditions [169–171]. Thanks to the digital nature of power electronics, a device can modify its operational point during its functioning and thus maximize its performance. The use of MPPT is not limited to solar, as it finds applications in other renewable, i.e. variable, sources: examples arise in wind power and optical power transmission [171].

Solar cells performance have a complex relationship between solar irradiance, load and temperature that shows a non linear efficiency. The MPPT method operates by sampling the PV output and searching for the right combination of load, current and voltage to attain the maximum power, under any weather condition. A wide range of solutions exists to find the optimal power point. Arguably the most common solutions fall under the *hill climbing* umbrella, a local optimisation technique that iteratively searches for a better power point making an incremental change to the current point. A survey of techniques is offered in [171, 172]. Whilst a local search method finds the optimal solution for stable weather conditions, e.g. a clear and

sunny day, it may return a sub-optimal solution under a less predictable, cloudy sky. A significant portion of related literature focuses in fact on partial shade conditions, which hold a highly complex, multi-modal relationship between power output and weather conditions. Variations of the perturb-and-observe methodology are most commonly employed; techniques evolved from the introduction of micro-controller based approaches [173] to include more advanced designs. A technique combining particle swarm optimisation and evolutionary algorithms is presented in [174]; the authors in [175] propose a two stage controller to find the maximum power point under rapidly changing solar irradiance. Newer solutions include feedforward [176], fuzzy control methods [177], along with neural networks based approaches [178].

The search for the MPP, especially using local optimisation methods, may require several iterations of the algorithm to find the operation point. In [179] the authors present a detailed theoretical and experimental comparison of two perturb-and-observe algorithms, in terms of system stability and performance characteristics. Following a change in input conditions (irradiance, cell temperature), the algorithms are able to track the MPP in few seconds: the process may take less than 1 second or 5 seconds according to the selected parameters and design choices. On the other hand, [174] shows that with a customization of the perturb-and-observe algorithm parameters tailored to the dynamic behaviour of the specific PV system (inverter and PV array) the MPPT can be achieved in hundreds of milliseconds when using a sampling time in the order of tens of ms. Other techniques [180], as fuzzy control logic or exploiting neural architectures, are shown to compute the MPP in under a second albeit at a higher implementation cost.

The functioning illustrated in Section 3.2 assumes that when a PV system reconnects to the grid readily injects the maximum power: the reconnection delays account for a reconnection at 90% of the expected power output, as outlined in the following Section. We argue that differences in this last top 10% of power output may be discounted as perturbations of the MPPT implementations (the algorithm may easily perturb the current operation point and reach a combination of voltage and current that returns a lower power output, in the range of a handful

of percentage points, as witnessed in, e.g., [179]). The transitory dynamics between 0 to 90% of the power output are thus embedded into the OFF state: the modelling framework may be refined by adding a further PV system state, to explicitly describe the power ramp up. Improvement of the modelling framework are matter of future efforts and discussed in detail in Chapter 8.

3.3 Parameters of Real Devices

Every device that is connected to an electricity grid must abide by tight regulations issued by single countries and, in Europe, by the ENTSO-E [26]. As such, PV systems' manufacturers must provide devices that fulfil the connection requirements in terms of both the working interval $\mathcal{I}_f$ and the time delay τ_r (see Section 3.2). European regulations [120] from 2017 outline a working frequency $\mathcal{I}_f = [47.5, 51.5]$ Hz and a time threshold τ_r that depends on devices' maximum power output. However, in previous years countries have drafted different legislations, sometimes changing them year by year, providing a significant source of population heterogeneity.

Furthermore, each device's dynamic behaviour heavily depends on the performance of its frequency measurement.

A widely common manner to measure the frequency is by implementing a phase-locked loop (PLL): a feedback system that returns a signal whose phase is related to the one of the input signal. A discussion about the various implementations and possible improvements of the PLL is out of the scope of this dissertation; the interested reader may find further details in, e.g., [181]. The performance of the PLLs varies in presence of disturbances, as outlined in [182] where a comparison among eight PLL implementations is presented, testing frequency and voltage steps, DC offsets and additive noise. Of particular interest for this work is the performance comparison in presence of noise and frequency jumps. Whilst under nominal conditions all the PLLs are able to measure the frequency with no issues, during the step transient only two implementations show a reliable frequency measurement, out of the eight PLLs tested. From the reported tests, the discussion

	Amplitude	Time	Measurement
Disconnection	0.01 Hz	5 s	Init. Step → Zero Current Output
Reconnection	0.02 Hz	50 s	Init. Step → 90% P_{nom}

Table 3.2: Step Test Settings.

also highlights how two PLLs are not reliable - they amplify the noisy signal - in presence of additive noise.

Tests Setup The tests outlined in the following encompass disconnection and reconnection thresholds together with disconnection and reconnection times. A key question is how to devise a meaningful and informative experiment that sheds light upon the internal behaviour of an inverter. The tests consider two types of frequency signals: a *frequency ramp* and a *frequency steps* input. The *frequency steps* experiment injects a constant frequency signal for a predefined amount of time that depends on the type of test considered. After the interval has passed, the frequency increases (overfrequency test) or decreases (underfrequency test) by a selected amount, that depends on the type of test considered. The disconnection and reconnection parameters are reported as follows. The disconnection frequency is measured when the inverter current output reaches zero. The disconnection time is computed as the interval between the moment in which the current reaches zero and the beginning of the corresponding frequency step. A similar procedure is carried out for the reconnection parameters. The reconnection frequency is measured as the inverter power output reaches 90% of its nominal value. The reconnection time is computed as the interval between that moment and the beginning of the frequency step. The step test settings are summarised in Table 3.2.

On the other hand, the *frequency ramp* experiment injects a constantly increasing (overfrequency test) or decreasing (underfrequency test), with a rate of 0.01 Hz/s for both tests. Since these tests cannot unambiguously define a disconnection or reconnection time, they are used solely as a validation of the frequency step tests. As such, we report in the following Section the results of the frequency ramp experiments.

3.3.1 Disconnection and Reconnection Tests

From a single device viewpoint, individual components as the inverter's frequency meter or its internal clock critically affect the device quality behaviour. Our industrial partners at RTE have tested [166, 167] the performance of commonly (in France) installed household solar devices. A total of 26 solar inverters, belonging to 13 different models (2 inverters per model) have been investigated both in an underfrequency ($f(k) < 50$ Hz) scenario and in an overfrequency scenario ($f(k) > 50$ Hz). The tests, in both scenarios, have evaluated four parameters: disconnection frequency, disconnection time, reconnection frequency, and reconnection time[1]. Figure 3.2 depicts the distribution of the measured reconnection (bottom plots) and disconnection frequency (top plots) considering an underfrequency scenario (left plots) and an overfrequency scenario (right plots). Similarly, Fig. 3.3 shows the distribution of time delays. The devices have different ages and consequently observe different regulations. Whilst they all present a 47.5 Hz (nominal) underfrequency reconnection and disconnection threshold, the overfrequency threshold is divided into three batches: 1) 50.2 Hz complying to norm DIN VDE 0126-1-1; 2) 50.4 Hz complying to norm FR 2013; 3) 50.6 Hz complying to norm FR 2014.

Results of the underfrequency tests report that the variance on the disconnection is limited, as all the inverters disconnect within a range ± 0.01 Hz around the nominal threshold; on the contrary, disconnection times range from 120 ms up to 700 ms. Similarly, the tested devices have reconnection frequency around 47.5 Hz with variations of ± 0.02 Hz. The reconnection times appear to be scattered among the tested devices and present high variance during the repetition of the experiments: values range from 19 s to 68.5 s. Notably, an outlier registers a reconnection frequency of 49.55 Hz and disconnection frequency of 49.0 Hz.

Results of the overfrequency tests show that the variance on the disconnection is limited, as all the inverters disconnect within a range ± 0.01 Hz around the nominal threshold; on the contrary, disconnection times present variations between 65 ms to more than 3 seconds. On the other hand, the tested devices respect

[1]Underfrequency results only account for 13 devices.

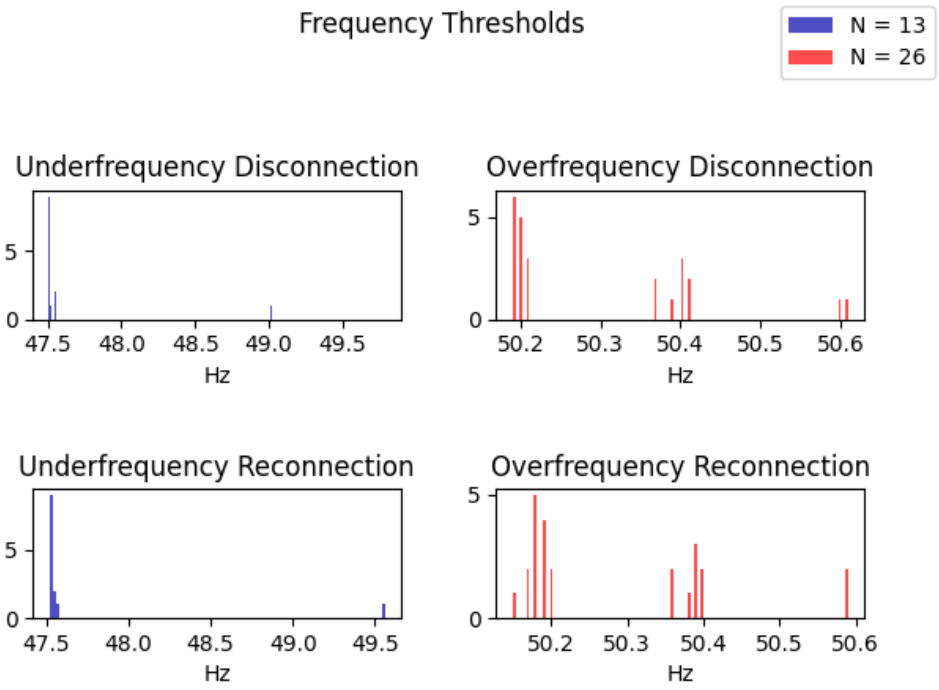

Figure 3.2: Distribution of measured frequency thresholds.

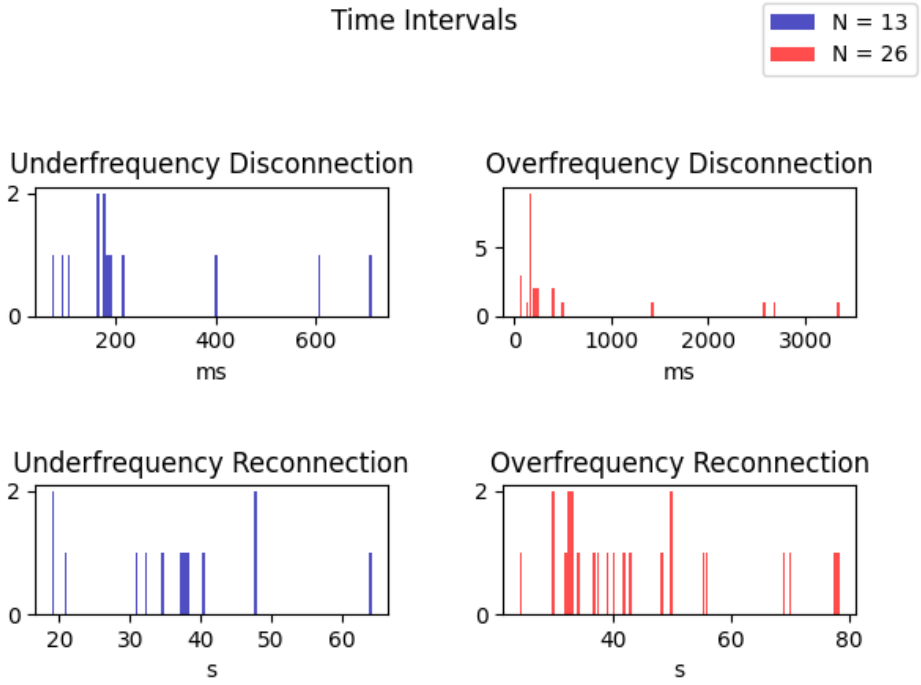

Figure 3.3: Distributions of measured time intervals.

three different legislations regarding the reconnection threshold but generally their reconnection frequency results within 0.01 and 0.02 Hz below the disconnection threshold. Notably, the reconnection times appear to be scattered among the tested devices and present high variance during the repetition of the experiments: values range from 24 s to 78.5 s.

3.3.2 Generalising Distributions

The first contribution of this work is the modelling of a large population of solar devices: hence, we generalise the obtained frequency thresholds and time intervals to encompass the behaviour of the European solar population. As an example, let us focus on the reconnection in overfrequency scenario: the other three settings (reconnection in underfrequency and disconnection in under-/overfrequency) can be similarly analysed.

We propose two distribution fittings for the frequency threshold using a Gaussian and a χ^2 distribution. A Gaussian distribution classically models measurement noise: it depicts a situation where all devices have a reference threshold value – in this scenario three values, 50.2 Hz, 50.4 Hz and 50.6 Hz – but the devices reconnect when the frequency is *around* the nominal value, in view of imperfect frequency measurement. To best fit the real data, we use three Gaussian distributions with average values $\mu_1 = 50.18$ Hz, $\mu_2 = 50.39$ Hz, $\mu_3 = 50.59$ Hz and variance $\sigma_1^2 = 5 \cdot 10^{-4}$, $\sigma_2^2 = 10^{-3}$ and $\sigma_3^2 = 10^{-3}$ respectively, as shown in Fig. 3.4. The parameters of average and variance μ_i, σ_i, $i = 1, 2, 3$, are obtained via a brute force search as the parameters set that minimises the mean squared error [183] between the Gaussian distribution and the histograms' values. Practically speaking, we have chosen the parameters that minimise

$$L(x) = \frac{1}{T} \sum_{i=1}^{T} (\Phi(x_i) - h(x_i))^2, \tag{3.1}$$

where $\Phi(\cdot)$ represents the Gaussian probability distribution function, $h(\cdot)$ represents the histogram's value, and x_i denotes the position of the individual histogram bar, i.e. the histogram's x-axis values. T is the total number of samples: in the case

under consideration is 26, i.e. the number of tested inverters. Notice that by the central limit theorem [184] we expect the distribution of a large number of devices to converge to a Gaussian distribution. The second candidate, a χ^2 distribution, models a *minimum performance* requirement. The grid connection code [185] imposes a *minimum* working interval, hence devices can remain active within a wider interval. In other words, devices can disconnect/reconnect at higher frequency value after an overfrequency deviation and disconnect/reconnect at lower frequency values after an underfrequency event. Notably, χ^2 distributions have support over a semi-infinite interval $[\lambda, +\infty)$ and its shape is defined with *degrees of freedom* φ, rather than with the variance. We defer the details of such probability distributions to [184]. To fit the real data, we use three χ^2 distributions with initial points $\lambda_1 = 50.14$ Hz, $\lambda_2 = 50.35$ Hz, $\lambda_3 = 50.45$ Hz and degrees of freedom $\kappa_1 = \kappa_2 = \kappa_3 = 4$ respectively, as shown in Fig. 3.5. Similarly to the Gaussian case, the parameters λ_i, κ_i, $i = 1, 2, 3$, are obtained via a brute force search as the combination that minimises the mean squared error between the distribution and the histogram's value.

The number of tested devices might be insufficient to definitively determine which distribution better represents the data, hence we develop a Markov model (see Chapter 4) that allows using any frequency threshold distribution. Further, we deploy tests to detect the risk of load shedding (see Chapter 5) with both Gaussian and χ^2 distributions and compare the results. We will add a Uniform distribution modelling from Chapter 6 to simplify the formal verification techniques.

We generalise the population thresholds using continuous distributions. In particular, we will consider a Gaussian distribution, that models noisy frequency measurements, a χ^2 distrobution, that models a minimum operational interval of frequency, and a uniform distribution, that models a general uncertainty.

Disconnection and Reconnection Times Disconnection times (cf. Fig. 3.3, top left and right) present almost all entries below 200 ms, with some significant outliers up to 3000 ms. For simplicity, in Chapter 4 we will model the disconnection as an immediate reaction to a frequency deviation, i.e. as if all disconnections happen in less than one time step – that is usually set to 200 ms. Such a framework

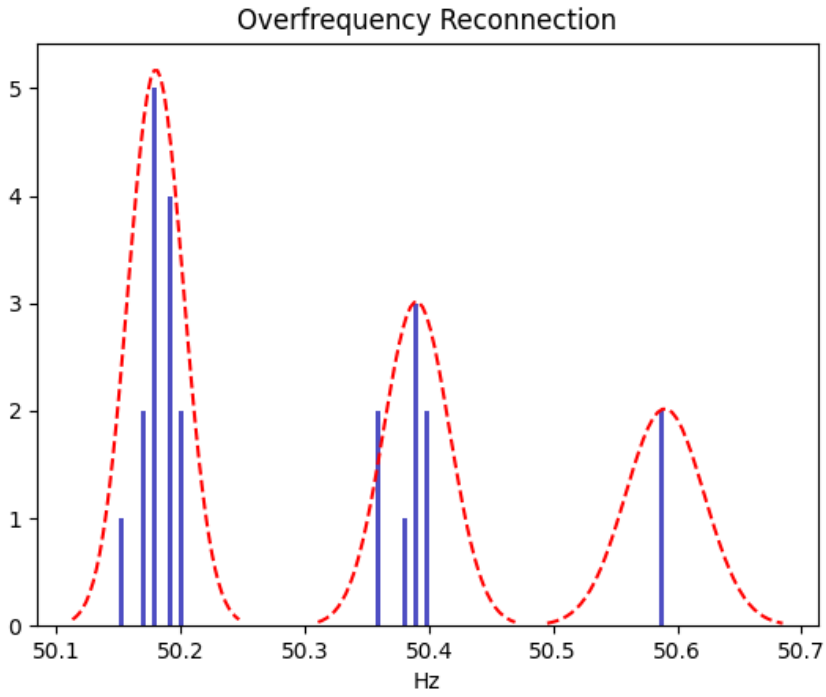

Figure 3.4: Frequency thresholds with three Gaussian approximations (red).

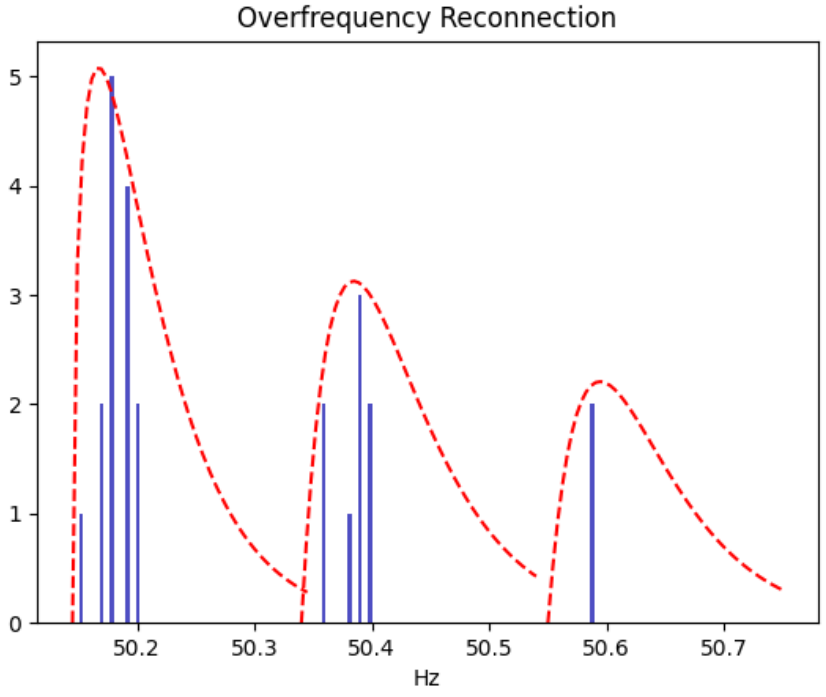

Figure 3.5: Frequency thresholds with three χ^2 approximations (red).

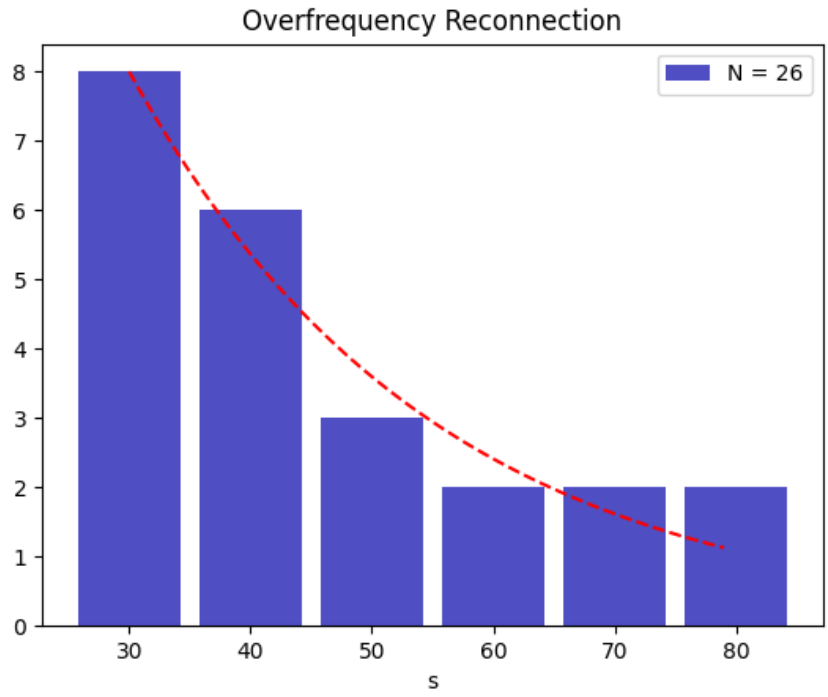

Figure 3.6: Reconnection delays with an exponential approximation (red).

models solar devices complying with the grid connection code [120, 185], which imposes a quick disconnection following a significant frequency deviation. In Chapter 4 we will briefly discuss a model extension encompassing a more complex disconnection pattern for completeness.

Interestingly, the European grid connection code [120, 185] delegates to the single TSOs the choice of the most appropriate automatic reconnection delays. As an example, in the UK the Energy Network Association (ENA) recommends the standard EREC G83 [186], which defines a minimum reconnection delay of 20 seconds, without further specifications (e.g., there is no mention of a possible maximum reconnection delay, nor a randomised reconnection delay). In Italy, the standard CEI 0-21 (V1) [187] recommends a minimum reconnection delay of 30 seconds, similarly with no further detail.

The distribution of reconnection times (cf. Fig. 3.3, bottom left and right) is more scattered and its interpretation is not trivial, as (almost) every device presents a different reconnection delay. Further, rather wide gaps prevent the

selection of a uniform distribution; as an example, the reconnection in overfrequency (Fig. 3.3 bottom right) presents the majority of entries between 30 and 40 seconds with the distribution fading towards 70 seconds. We thus decide to cluster the data, grouping the reconnection delays around tens of seconds, as shown in Fig. 3.6. Reconnection delays of the underfrequency scenario present a similar shape – a depiction is omitted for brevity. The reconnection delays result shaped as an exponential decay, as classically suggested by queueing theory [183]. Notice that this distribution can translate a minimum requirement: the reconnection delay must be *at least* τ_r seconds. We will later use the geometric distribution – which models a discrete exponential decay – to describe the reconnection delays in the Markov models introduced in Chapter 4.

Correlation between Thresholds and Reconnection Times Insofar, we have dealt with frequency thresholds and time delays as separate quantities. A further insight on their distribution is given by their covariance. Intuition suggests that devices built with poor quality components may also have poorly implemented firmwares, leading to actual disconnection/reconnection thresholds distant from the nominal values. In such cases, a larger reconnection delay may be set in order to compensate for the noisy frequency measurement. Similarly, good components result in a frequency threshold and a reconnection delay close to the nominal values. We consider the Spearman's rank correlation ρ_c, which is a robust evaluation of the relationship between two variables [183, 184]. The quantity ρ_c belongs to the interval $[-1, 1]$ and it is zero when there is no tendency for one variable to either increase or decrease when the other variable increases; it is 1 (resp. -1) when each of the variables is a perfect monotone increasing (decreasing) function of the other. By and large, a coefficient ρ_c below -0.5 or above 0.5 indicates a notable correlation, and a correlation below – in absolute value – those values suggests a less notable relationship between the two variables. We will not discuss further any notion on the correlation; for a detailed description of the correlation coefficients, please refer to [184]. We compute ρ_c for the four pairs obtained combining the disconnection/reconnection thresholds

Scenario	ρ_c
Disconnection, Underfrequency	0.05
Disconnection, Overfrequency	0.28
Reconnection, Underfrequency	0.01
Reconnection, Overfrequency	0.19

Table 3.3: Correlation coefficient between frequency thresholds and time delays.

with the disconnection/reconnection delays, whose evaluation is reported in Table 3.3. All values of ρ_c are positive, which indicates a tendency to have longer time delays in presence of larger frequency thresholds; all values are, however, close to zero, suggesting a weak relationship between the two quantities. In view of the low correlation values, we build the Markov model (see Chapter 4) assuming independence between frequency thresholds and time delays.

3.4 Dispersed Generation in Europe

An investigation on the load shedding risk in presence of a significant penetration of renewable energy source is outlined in the ENTSO-E Dispersed Generation [41] and Assessment of System Security [188] reports. We present here the risk assessment evaluation, adding data analysis and the main conclusions.

3.4.1 Dispersed Generation and Frequency Quality

The interconnected system of Continental Europe is composed of 25+1 state members[2], it extends from Portugal to Poland and from Denmark to Turkey, and feeds a load between 220 GW to 440 GW[3]. This large system is operated in a synchronous way, meaning that, when we neglect phenomena with time constant smaller than a second, the frequency is identical everywhere.

In the last decade, the penetration of distributed energy sources has increased significantly all over Europe. Mostly of the renewable type, these generating units are connected to distribution systems and have been subjected to connection

[2]Turkey is an observer member.

[3]The load demand between 220 GW and 440 GW represent a low-load and high-load scenario from 2016. Figures from 2019 indicate a low-load scenario of 300 GW and a high-load scenario of 550 GW.

requirements compliant with the historical planning and operating principles of distribution systems, which were designed for passive loads. In areas where the amount of dispersed generation has increased, standard electrical faults cannot be efficiently managed by the existing protection schemes: this may lead to islanding or undefined system conditions. Safety is a major concern, thus monitoring sub-networks connection to the synchronous area is crucial. In particular, TSOs detect a sub-network disconnection by measuring frequency deviations: whenever the frequency of a sub-network significantly deviates from the reference value, the control routine triggers the disconnection of generation units to avoid keeping energized isolated networks. However, if applied simultaneously to a large number of units, unique frequency thresholds can jeopardize the security of the whole interconnected system in case of frequency deviations.

Nowadays renewable sources frequently provide more than 20% of the system demand. At the same time [5], the quality of frequency in the Continental Europe Synchronous Area has constantly decreased: the system frequency deviates from the reference value more often and for longer time periods. We use the UK grid's frequency values from Gridwatch [189] that stores frequency measurements every 10 minutes. Figure 3.7 depicts the monthly number of frequency deviations larger than 75 mHz, i.e. whenever the frequency registers values below 49.925 Hz or above 50.075 Hz. The spike around the beginning of a new year indicates that a great number of deviations happens during the winter months, when the demand is higher. Similarly, Fig. 3.8 shows the number of deviations per year – data for the year 2020 are updated to May. Since 2014, the number of deviations has steadily grown, with a light decline in 2019: we can interpret this figure as the instability brought by the increasing renewable penetration in the power grid. From the control viewpoint, frequency deviations are expensive: exceeding 49.9 Hz or 50.1 Hz corresponds to an activation of a significant portion of the primary control reserve, up to 50%.

We may analyse the historical frequency signal to gather useful information about the grid. In particular, we are interested in finding the average value and

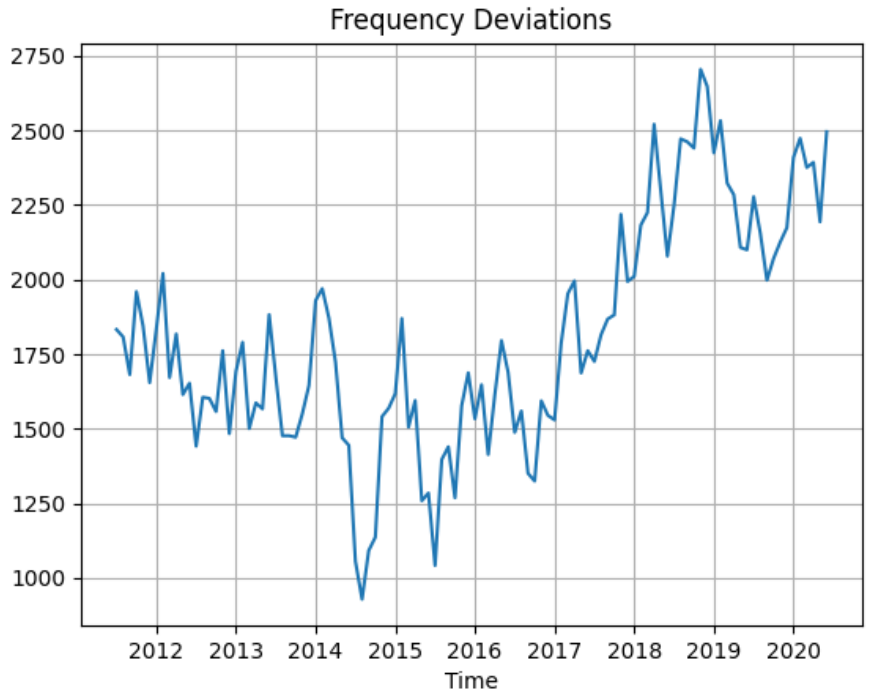

Figure 3.7: Number of significant frequency deviations per month.

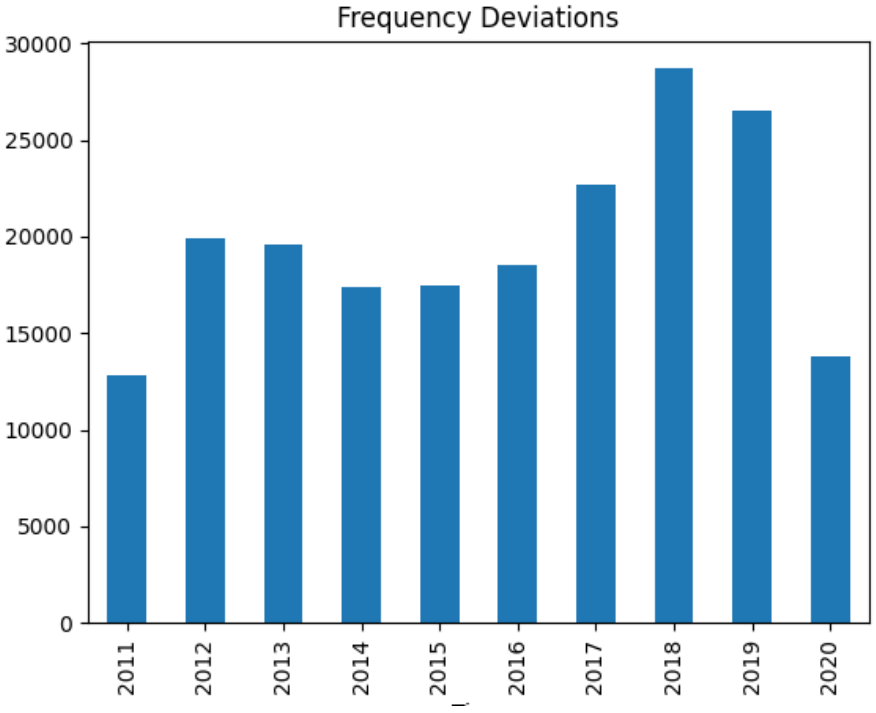

Figure 3.8: Number of significant frequency deviations per year.

variance of the grid frequency over the last years. We then employ a maximum likelihood method [183, 184] to extract this information. We find

$$\mu_f = 50.0011, \qquad \sigma_f^2 = 0.025, \tag{3.2}$$

witnessing that the frequency is slightly above the reference, with a rather small standard deviation. We will use this value of variance in the case studies in Chapters 5, 6, 7.

3.4.2 Dispersed Generation and Power at Risk

In several European countries, a significant portion of distributed generation units – solar and wind in particular – have protection settings for automatic disconnection from the electricity grid which do not comply with the standard disconnection limits of the transmission system. The disconnection/reconnection settings for distributed generation are in the range [50.2, 50.3] Hz for overfrequency and within [49.7, 49.5] Hz for underfrequency, while the (current) mandatory working interval spans between 47.5 and 51.5 Hz. A few examples: up until 2014, Germany has reported 14000 MW of PV with overfrequency automatic disconnection setting of 50.2 Hz and 189900 MW of installed wind generation with underfrequency disconnection setting of 49.5 Hz. Italy has reported 11500 MW of PV with disconnection settings of 50.3 Hz and 49.7 Hz for over and underfrequency respectively. The risk of serious system disturbances due to an uncoordinated disconnection of distributed generation units cannot be excluded. Under such great stress conditions, the system balance might be restored only by the activation of large scale underfrequency load shedding.

Table 3.4 shows the status of solar dispersed generation in ENTSO-E area. The second column reports each country's total solar devices with non-compliant disconnection thresholds in MW. The subsequent columns contain the amount of solar generation per disconnection threshold. In recent years, European countries have started a *retrofit* program, upgrading solar (and other renewable) sources to comply with existing regulations. The total solar generation with non-compliant thresholds after the retrofit program is displayed in the *After Retrofit* row. As

Country	Total Capacity [MW]	50.5 [Hz]	50.3 [Hz]	50.2 [Hz]	49.8 [Hz]	49.7 [Hz]	49.5 [Hz]
Germany	24800	–	–	14000	–	–	–
Italy	14300	–	11500	–	–	11500	–
Spain	4047	30	–	–	–	–	–
France	2500	–	–	2500	75	–	500
Czech	1900	–	–	950	–	–	–
Belgium	2225	600	–	1100	–	–	–
Greece	1000	1000	–	–	–	–	1000
Slovakia	512	–	–	512	–	–	–
Portugal	155	8	–	79	79	–	8
Denmark	290	–	–	6	–	–	–
Poland	5	–	–	5	–	–	–
Hungary	4	–	–	4	–	–	–
Total	51738	1638	11500	19156	154	11500	1508
After retrofit	16556	1638	2300	8656	154	2300	1508
Total renewables after retrofit	30088	2670	2300	8656	154	2300	14008

Table 3.4: Power at risk in Europe.

mentioned, other renewable sources – especially wind generation – show a narrow threshold issue. The total renewable generation with non-compliant thresholds after the retrofit program is finally reported in the last row. A total of more than 16 GW of solar and more than 30 GW of renewable sources are connected to the main electricity grid with non-compliant disconnection thresholds, hence are considered *power at risk*. In particular, the amount of generation from PV and wind tripping at 50.2 Hz or 49.5 Hz causes concerns, and shows the need for further retrofit programs.

3.4.3 Classification of Contingencies

Whilst the overall quality of the network operations in the Continental Synchronous Area has notably decreased, reaching frequency values of 50.2 Hz or 49.5 Hz is highly unlikely during normal operations. Nevertheless, frequency may reach such (and more extreme) values in the instants following a network incident. Incidents are classified according to Operational Security Network Code [190] of ENTSO-E as:

- Normal incidents,

- Exceptional incidents,
- Out-of-range incidents.

Normal incidents are triggered by a single system failure. They must be managed in a secure manner: in other words, the function of the electric network must continue without violation of any technical limit. In case of exceptional contingencies, the interruption of the energy supply is accepted; however, the integrity and stability of the network must be ensured. Out-of-range incidents may have severe consequences on the network's infrastructure and might be not controllable for the system[4].

This work studies in particular the behaviour of the network after a *normal* incident, which is related to power imbalances and frequency control. Normal incidents are characterised by:

- Loss of load $\leq$ 2 GW, or
- Loss of generation $\leq$ 3 GW,

which might be caused by single events like: *a*) trip of the HVDC link between France and Great Britain, or *b*) busbar failure with generation loss. The combination of two or more such events is considered as *exceptional* incident – a loss of load greater than 2 GW, or tripping of generation between 3 GW and 6 GW. An example of out-of-range incident is represented by multiple contingencies leading to system splitting with high power imbalance in different areas.

3.5 Concluding Remarks

This Chapter has defined the PV system as an overall unit composed by solar panels, inverters and meters, which represents the main object of study of this dissertation. We have described its ideal operations in response to frequency deviations and tested the actual behaviour of several solar units. We have then generalised the test results to a large population of devices and have outlined the foundations of the modelling of a heterogeneous population. Building upon this, Chapter 4 will offer our modelling

[4]In such a case, a so-called *Defence Plan* is implemented in order to avoid a total blackout.

framework for a large and heterogeneous population of solar devices. These models support any disconnection and reconnection distribution to counterbalance the limited data gathered from our partners' tests. As future effort, we may think of sharpening the experimental results on disconnection and reconnection parameters. Beside increasing the number of device tested, we propose to repeat the experiments for every device, to understand the stochastic component per single PV system. Further, we are interested in testing new against old PV systems to offer an insight on the performance degradation of the internal components; experiments under different shading patterns may reveal a correlation between these parameters and weather conditions. Finally, corroborated by reports about the power at risk in Europe, we have provided a numerical analysis of the quality of the grid's frequency following an increased penetration of renewable energy sources.

4

Markov Models of Photovoltaic Systems

Contents

4.1 A Markov Model for a Population of PV Systems

Following the operational description of a solar device, we present our Markov modelling. As outlined in Chapter 2, the behaviour of a single system depends on its sampling of the frequency signal: this feature suggests a discrete-time modelling framework. In view of the limited time-scale that interests our study, we assume constant power production over time and over different PV systems, as well as population homogeneity, i.e. all devices are characterised by the same parameters. This represents a limitation of our modelling: we aim at extending

our models by adding a weather conditions based, stochastic description of the power output, as per [191].

We start our discussion illustrating the modelling of a simple homogeneous, delay-free population; we later introduce heterogeneity in the disconnection and reconnection thresholds; finally, we equip the models with a timed reconnection structure, as reported in Chapter 3. The architectures for a delay-free homogeneous and heterogeneous population are necessary steps for the derivation of the final model; nevertheless, we will solely consider the heterogeneous, delayed model in the rest of this work. We initially outline an $(n+2)$-state model structure, where n is the number of reconnection time steps needed to describe the whole population. We are then able to provide a reduced structure by using a time-varying coefficient: the final model structure is composed of only three states. This lumping comes at a price: the Markov model becomes time-varying, with the time-dependent coefficient that can be estimated from previous frequency measurement.

Last, we test our models against an explicit model of a heterogeneous population of PV systems. We simulate N devices, individually implemented, and compare their frequency response to disturbances to the response of our models. As $N \to \infty$ the signals converge until they exactly overlap, witnessing the reliability of our framework and justifying the use of such models.

4.2 A Homogeneous Population Without Delays

The model describing the interconnection between PV systems and the power grid relies on two quantities: the network frequency and the power output coming from a population of PV systems. At first, we formalise the solar power output. As an early simplifying assumption, we hypothesise population homogeneity: every device is a perfect replica, showing the same behaviour in time and is characterised by the same parameters. This allows us to define a quantity P_{PV} expressing the weighted power production of the whole population as

$$P_{PV} = \frac{1}{N} \sum_{i=1}^{N} P_{PV,i},$$

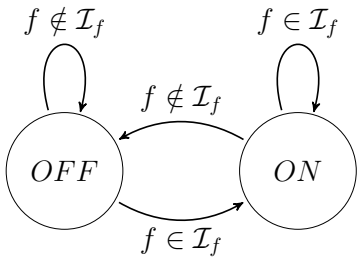

Figure 4.1: A Markov chain for a homogeneous population.

where N is the number of PV systems in the population, and $P_{PV,i}$ is the i-th system's power output. Consider the normalised power production $R(k)$, at time k, as

$$R(k) = \frac{1}{NP_{PV}} \sum_{i=1}^{N} q_i(k) P_{PV,i},$$

where $q_i(k) \in \{0, 1\}$ denotes whether the i-th device at time k is in the OFF or ON state, respectively.

The behaviour of q_i can be described by a Bernoulli distribution: q_i is a discrete random variable which takes value 1 with probability p – that depends on the network frequency – and takes the value 0 with probability $1 - p$.

In view of the population homogeneity, all the devices behave in a synchronous way: either $q_i = 1$ or $q_i = 0$ $\forall i$. Such a population naturally induces a dtMC with two states, ON and OFF as per Fig. 4.1. As a result, also $R(k)$ can only take value 1 or 0 accordingly

$$R(k) = \begin{cases} 1 & \text{if } q_i = 1 \ \forall i, \\ 0 & \text{if } q_i = 0 \ \forall i, \end{cases}$$

thus also $R(k)$ can be described by a Bernoulli random variable. We are interested to define the average power output of the population of PV systems. To this end, let us introduce a new variable

$$x(k) = \mathbb{E}[R(k)]. \tag{4.1}$$

where $\mathbb{E}[\,\cdot\,]$ denotes the probabilistic expectation operator. By definition of expectation,

$$\begin{aligned} x(k) = \mathbb{E}[R(k)] &= \mathsf{P}[R(k) = 1] \cdot 1 + \mathsf{P}[R(k) = 0] \cdot 0 \\ &= \mathsf{P}[R(k) = 1], \end{aligned} \tag{4.2}$$

hence $x(k)$ defines the expected value of $R(k)$ at time k, and additionally denotes the probability of being in the ON state at that time. Notice that $x(k)$ is a function of time, as $R(k)$. We are interested in formally writing its dynamical evolution with respect to its value in the previous time instant. To this end, we may use the law of total probability

$$\begin{aligned} x(k+1) &= \mathsf{P}[R(k+1)=1] \\ &= \mathsf{P}[R(k+1)=1 \,|\, R(k)=1] \cdot \mathsf{P}[R(k)=1] \\ &\quad + \mathsf{P}[R(k+1)=1 \,|\, R(k)=0] \cdot \mathsf{P}[R(k)=0]. \end{aligned} \tag{4.3}$$

Let us now analyse each term of the right-hand side of Eq. (4.3). The term $\mathsf{P}[R(k+1)=0 \,|\, R(k)=1]$ represents the probability of transitioning towards the OFF state ($R(k+1)=0$) from the ON state ($R(k)=1$); we will denote it as

$$a(k) := \mathsf{P}[R(k+1)=0 \,|\, R(k)=1].$$

On the contrary, the term $\mathsf{P}[R(k+1) = 1 \,|\, R(k)=1]$ represents the probability of switching towards the ON state ($R(k+1)=1$) from the ON state ($R(k)=1$); we may denote this probability $(1-a(k))$, as it is the complement to one of $a(k)$. Similarly, we denote

$$b(k) := \mathsf{P}[R(k+1)=1 \,|\, R(k)=0]$$

the probability of transitioning towards the ON state ($R(k+1)=1$) from the OFF state ($R(k)=0$). Using Eq. (4.2) and noting that

$$x(k) = \mathsf{P}[R(k)=1] \quad \Longrightarrow \quad \mathsf{P}[R(k)=0] = 1 - x(k),$$

Equation (4.3) can then be written as

$$x(k+1) = (1-a(k))x(k) + b(k)(1-x(k)). \tag{4.4}$$

This relation describes how the probability of being ON gets updated at time k. In the adopted framework, the transition probability ON-to-OFF (quantity $a(k)$) and OFF-to-ON (quantity $b(k)$) are governed by the value of the network frequency,

namely whether or not $f(k)$ is within the working interval $\mathcal{I}_f$. Recall that we are assuming population homogeneity, thus $a(k)$ and $b(k)$ can only evaluate to 0 or to 1:

$$f(k) \in \mathcal{I}_f \implies \quad a(k) = 0, \quad b(k) = 1,$$

$$f(k) \notin \mathcal{I}_f \implies \quad a(k) = 1, \quad b(k) = 0.$$

Whilst quantities $a(k)$ and $b(k)$ are deterministic in a homogeneous population, this assumption is restrictive in the case of a heterogeneous population. We will encapsulate the population heterogeneity within terms $a(k)$ and $b(k)$: specifically, we assume different PV systems to be subjected to different disconnection and reconnection thresholds, in view of different ages, manufacturing processes, weather conditions, etc. This results in two modelling choices, as illustrated in Section 3.3. We either assume that the different thresholds in a population are distributed according to a Gaussian or, alternatively, we consider a χ^2 distribution. Our modelling evidently supports other choices: in fact, our study can encapsulate the use of any probability distribution that might be empirically fit from population data.

4.3 A Heterogeneous Population Without Delays

Section 3.3 has outlined the behaviour of real solar units as a function of the (simulated) network frequency. From these tests, we have deducted the disconnection and reconnection thresholds for each device: we have reported that, in general, each soalr system presents different frequency thresholds, disconnecting or reconnecting at different values of frequency. Our modelling framework subsumes a large population of devices, where the following holds:

Assumption 1 *Frequency thresholds over the whole population are distributed according to known continuous probability distributions.*

In other words, we assume that the sources of heterogeneity, as age or components quality, introduce a continuous disturbance on the frequency measurement of the solar devices. In practical terms, disconnection and reconnection distributions

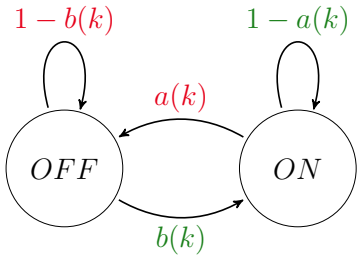

Figure 4.2: A time-varying Markov chain for the population dynamics.

define the transition from the ON state, i.e. $R(k) = 1$, to the OFF state, i.e. $R(k+1) = 0$, and vice versa.

We introduce $a(k)$ and $b(k)$ as the ON-to-OFF and OFF-to-ON transition probability, respectively (cfr. Fig. 4.2, transition towards the ON state are coloured in green, towards the OFF state in red). Notably, $a(k)$ ($b(k)$) encompasses two different disconnection (reconnection) scenarios: an over-frequency scenario, when $f(k) > 50$ Hz and an under-frequency scenario, when $f(k) < 50$ Hz. Notice also that $a(\cdot)$ and $b(\cdot)$ are not binary, as opposed to the homogeneous case: their values are now computed by integrating probability distribution functions, in the following way.

Let us consider the over-frequency, disconnection scenario. Whenever $f(k)$ evaluates to a value greater than $f_0 = 50$ Hz, solar units compare the frequency value to their disconnection threshold: if $f(k)$ exceeds the disconnection threshold, the device must disconnect. On the other hand, if the disconnection threshold is greater than $f(k)$, the device may remain active. From the distribution perspective, the value $f(k)$ is used as one of the extrema of the integral: we, in fact, compute the amount of devices that have thresholds *up to* the current frequency value. For a pictorial depiction, see Fig. 4.3. In this way we obtain the portion of PV systems that are enabled to transition to the other state. Formally,

$$a(k) = \begin{cases} \int_{-\infty}^{f(k)} p_o^d(u)du & \text{if } f(k) > f_0, \\ \int_{f(k)}^{+\infty} p_u^d(u)du & \text{otherwise,} \end{cases} \tag{4.5}$$

$$b(k) = \begin{cases} \int_{f(k)}^{\infty} p_o^r(u)du & \text{if } f(k) > f_0, \\ \int_{-\infty}^{f(k)} p_u^r(u)du & \text{otherwise,} \end{cases} \tag{4.6}$$

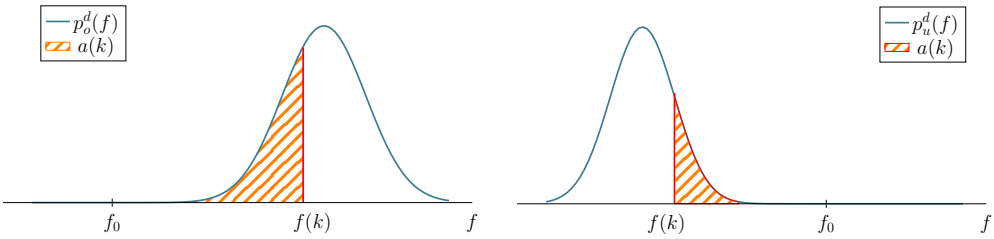

Figure 4.3: Representation of $a(k)$ in over-frequency (left) and in under-frequency (right).

where p_i^j denotes the probability distribution function describing the frequency thresholds; the superscripts d and r indicate the disconnection and reconnection, respectively; the subscripts o and u further denote over- and under-frequency, respectively. Note that intuitively $a(k)$ is computed as the part of the integral that is closer to f_0, whereas $b(k)$ is the part that is further away from f_0. The evaluation of $a(k)$ in under- and over-frequency is shown in Fig. 4.3. As an alternative, $a(k)$ and $b(k)$ can be expressed via corresponding cumulative distributions.

The model of a heterogeneous population without delays can be subsumed by Eq. (4.4), where now $a(k)$ and $b(k)$ take values within the whole interval $[0, 1]$. Equation (4.4) may be represented by a two-state dtMC with time varying transition probabilities, as depicted in Fig. 4.2. Note that $a(\cdot)$ and $b(\cdot)$ are function of the current frequency $f(k)$: to ease notations, we will denote them as $a(k)$ and $b(k)$ instead of $a(f(k))$ and $b(f(k))$.

The value $x(k)$ can be alternatively interpreted as the *portion of soalr systems* that are ON at time k. This modified point of view brings us a jump in perspective: from a model of a *single* PV system to an *aggregated* model for the population.

In general, the four distributions p_j^i, $i = d, r$, $j = u, o$ can belong to different families of probability distributions. We may think of a special scenario, where the reconnection and disconnection are symmetrical, i.e. belonging to the same distribution. If this is the case, we are allowed to substitute $b(k)$ – the reconnection probability – with $1 - a(k)$ – the complement of the disconnection. Equation (4.4) simplifies to

$$x(k + 1) = 1 - a(k), \tag{4.7}$$

indicating that the portion of active population is completely identified by the frequency threshold distribution.

4.4 A Heterogeneous Population With Delays

We now introduce a framework to encompass delays in the population model: as observed in practice, PV systems cannot reactivate instantaneously but must wait a delay. We generalise the discussion introduced in the previous section by allowing for possibly random delays. We assume to know the delay probability distribution function, and introduce transition values τ_i: value τ_i represents the probability of switching to state ON, given that the PV system has been *waiting* for i time instants (which means that $f(k) \in \mathcal{I}_f$ for i time instants, from $f(k-i)$ until the current value $f(k)$).

In other words, assume a device is disconnected, i.e. in the OFF state, at time k. It samples the frequency and may reconnect – according to the thresholds condition – with probability $b(k)$. The soalr device however cannot directly activate because it must wait a delay τ. Hence, with probability $b(k)$ the PV system will go into a waiting state, say w_1. At the next time step $(k+1)$, the PV system samples again the frequency. With probability $b(k)$ it passes the threshold condition, and with probability τ_1 the PV system may reconnect directly after 1 time step. The total reconnection probability thus becomes $\tau_1\ b(k)$. Similarly, in waiting state w_i a PV system may reconnect with the grid with probability $\tau_i\ b(k)$.

We generalise this idea and utilise n waiting states, and in particular $n \gg 1$, defined as $w_i,\ i = 1, \ldots, n$, representing the i-th time step in which $f(\cdot) \in \mathcal{I}_f$ with the device still OFF. Further, backed by the measurement presented in Section 3.3, we assume that the τ_i delays are distributed according to a geometric distribution: this is often used to characterise arrival processes or waiting-time random variables. This distribution has the property that $\forall i, \tau_i \geq \tau_{i+1}$ and $\sum_i \tau_i = 1$.

Let us focus on the i-th delay state. A PV system that is waiting for i time instants can either re-activate or keep waiting, if the frequency is close to the nominal value, or go back to the OFF state, where the reconnection counter is reset.

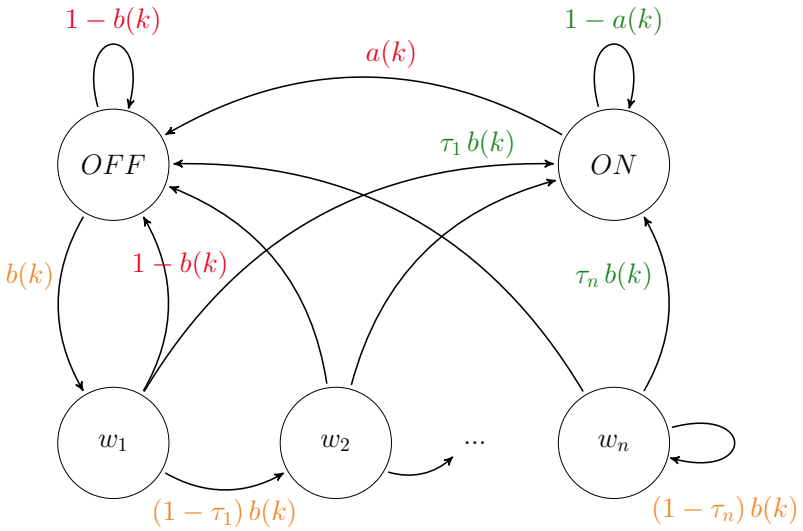

Figure 4.4: A Markov model for the aggregated dynamics.

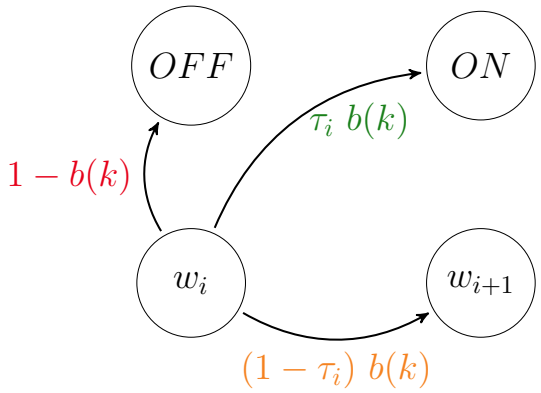

Figure 4.5: Outgoing transitions of a single waiting state.

In the dtMC framework we identify three outgoing transitions: one towards the ON state (reactivation), a second towards state w_{i+1} (keep waiting), and one back to the OFF state (cf. Fig. 4.5). The probability associated with the latter transition is $1 - b(k)$, which is the $f(k) \notin \mathcal{I}_f$ probability. The first outgoing probability is $\tau_i\ b(k)$: τ_i is the delay probability value to transition from state w_i to state ON, multiplied by $b(k)$, namely the probability of $f(k) \in \mathcal{I}_f$. The transition towards w_{i+1} requires two conditions to hold: the thresholds requirement and the opposite of the reconnection, i.e. the PV system keeps waiting. The transition probability thus results in $b(k) \cdot (1 - \tau_i)$, as a product of the two conditions.

Recall that, to allow for a reconnection, the condition $f(k) \in \mathcal{I}_f$ must be satisfied

at all time instants, corresponding to the condition on the counter $\tau \leq \tau_r$. At a result, every transition towards the ON state or the next waiting state is multiplied by $b(k)$.

Differently from the delay-free model in Eq. (4.4), we compute (refer to Fig. 4.4) the probability of switching ON – that also represents the portion of activate solar devices – by adding the transitions from the waiting states

$$x(k+1) = (1 - a(k))\, x(k) + b(k) \sum_{i=1}^{n} \tau_i w_i(k). \tag{4.8}$$

The term $(1 - a(k))\, x(k)$ denotes the portion of active solar devices that will remain active at the next time step; terms $b(k)\, \tau_i\, w_i(k)$ denote the portion of waiting solar devices – $w_i(k)$ – that with probability τ_i will transition towards ON, provided that they sample the frequency within the working interval – $b(k)$. The i-th waiting state has one incoming transition from the previous waiting state, holding

$$w_i(k+1) = (1 - \tau_i)\, b(k)\, w_{i-1}(k), \quad i = 2, \ldots, n-1, \tag{4.9}$$

whereas the first waiting state is preceded by the OFF state,

$$w_1(k+1) = b(k)\, OFF(k) = b(k) \cdot \left(1 - x(k) - \sum_{i=1}^{n} w_i(k)\right), \tag{4.10}$$

where $OFF(k)$ denotes the portion of solar devices in the OFF state; recalling that the solar devices can either be ON, OFF, or waiting, namely that x, w_i and OFF sum to one, we substitute $OFF(k) = 1 - x(k) - \sum_i w_i(k)$.

We may be interested in limiting the model's dimensionality, that depends on the cardinality of the waiting states. To this end, we set a fixed number of waiting states, n, and add a self-loop to w_n. Recalling that τ_j represents the transition probability to the ON state, and $(1 - \tau_j)$ represents the probability of remaining in a waiting state, we equip the n-th waiting state, besides the transitions towards ON and OFF states, with a self-loop encompassing the transitions to and from w_{n+i} states. Formally,

$$w_n(k+1) = b(k)\, (1 - \tau_{n-1})\, w_{n-1}(k) + b(k)\ (1 - \tau_n)\ w_n(k). \tag{4.11}$$

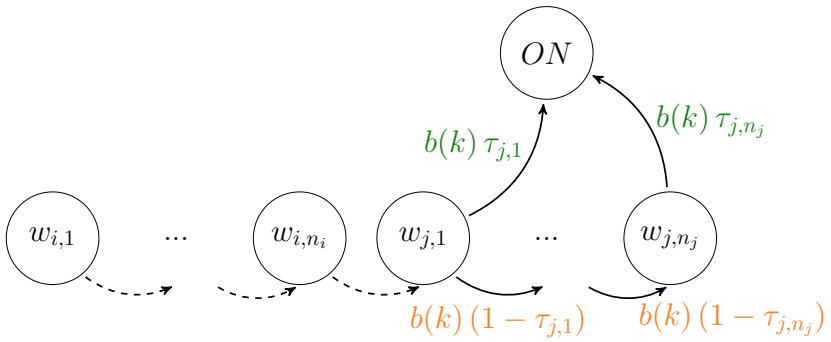

Figure 4.6: Delayed reconnection.

We now bring together Eqs. (4.8), (4.10), (4.9), (4.11) and summarise the dynamics of the Markov chain in Fig. 4.4 as

$$\begin{cases} x(k+1) = (1-a(k))x(k) + b(k)\sum_{i=1}^{n}\tau_i w_i(k), \\ w_1(k+1) = b(k)\left[1 - x(k) - \sum_{i=1}^{n} w_i(k)\right], \\ w_i(k+1) = b(k)(1-\tau_{i-1})w_{i-1}(k), \quad i = 2,\dots n-1, \\ w_n(k+1) = b(k)\left[(1-\tau_{n-1})w_{n-1}(k) + (1-\tau_n)w_n(k)\right], \end{cases} \tag{4.12}$$

offering an analytical representation of a heterogeneous population model.

We have tacitly assumed that the probability distributions of frequency thresholds and time delays are independent. As shown in Chapter 3, the correlation between these two quantities is generally small yet not zero. A more formal design might model the correlation and correct the definition of τ_i $b(k)$ by computing complex integrals of joint discrete-continuous probability distributions. For simplicity we consider independent probability distributions, knowing that the approximation is negligible and this assumption is backed by measurement on real devices.

Reconnection Interval In practice the reconnection interval is usually divided into two sub-intervals: during an early sub-interval no device is allowed to reconnect, whereas during a following sub-interval the reconnection happens stochastically. Recall the reconnection times presented in Chapter 3. Between 0 and 20 seconds, no device reconnects, whereas after 20 seconds the reconnection can be modeled as a random variable. Figure 4.6 depicts the waiting states in such a scenario. Let

us denote the first and the second sub-intervals reconnection probabilities with the subscripts i and j, respectively.

The first sub-interval is modeled by setting τ_i to zero; these states only have two outgoing transitions, towards the next waiting state – depicted with a dashed line, with probability $b(k)$ – or towards OFF – with probability $1 - b(k)$, not represented in Fig. 4.6. We thus model a deterministic delay of n_i time steps.

The second sub-interval is modeled by setting the values of τ_j according to the probability distribution of interest; these states have the three outgoing transitions towards ON, OFF and the next waiting state. Generally we expect a framework composed of a minimum deterministic delay of n_i time steps with no reconnection, and a subsequent interval of n_j time steps when the devices might reconnect. The delay states needed to encompass this scenario is a total of $n_i + n_j$. In conclusion, we can easily model a delayed reconnection splitting the reconnection probabilities as $\tau_i = 0$, for $i = 1, \dots n_i$, and $\tau_j > 0$, for $j = 1, \dots n_j$.

4.5 Simplification of the Delayed Model

A clear drawback of the presented model is the number of states needed to represent the reconnection procedure, which is dependent on the reconnection time itself and on the time discretisation. In a discrete time framework, the number of waiting states is $n = t_M/h$, where t_M is the maximum reconnection time (given in seconds) and h is the time discretisation step (in seconds). As an example, a $t_M = 100$ seconds with a time discretisation of $h = 0.01$ seconds gives $n = 10^4$ states. A model with thousands of states can be simulated or handled by model checking techniques, however is arguably hard to be analytically examined by a human user.

Towards a simplified analysis of the dynamics of the model with delays, we aggregate the n waiting states into a single location (denoted by WAIT in Fig. 4.7), thus representing the portion of devices that are waiting to be reconnected to the main grid. The three states clearly denote the PV system's possible configurations.

Analytically, we associate the aggregated location to a new variable

$$y(k) := \sum_{i=1}^{n} w_i(k), \tag{4.13}$$

symbolising the portion of PV systems which are not yet reconnected, i.e. whose measured frequency dwells within the interval $\mathcal{I}_f$ for a time period insufficient to reactivate. To formulate the corresponding dynamics, we sum the dynamical evolution of $w_i(k)$ from Eq. (4.12) as

$$\begin{aligned} y(k+1) = \;& w_1(k+1) + \cdots + w_i(k+1) + \cdots + w_n(k+1) \\ = \;& b(k)\left[1 - x(k) - \sum_{i=1}^{n} w_i(k)\right] \\ & + b(k)(1-\tau_{i-1})w_{i-1}(k) \\ & + b(k)\left[(1-\tau_{n-1})w_{n-1}(k) + (1-\tau_n)w_n(k)\right] \\ = \;& b(k)(1-x(k)) - b(k)\sum_{i=1}^{n} \tau_i \; w_i(k). \end{aligned} \tag{4.14}$$

The last term $\sum_{i=1}^{n} \tau_i \; w_i(k)$ represents an issue for our framework: in order to offer a closed-loop model, we need a term which is proportional to $y(k) = \sum_i w_i(k)$. However, we cannot easily split terms τ_i from the corresponding $w_i(k)$. Hence, we resort to a cosmetic transform and rewrite $\sum_i \tau_i \; w_i(k)$ as a function of $y(k)$

$$\sum_{i=1}^{n} \tau_i \; w_i(k) = \varepsilon(k) \sum_{i=1}^{n} w_i(k), \tag{4.15}$$

namely

$$\varepsilon(k) := \frac{\displaystyle\sum_{i=1}^{n} \tau_i w_i(k)}{\displaystyle\sum_{i=1}^{n} w_i(k)}. \tag{4.16}$$

introducing a new term $\varepsilon(k)$, which represents the weighted average of the $w_i(k)$. Notice that, by construction, $\varepsilon(k) \in [0,1], \; \forall k$. Equation (4.13) may be written as

$$y(k+1) = b(k)(1-x(k)) - b(k) \cdot \varepsilon(k) \cdot y(k), \tag{4.17}$$

and similarly we may rewrite the first equation of the system in (4.12)

$$x(k+1) = (1-a(k))x(k) + b(k) \; \varepsilon(k) \; y(k). \tag{4.18}$$

Three-state Model After the definition of $y(k)$ in Eq. (4.13) and the introduction of the time-varying term $\varepsilon(k)$ in Eq. (4.16), we are ready to formally study the three-state Markov model shown in Fig. 4.7. The three locations define a dynamical system

$$\begin{cases} x(k+1) = (1-a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1-x(k)) - b(k)\varepsilon(k)y(k), \end{cases} \tag{4.19}$$

where $x(k)$ and $y(k)$ are the portion of population in the ON and WAIT state, respectively. Notice the absence of the OFF portion, which is redundant: it can be subsumed by computing $1 - x(k) - y(k)$ at any time k.

This rewriting presents a nice feature. If we sum $x(k+1)$ with $y(k+1)$ we get

$$x(k+1) + y(k+1) = (1-a(k))x(k) + b(k)(1-x(k)),$$

which is exactly the expression of Eq. (4.4), i.e. the probability of being ON in the delay-free reconnection model. Our modelling is indeed coherent: if we had to separate the OFF state from the ON and WAIT states – *de facto* having a system with two locations, "devices that measure $f \notin \mathcal{I}_f$" and "devices that measure $f \in \mathcal{I}_f$" – we would re-obtain the delay-free model.

This new model is smaller, and prone to explicit analysis. However, in general we do not know the exact value of $\varepsilon(k)$ and thus we need to extrapolate it from measurement of $x(k)$, $y(k)$ or $f(k)$. We can show that, solely knowing the values τ_i, we are able to estimate $\varepsilon(k)$. A sketch of the procedure is as follows. The quantity $\varepsilon(k)$ depends on states $w_i(k)$, which in turn depends on $x(k)$. We also know that the expressions of $x(k)$ and $f(k)$ (a model for the frequency dynamics will be introduced in the next Chapter) are intertwined: from solely the values of $f(k)$ we may estimate $x(k)$, which helps finding the evaluations of $w_i(k)$, hence $\varepsilon(k)$. The full estimation algorithm, starting from the joint model of $x(k)$, $y(k)$ and $f(k)$, is detailed in Appendix A.1.

Remark 2 *From the TSO perspective, the only observed quantity of this modelling framework is the network frequency. The observability and estimation from the*

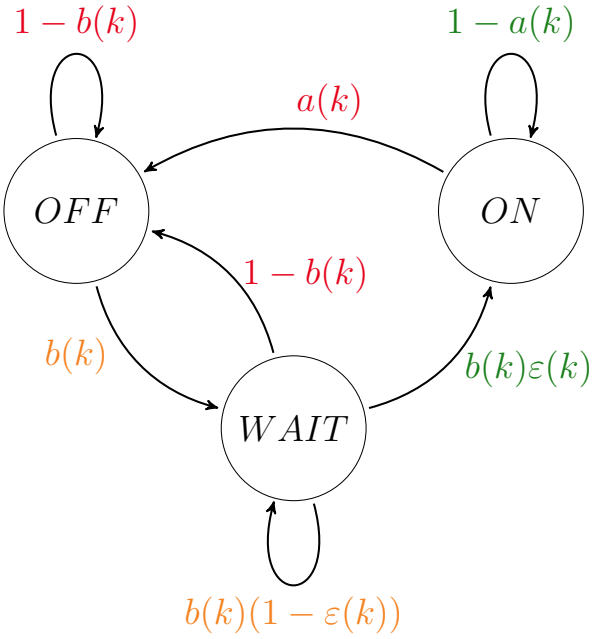

Figure 4.7: Simplified Markov model: states w_i are lumped into state *WAIT*.

value of $f(k)$ are crucial in practice, as they allow reconstructing a truthful image of the aggregated solar power production, hence guaranteeing a precise monitoring of the power grid. □

4.6 Aggregation of Population Clusters

In this Section we propose two approaches to describe a set of populations of PV systems, and analyse their differences.

For simplicity and for brevity, we will consider a population composed by two clusters. Generalisations follow similarly. At first, let us define them separately and let us denote x_1 and x_2 as the probability of being ON of the two clusters, respectively. In the attempt of maintaining the discussion clear and effective, we consider a reduced version of the equations presented above. Specifically, recall from Eq. (4.19) the dynamics of $x(k)$: we will outline our reasoning solely with the mock expression $x(k+1) = a(k)x(k)$. Our analysis can easily be adapted and extended to

the full dynamics of Eq. (4.19). A system of independent clusters may be written as

$$\begin{cases} x_1(k+1) = a_1(k) \cdot x_1(k), \\ x_2(k+1) = a_2(k) \cdot x_2(k). \end{cases} \tag{4.20}$$

Let us define the total power output of the clusters as c and assume we normalise it to 1; we denote c_1 and c_2 as the normalised power output of cluster 1 and cluster 2, respectively, so that $c_1, c_2 \in [0, 1]$ and $c_1 + c_2 = c = 1$. Considering the dynamics of Eq. (4.20), let us define the weighted sum $\bar{x}$ as

$$\bar{x}(k) = c_1 x_1(k) + c_2 x_2(k), \tag{4.21}$$

where $\bar{x}(k) \in [0, 1]$, $\forall k$, by construction. Notice that the heterogeneity term $a(k)$ is kept separated between the two clusters.

Instead of considering the two heterogeneity components as separate clusters, we may want to use a normalised heterogeneity term as

$$c_1 a_1(k) + c_2 a_2(k), \tag{4.22}$$

and study the behaviour of the aggregated dynamics x_T as

$$x_T(k+1) = (c_1 a_1(k) + c_2 a_2(k))\, x_T(k). \tag{4.23}$$

We define the aggregation error $e_A(k)$ as the $\mathcal{L}_1$ norm of the difference between $\bar{x}(k)$ and $x_T(k)$,

$$e_A(k) = |x_T(k) - \bar{x}(k)|, \tag{4.24}$$

and aim at studying its evolution in time

$$\begin{aligned} e_A(k+1) &= |x_T(k+1) - \bar{x}(k+1)| \\ &= |\ (c_1 a_1(k) + c_2 a_2(k))\, x_T(k) - c_1 x_1(k+1) - c_2 x_2(k+1)\ | \\ &= |\ (c_1 a_1(k) + c_2 a_2(k))\, x_T(k) - c_1 a_1(k) x_1(k) - c_2 a_2(k) x_2(k)\ |\,, \end{aligned}$$

where we have used the relations in Eq. (4.23). To ease the notation, define $c_1 a_1(k) := \gamma_1(k)$ and $c_2 a_2(k) := \gamma_2(k)$, so that $e_A(k+1)$ becomes

$$|\gamma_1(k)\,(x_T(k) - x_1(k)) + \gamma_2(k)\,(x_T(k) - x_2(k))| \leq \gamma_1(k) + \gamma_2(k).$$

In order to have a null aggregation error, the aggregated dynamics $x_T(k)$ must be equal, at all times, to the weighted average of the two population clusters, as

$$x_T(k) = \frac{c_1 a_1(k) x_1(k) + c_2 a_2(k) x_2(k)}{c_1 a_1(k) + c_2 a_2(k)}. \tag{4.25}$$

The aggregated dynamics $x_T(k)$ are a useful expression of a clustered population of PV systems, especially in presence of a high number of clusters.

4.7 Experimental Evaluation of the Population Models

In conclusion of this Chapter, we set up rounds of simulations comparing the two population models (the n-waiting-state model, which we refer to as $(n+2)$-state model, in Eq. (4.12) and the three-state in Eq. (4.19)). As a benchmark, we simulate an *explicit* model that we use as a ground truth: it is composed of N models of individual PV systems, where each of the N PV systems has been given four different frequency thresholds (disconnection and reconnection in over- and under-frequency) and a time delay (as a number of time steps the device needs to wait before turning to state ON) between 20 and 60 seconds. This is slightly in contrast with data outlined in Chapter 3, where we show that the delays may be aggregated as a geometric distribution starting after 30 seconds, although this allows for a more rapid comparison of the models. Further, all the values included in these simulations have settled after 20 seconds, thus we do not lose any transitory behaviour. The threshold and reconnection parameters have been generated according to different probability distributions for the population, which are then used in the population models of Eq. (4.12) and Eq. (4.19). We have set $N = 10^6$, whereas the distributions of frequency thresholds as Gaussian, and that for the time delay as geometric, whose details are reported in Table 4.1.

In order to validate our modelling framework, we compare the evolution of $x(k)$ of the $(n+2)$-state model, the three-state model and the explicit model in an open-loop fashion using a controlled frequency signal. To this end, we simulate a disconnection and a reconnection scenario. In particular we are interested in finding possible

Parameters	Distribution
Underfrequency, Disconnection	$\mathcal{N}(49.5, 0.1)$
Underfrequency, Reconnection	$\mathcal{N}(49.5, 0.1)$
Reconnection delays	$\mathcal{G}(0.25)$

Table 4.1: Parameters of the distributions.

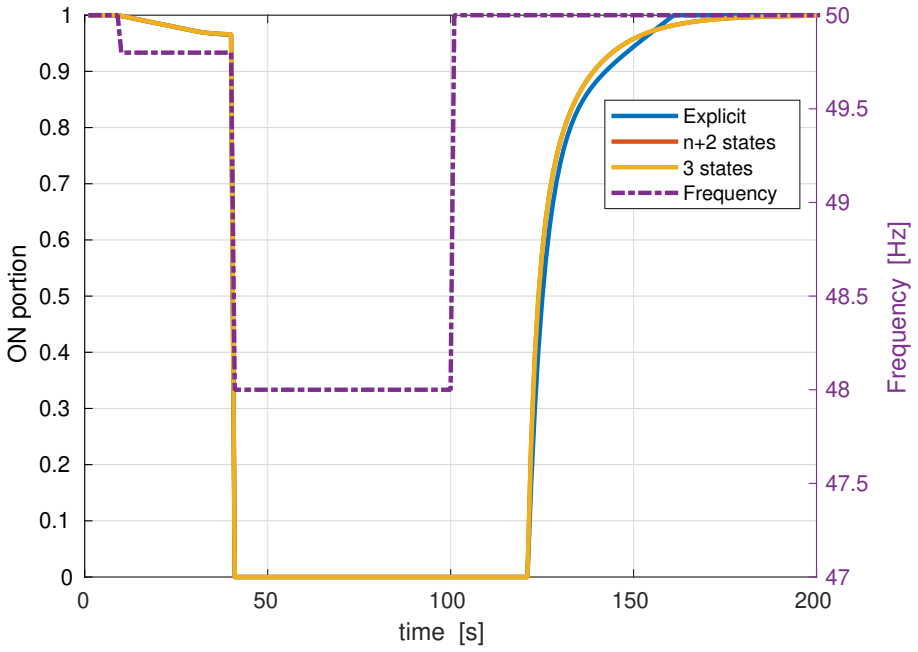

Figure 4.8: Response of the explicit model (blue), the $(n+2)$-state model (red), and for the 3-state Markov model (yellow) to the frequency signal (purple, dashed).

discrepancies among the models whenever a portion of PV systems disconnect or reconnect. We ought to recall that models in Eqs. (4.12), (4.19) are frequency dependent in the $a(k)$ and $b(k)$ terms by the relationship in Eqs. (4.5) and (4.6).

Figure 4.8 shows the power response of the explicit model, as well as the $(n+2)$-state and 3-state models, together with the injected frequency signal. The frequency evaluates to 50 Hz for the first 10 seconds, 49.5 Hz between 11 and 40 seconds (causing a partial disconnection of the population), 48 Hz between 41 and 100 seconds (resulting in a total disconnection of the population), and increases back to 50 Hz from 101 seconds on.

First, we note the perfect overlapping between the $(n+2)$-state and the 3-state models, as they are analytically equivalent. We then show the small discrepancies

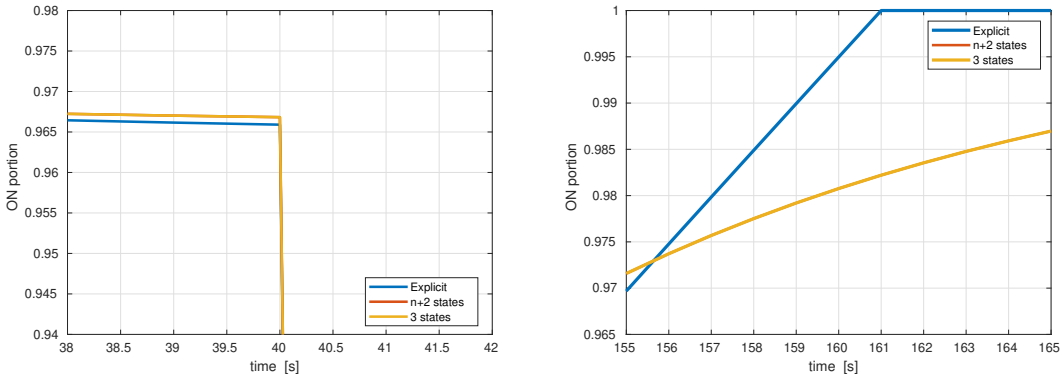

Figure 4.9: Zoom at the end of the first frequency drop (left) and at the end of the reconnection stage (right).

between the explicit and the two abstract models in Fig. 4.9. Specifically, Fig. 4.9 (left) zooms in on the difference between the $x(k)$ signals after the first disconnection period, which evaluates to around 10^{-4}, whereas on the right depicts the gap at the end of the reconnection process, evaluating at around 0.017. These modelling gaps derive from our approximation: disconnection and reconnection thresholds are defined over a *continuous* domain, whereas the explicit model is characterised by a large-but-finite number of devices. Further, the reconnection of the abstract models shows an exponential rate – in view of the term $(1 - \tau_n)w_n$ in the last waiting state in Eq. (4.12), whereas the explicit model displays a less smooth behaviour. In the explicit model all devices deterministically reconnect after 40 seconds, thus reaching $x(k) = 1$ after a finite time horizon (see Fig. 4.9 (right) where at 141 seconds $x(k)$ becomes 1); the abstract models show instead an exponential convergence to $x(k) = 1$, thus resulting in a continuously decreasing gap between the power signals. From the practical point of view, these discrepancies entail an error in percentage smaller than 2% on the population power output.

In terms of computational complexity, the explicit model is much heavier to simulate: a single run takes around 20 seconds against 0.03 seconds of the population models, resulting in a speedup of three order of magnitude. The explicit model is much heavier to run because at each time step each solar unit compares the current value of the network frequency with its frequency thresholds (and delay counter), and disconnects (or reconnects) accordingly. The Markov

model, on the other hand, computes a single integral per time step, relieving the computational effort significantly.

4.8 Concluding Remarks

In this Chapter we have presented a modelling framework for the aggregation of a heterogeneous population of photovoltaic systems, which also features a delayed reconnection procedure. We interpret these models as dtMCs with time-varying transition probabilities, which depend on the value of the network frequency. We have proposed a $(n+2)$-state model, where n is the (maximum) number of time instants needed to reconnect and the other 2 states represent the connected (ON) and disconnected (OFF) status. We have shown that this model can be reduced into an analytically equivalent three-state model.

We have provided experiments comparing the two dtMC models against an explicit system, where N photovoltaic systems were singularly simulated: a disconnection and reconnection scenarios have shown the computational efficiency (the dTMCs are three order of magnitude faster to simulate) and reliability (maximum absolute value error of 2%).

The dtMC models are limited by the assumption of a constant power output by the PV systems. Future efforts may extend this framework to support variable power outputs, modelling uncertainties such as weather, occlusions, within the stochastic setting [191].

Insofar we have dealt solely with the aggregated solar power. The next Chapter introduces the frequency dynamics, which will be coupled with the solar aggregation, and applies a proportional control scheme to transform the household PV systems into an ancillary service for extreme events.

5

Network Dynamics and Decentralised Control

Contents

One of the goals of the dissertation, after proving the soundness of model design, is asserting the impact of a large population of PV systems connected to the electricity network. As mentioned in Chapter 2, the addition of renewable energy sources reduces the amount of global inertia of the power grid: we analyse the effects of a smaller inertia coefficient in terms of frequency oscillations, especially focusing on the time instants after a normal incident[1].

To this end, we devise a closed-loop model intertwining the solar power output, derived from the models in Chapter 4, with the power grid model presented in the following. We evaluate the closed-loop characteristics varying the amount of PV

[1]The definition of normal incident is given in Section 3.4.3.

systems in the network. We propose a sensitivity analysis of its inertia coefficient and evaluate the step response in terms of the maximum oscillation and settling time. By and large, lowering the amount of inertia in the power grid modifies the poles of the system, increasing both their negative real part and their imaginary part. Whilst the first is a desirable feature, the second carries concerns on the electric network stability. We then study the PV systems switching after a generation loss incident. Solar devices are subjected to safety rules that lead to a chain of disconnections, jeopardising the reliability of the grid. Certain distributions of working intervals may prevent such events to happen, acting as an added inertial system.

5.1 Power Grid Dynamics

The model for the electric network that we utilise in the following has been provided by our research partners at RTE, in the form of a Simulink [134] block diagram. It is the same network model used in the ENTSO-E report [41] regarding the effect of dispersed generation in the European synchronous area. This model is effectively used to represent the network in a simple, yet effective, manner. It follows the guidelines of *Target Performance* outlined in [192], specifically designed to study a contingency scenario. In particular, it has been validated against field measurements and behaviour of the currently available figures of frequency dynamics in Continental Europe. The principles of the the ENTSO-E Policy on Load Frequency Control [27] are taken as basis to ultimately construct the network model and, where uncertainties subsist, realistic but pessimistic assumptions are leveraged, again according to the policies in [27, 190, 192].

The network model represents a geographically widespread power grid with several energy sources and load characteristics. Its parameters are calibrated in order to represent an average behaviour over the whole network. As a typical framework in power systems, we outline and analyse a *continuous-time* model of the electric network. Note that the continuous-time feature of the grid model is in contrast with our discrete-time population models: we first present a description and the analysis of the network in continuous-time, and later introduce a time discretisation

when connecting it to the aggregated model of the population of PV systems. The discrete-time version will be employed in the experiments of Section 5.4.

From the power grid block diagram, we derive the corresponding transfer function. We obtain a second-order transfer function, denoted $G(s)$, which relates the photovoltaic power output (its input) to the network frequency (its output). In particular, the input signal is the power *deviation*,

$$\Delta P_{PV}(t) = P_{PV}(t) - P_0, \tag{5.1}$$

defined as the difference between the current power output $P_{PV}(t)$ and a reference solar power value P_0. Naturally, we define P_0 as the solar power output that is related to the nominal frequency f_0. The output of $G(s)$ is the frequency change

$$\Delta f(t) = f(t) - f_0, \tag{5.2}$$

representing the difference between the current value of frequency $f(t)$ and the reference value f_0. The transfer function analytically results in

$$G(s) = \frac{s+1}{T_L s^2 + (k_a + T_L)s + (k_a + k_{PU})}, \tag{5.3}$$

where the quantity T_L represents the time to launch (related to the inertia of the system), whereas k_a and k_{PU} are gains of the load self regulation and of the primary control unit, respectively.

These three quantities encompass a more realistic approximation of the network dynamics over a short timescale (seconds). Specifically [193], (i) the transfer function includes k_{PU}, representing the primary control of the grid, which reacts to frequency deviations; (ii) parameter k_a encodes a variation of load in response to frequency deviations. The self regulation of the load is the sensitivity of consumers demand to variations in the system frequency (e.g., a decrease of the frequency results in a decrease of the load), and is generally expressed as %/Hz; (iii) finally, the time to launch is the time that a device needs to accelerate from zero to the working speed.

The network model parameters are chosen according to the guidelines in the Appendix of [27]: the self-regulation of the load k_a is assumed to be 0.01 s, the

Parameter	Reference Value
k_a	1 % / Hz
k_{PU}	15 $\cdot 10^3$ MW/Hz
T_L	10 s

Table 5.1: Parameters' value for the power grid model.

average network primary control gain equal to $k_{PU} = 15000$ MW/Hz, and a time to launch equal to 10 seconds, as reported in Table 5.1. Notice that we assume the lowest value of k_{PU} and T_L suggested in [27] and [192], to simulate a fast network with little primary control, i.e. the worst case scenario.

Figure 5.1 depicts the block diagram representation of $G(s)$. Quantities $f(t)$ and $f_{pu}(t)$ represent the frequency and the frequency per unit[2], respectively. $P_m(t)$ denotes the power coming from traditional machines, whereas $P_{PV}(t)$ is the solar power. The PV model is represented by the orange block.

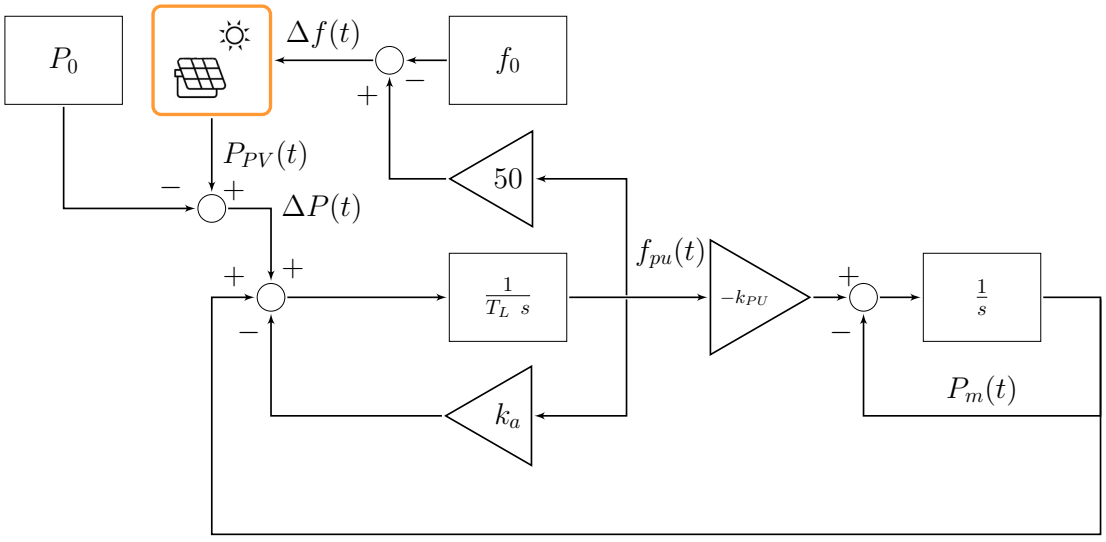

Figure 5.1: Network block diagram representation.

Renewable-dependent Dynamics Models of the national or European grid [10, 194, 195] are developed assuming the presence of synchronous machines. In that context, the total moment of inertia associated to synchronous machines is considered to be constant. Whilst this assumption is justified when conventional

[2]A *per unit* quantity is denoted as the value of the quantity of interest divided by its reference value. In this case, $f_{PU}(t) = f(t)/50$.

power sources are connected to the power grid, the presence of renewable energy sources render the concept of inertia harder to grasp and to handle. For instance, the notion of *synthetic inertia* [196] is introduced when wind generation is under study (here potential energy is stored in their blades). However, evidently PV systems have no moving parts, which implies they come with no mechanical inertia at all: if solar irradiance suddenly subsides, a solar device immediately stops generating electricity. The opposite is also true: whenever the sun shines onto a PV system, power is immediately generated.

As a result, when the penetration of solar power sources in a network is significant, their presence must be accounted for, both in terms of power production and of reduced inertia. We include this contribution in the transfer function $G(s)$, which we regard as a function of the amount of the *conventional power* generated by synchronous machines. In particular, we assume a linear relation between the time to launch T_L and the amount of conventional power generated in the network. Intuitively, the more conventional generators are present in the network, the more inertia governs the frequency evolution and the larger the starting time is. We formalise this idea as

$$T_L = \frac{CP}{k_T},$$

where CP represents the conventional power (expressed in MW) and k_T is a constant value that defines the ratio between T_L and CP. The report [27] evaluates $T_L = 10$ seconds for a network with S equal to $150 \cdot 10^3$ MW, (150 GW). Assuming that with no conventional sources the value of T_L is null, we obtain the value of $k_T = 15 \cdot 10^3$, considering the network load is expressed in MW. Naturally, the quantity CP dwells within 0 and S, the total load of the network.

We substitute T_L in the expression of $G(s)$, obtaining

$$G(s, CP) = \frac{(s+1)}{\frac{CP}{k_T}s^2 + \left(k_a + \frac{CP}{k_T}\right)s + (k_a + k_{PU})}. \tag{5.4}$$

In the following we will compute the discrete-time equivalent of $G(CP, s)$ to model the frequency deviations of the electricity network. We bear in mind that

this modelling choice brings a few limitations: as mentioned, it does not account for local changes in frequency, nor for the topology of the network, aiming instead at modelling a large grid network. In addition, the main limitation of our network model is the homogeneous description of all energy sources in the network. This particular model accounts for only one equivalent synchronous generator to represent all the synchronous (but dynamically slightly dissimilar) power plants: different power plants, say nuclear, coal, hydroelectric, have different characteristics and T_L coefficients, which are disregarded in this aggregated design. Finally, we have assumed that PV systems bring no inertia to the grid. Other renewables, e.g. wind power, may have a different impact and need to be carefully modelled. More detailed models exist (among others, dynamic tools as Eurostag [197] provide more complex designs), however they come with higher computational costs.

Nevertheless, the formulation of Eq. (5.3) is in practice utilised by TSOs – in particular by our partners at RTE – and it is considered a reliable model for the continental electricity grid, as the frequency can be considered homogeneous across countries. For this reason, we will use Eq. (5.4) for the following analysis of a renewable-dependent network.

5.2 Influence of Solar Penetration: Root Locus Analysis

This section studies how the ratio CP/S influences the stability of the electricity network. In particular, we choose a technique named *root locus analysis* [198] that is common in control theory.

The performance – stability and transient behaviour – of a (feedback) system is directly related to the location of the roots of the transfer function. The root locus technique is a graphical and powerful tool for the analysis and design of controlled systems. The root locus plot shows the locations and how the roots of a transfer function move as a single parameter varies. In fact, the root locus provides a sensitivity measure of the roots of the system to the variation of a parameter of interest.

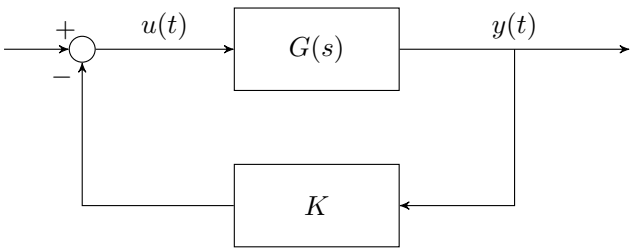

Figure 5.2: Schematic representation of the feedback loop.

To apply the analysis, we consider the feedback interconnection of $G(s)$ with a hypothetical proportional control – to be later presented in this Chapter – under different ratios of CP/S. Practically speaking, we study the locations of the poles of the feedback system, namely the roots of the denominator of the overall transfer function. For a system shown in Fig. 5.2, the overall transfer function results in

$$T(s, CP) = \frac{G(s, CP)}{1 + KG(s, CP)}, \tag{5.5}$$

where K represents the controller's proportional gain. Hence, the closed-loop poles are the roots of the equation $1 + KG(s, CP) = 0$, or $G(s, CP) = -\frac{1}{K}$.

The root locus plot of the feedback system in Eq. (5.5) shows the locations, in the frequency domain, of the poles while varying the parameter K. By visualising the position of the poles of this system, we are able to infer the behaviour of the power grid following a change in its input signal, i.e. the network frequency. Three characteristics are fundamental for this analysis: the stability of the poles and their real and imaginary part. Stability requires that the system's poles dwell in the left half of the complex plane, i.e. have a negative real part. In practice, a stable pole guarantees that the system (the power grid, in our case) will eventually stabilise its output (the network frequency) to a finite value after an input (a generation loss in case of an incident) is entered. Generally speaking, the real part of a pole influences the convergence time, namely how quickly the system stabilises around a value, whereas the imaginary part is related to the oscillations of a system: a large imaginary component implies wide oscillations of the output signal. Ideally, the feedback system's poles should aim at a large negative real part,

to accelerate the convergence, and a small imaginary part to keep oscillations at a minimum. We study the effect of a population of renewables on the network's poles by considering three conditions: (i) as a reference, we first outline the poles' location with $CP/S = 1$, i.e. with no renewables in the grid; (ii) we then analyse the poles with 20% renewable penetration, i.e. the ration $CP/S = 0.8$, that represents the current (or the short term objective) penetration in many European countries; (iii) we finally consider a longer term objective and a more significant renewable penetration as $CP/S = 0.5$, with 50% of renewable sources. As mentioned in Chapter 2, predictions foresee that 63% of the electricity demand will be met by renewable sources in the next few years [5].

Figure 5.3 shows the root locus under the three different conditions $CP/S = 1$, i.e. no renewable power in the network, $CP/S = 0.8$, and $CP/S = 0.5$, considering a global network load $S = 220$ GW. First of all, we highlight that the function $T(s, CP)$ is stable in every configuration, i.e. its poles are in the negative real half plane. The increase of K stabilises the interconnection of Eq. (5.5) by moving the poles towards a location with a higher real negative part and smaller imaginary part. On the other hand, the decrease of CP/S increases both the real negative part and the imaginary part: this translates to a faster convergence to stability but wider oscillations. The latter consequence represents an issue for the power grid: wider oscillations of the frequency may reach values that cause load shedding.

Another graphical tool to assess the stability properties of a system is the step response. We consider $G(s, CP)$, inject a step signal and evaluate the output in the time domain. We compute the step response considering three CP/S ratios, i.e. 0%, 20% and 50% renewable penetration in the grid, as depicted in Fig. 5.4. The effect of renewables on the network dynamics is quite clear: the more renewable power, the faster and more oscillatory the system becomes, leading to overshoots (notice the signal with 50% renewable power). In fact, the step response with 50% solar penetration shows a rise time - the time needed to go from 10% to 90% of the final value - of 2.22 seconds; with 20% solar penetration the response is much smoother, with a rise time of 4.30 seconds; finally, the step response without renewables results

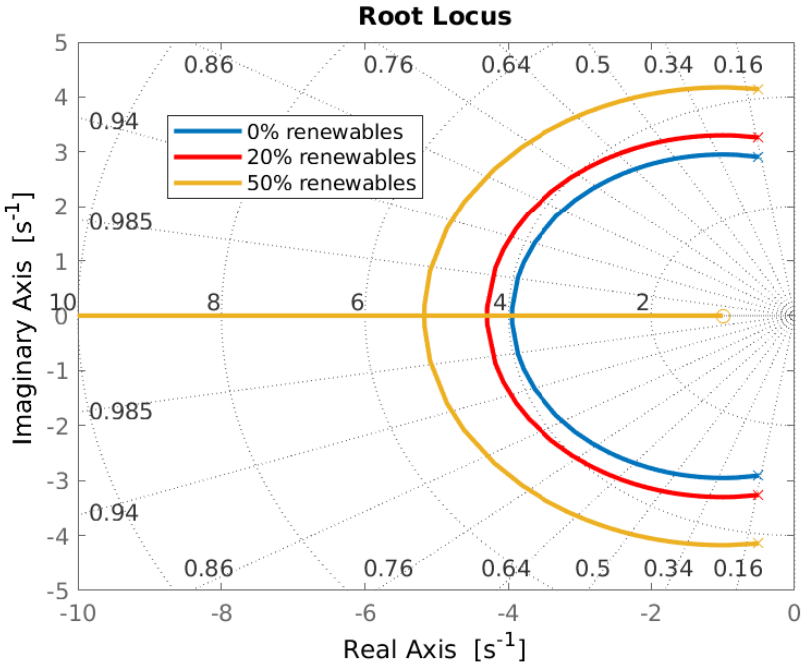

Figure 5.3: Root locus with 0% (blue), 20% (red) and 50% (yellow) renewable power.

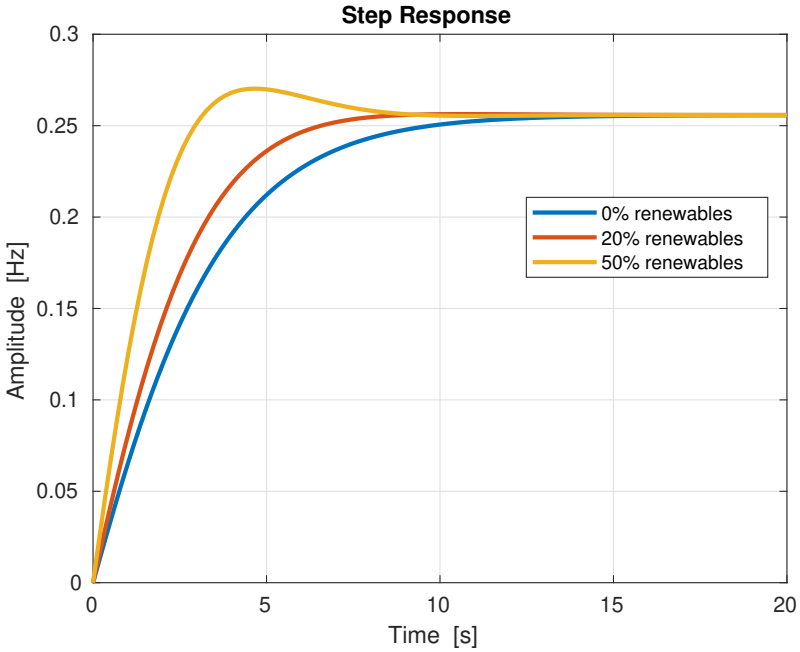

Figure 5.4: Step response in presence of three different solar penetrations.

the smoothest, with a rise time of 5.92 seconds. Most importantly, oscillations could eventually cause issues in the electric grid setting: the grid operators should avoid wide variations of the frequency signal, as this brings the network to its operational limits and potentially causes load shedding. On the other hand, Fig. 5.3 illustrates that also the real part increases in absolute value, meaning a faster convergence. As a measure of convergence we compute the settling time [198]: roughly speaking, it is the time elapsed from the application of a step input until the system's output has entered and remained within a specified error band of the output reference value. For this study, we use an error band of 5%, namely we account for the time at which the output signal enters the interval $[0.95, 1.05]$. Increasing inertia-less sources reduces the settling times: we report times of 7.9 seconds, 5.6 seconds, and 5.4 seconds for the network without renewables, with 20% and 50% renewables respectively.

Overall, our analysis suggests that the uncontrolled addition of renewable, inertia-less sources to the power grid may harm and jeopardise its operations. A more oscillatory and faster response is a direction that traditional control architectures deployed in power systems cannot cope with. As mentioned in Chapter 2, control actions for the power grid are designed assuming conventional power sources, hence having large time constants. New power sources need novel control strategies, able to cope with a faster and more oscillatory network. Let us consider, as an example, a network with 20% renewables – a percentage that is commonly achieved nowadays: compared to a conventional network, its rise time is approximately 27% shorter.

Remark 3 *Notice that the step response in Fig. 5.4 is hardly comparable to the following closed-loop simulation (see, e.g., Fig. 5.6 - 5.8). When the frequency reaches values close to the disconnection thresholds (a drop of 0.2 Hz), PV systems start to disconnect: in practical terms, this is similar to summing an additional step input to the analysis. A value that remains similar before the disconnection of the solar devices is the rise time: a network with lower inertia show a quicker response, and the addition of disconnecting PV systems deteriorates the overall network stability.*

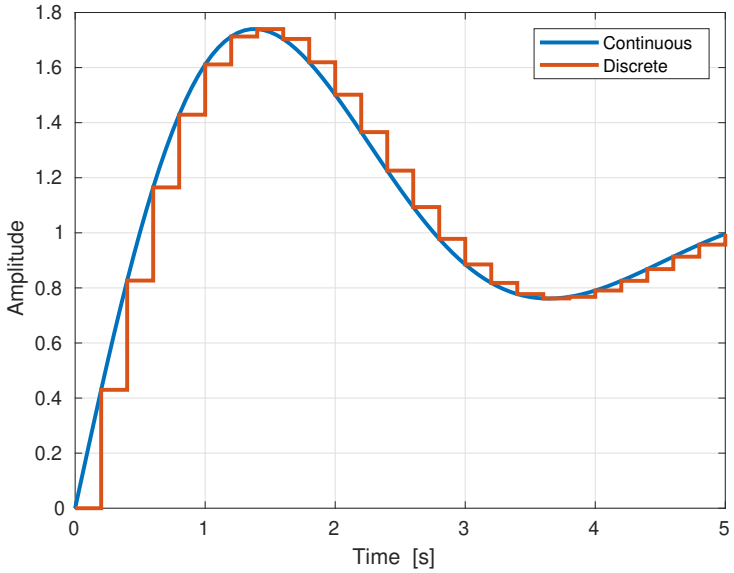

Figure 5.5: Time discretisation with step invariance response.

5.3 Closed-Loop Dynamics

We now connect the population of PV systems with the electric network in a feedback fashion, as shown in Fig. 5.1. For simplicity and clarity, let us focus on a network with a constant amount of conventional power, and refer to the transfer function as $G(s)$, rather than as $G(s, CP)$. Formally, we have that

$$\Delta f(t) = G(s)[P_{PV}(t) - P_0]. \tag{5.6}$$

We now translate $G(s)$ from Eq. (5.4) into its discrete version $G(z)$ via the *step response invariance method* [199]. This approach guarantees perfect matching of the continuous-time signal with the discrete-time signal at sampling times, as depicted in Fig. 5.5. We are thus able to preserve the exact network frequency signal after a step-like input as the one derived from a normal incident. Formally we obtain a second-order transfer function as,

$$G(z) = \frac{\beta_1 z + \beta_2}{z^2 + \alpha_1 z + \alpha_2}, \tag{5.7}$$

where α_1, α_2, β_1, β_2 are constants that depend on the chosen sampling time, and where z indicates the discrete time variable. In this work we select the same

sampling time as the one of the inverters (0.2 seconds). We can write the discrete-time equivalent of the previous equation as

$$\Delta f(k) = G(z)[P_{PV}(k) - P_0], \tag{5.8}$$

where $\Delta f(k) = f(k) - f_0$, and f_0 represents the nominal value of the network frequency, whereas $f(k)$ is the value of the frequency at time k; $P_{PV}(k)$ is the power output of the population of PV systems at time k and (as above) P_0 the power produced when $f(k)$ is at the equilibrium f_0. We set

$$P_{PV}(k) = \bar{P}Nx(k),$$

where $\bar{P}$ is assumed to be the constant power output of a single PV system, and N is the total number of systems. As such, we obtain a proportional relation between the total power output and the portion of PV systems in the ON mode. This assumption simplifies the formal analysis on the feedback models in Eqs. (5.9) and (5.10).

Let us denote $\Delta x(k) = x(k) - x_0$, where x_0 represents the portion of active PV systems when $f(k)$ is at the equilibrium f_0. We now embed the frequency evolution in Eq. (5.8) into the dynamics of the population models. Let us present the overall system considering the $(n+2)$-waiting-state model in Eq. (4.12), that holds

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \Delta\alpha_2 f(k-1) + \gamma_1 \Delta x(k) + \gamma_2 \Delta x(k-1), \\ x(k+1) = (1 - a(k))x(k) + b(k) \sum_{i=1}^{n} \tau_i w_i(k), \\ w_1(k+1) = b(k) \left[1 - x(k) - \sum_{i=1}^{n} w_i(k) \right], \\ w_i(k+1) = b(k)(1 - \tau_{i-1}) w_{i-1}(k), \quad i = 2, \dots n-1, \\ w_n(k+1) = b(k)[(1 - \tau_{n-1}) w_{n-1}(k) + (1 - \tau_n) w_n(k)]. \end{cases} \tag{5.9}$$

A simpler overall model is returned if we embed the frequency dynamics into the three-state model in Eq. (4.19) as

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \alpha_2 \Delta f(k-1) + \gamma_1 \Delta x(k) + \gamma_2 \Delta x(k-1), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)). \end{cases} \tag{5.10}$$

Notice that we have introduced constants $\gamma_i = \bar{P} \cdot N \cdot \beta_i$, $i = 1, 2$. The models described by Eqs. (5.9) and (5.10) are utilised in the following case studies. The feedback connection of the Markov chains and the process $G(z)$ can be shown to be locally asymptotically stable around the nominal frequency. The proof can be found in Appendix A.2.

5.4 Testing Distributions and Load Shedding Relation

Chapter 4 has outlined the modelling framework and has tested in an open-loop fashion the reliability of the proposed Markov population models against an explicit, single-device crafted model. This section instead uses the closed-loop dynamics offered in Eq. (5.10) – which considers the three-state population model coupled with the frequency dynamics – to evaluate the consequences of a normal incident (cf. Section 3.4) in a power grid with significant penetration of PV systems. In particular, we study the relationship between the distribution of frequency thresholds and the event of load shedding after an incident: we show that as the solar penetration grows, the chance of load shedding increases. However, heterogeneity can be exploited to enhance the network reliability and its resilience to incidents, much as we have argued earlier for stability analysis.

In line with the ENTSO-E requirements [27], we consider a network with demands of $S = 220$ GW and 440 GW, which model a low demand and high demand, respectively. We simulate an infeed loss incident of 3 GW and test the response of the network under different circumstances. On the contrary, as noted in [15, 41], the load loss incident has practically no load shedding risk during normal operations: therefore, this test is omitted. We are interested in testing the worst-case setup for the disconnection of devices: namely, we stress the network and attempt to attain the lowest value that $f(k)$ can reach in the electricity grid. To ensure this, the time-frame of the power loss tests is set to 20 seconds, which prevents any PV system from reconnecting – the reconnection of devices can only increase $f(k)$, adding production to the power imbalance.

In order to simulate a power generation loss, we inject a negative step in the frequency loop corresponding to the incident, similarly to the tests in Section 4.7. The simulations use a reduced inertia parameter that is proportional to the amount of solar penetration in the grid. To add plausibility to our case studies, the unpredictable solar power output is encompassed via additive zero-mean Gaussian noise ω_P. Furthermore, the natural noisy nature of the network frequency is described similarly with an additive Gaussian signal ω_f. Formally, we add ω_f to the dynamics of $\Delta f(k+1)$ and ω_P to the dynamics of $P_{PV}(k+1)$ in (5.10), as

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \alpha_2 \Delta f(k-1) + \beta_1 \Delta P_{PV}(k) + \beta_2 \Delta P_{PV}(k-1) + \omega_f(k), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)), \\ P_{PV}(k) = \bar{P}Nx(k) + \omega_P(k), \end{cases} \tag{5.11}$$

resulting in a stochastic hybrid system (cf. Section 2.2.2).

We further set the power production of a single PV system P_{MAX} to 3 kW. The variance of ω_P is set to 1% of P_{MAX}. The variance of ω_f is computed from network frequency data [189] using a Maximum Likelihood Estimation technique – outlined in Section 3.4 – the obtained value is equal to 0.025 Hz^2, as reported in Eq. (3.2). As we shall see, the injected noise has a small amplitude compared to the frequency and power signals, and will be negligible in the case of the 440 GW network. Time delays are modelled in accordance with disconnection regulations and the discussion in Section 3.3. Whilst the minimum and maximum reconnection delays are handled deterministically, the actual population reconnection times are modelled via a geometric distribution, as discussed previously.

According to requirements in [27], we set a boundary for the frequency value of 49.2 Hz: in real setups, if the network frequency trips below 49.2 Hz, an automatic procedure of load shedding (see Section 2.1.4) is activated. In our simulations, we check if the network frequency trips below this critical value, and assume that dedicated control systems activate a load shedding procedure: we therefore stop the simulation, and disregard the modelling of the load shedding procedure.

% PV	Average	Variance	Load Shedding
10	49.5	0.05−0.25	0.15−0.25
20	49.5	0.05−0.25	0.10−0.25
40	49.5	0.05−0.25	0.05−0.25
$\sim 10\%$	$49.8 - 49.0$	1−5	no

Table 5.2: Test results for 3 GW power loss scenario over 220 GW network load.

% PV	Average	Variance	Load Shedding
10	49.5	0.05−0.25	no
20	49.5	0.05−0.25	0.15−0.25
40	49.5	0.05−0.25	0.15−0.25
$\sim 5\%$	$49.8 - 49.0$	1−5	no

Table 5.3: Test results for 3 GW power loss scenario over 440 GW network load.

Simulations are implemented in MATLAB. The grid frequency is sampled at a rate of 0.2 seconds, consistently with the inverters' parameters and the requirements introduced in [120]. The study focusses on the consequences of an incident in a time frame of a few seconds, hence the simulation time is set to 20 seconds. After this time interval, we assume the primary and secondary network controls have kicked in and shall stabilise the signal $f(\cdot)$ around its nominal value f_0 (cf. discussion in Section 2.1.3).

Different thresholds distributions over $\mathcal{I}_f$ have different outcomes and bring to different conclusions about heterogeneity. We then test two scenarios concerning population thresholds: a *Narrow Interval Scenario*, encompassing a limited working frequency interval and Gaussian distributions, and a *Composite Scenario* with several χ^2 distributions. As mentioned in Section 3.3, Gaussian distributions accurately model inverter measurement noise, whereas χ^2 can be used to define a minimum performance setting, in terms of a minimum working interval.

Tables 5.2 and 5.3 refer to network load scenarios of 220 and 440 GW, respectively. They show the results of tests as a function of the percentage of solar penetration in the network, and of the mean and variance of the considered distributions. We report the occurrence of load shedding in the *Load Shedding* column of the

Tables, along with the value of the distributions' variance. In other words, the *Load Shedding* column indicates which values of variance cause the load shedding procedure to activate. In the following we describe the outcomes of the two scenarios of thresholds distribution.

Narrow Interval Scenario: In the *Narrow Interval Scenario*, we consider three percentages of penetration of solar energy production in the network: 10%, 20%, 40% (first three rows of Tables 5.2 and 5.3). We assume that the heterogeneity over the frequency intervals that model the reconnection/disconnection thresholds – in terms of ageing, manufacturing, performance deterioration, sensor noise, etc. – is well described by Gaussian distributions, which we consider over different values of mean and variance.

Thresholds average value is set to 49.5 Hz (underfrequency), resulting in a 500 mHz band around the nominal frequency of 50 Hz. We investigate five values of variance, with values ranging from 0.05 to 0.25 Hz^2 with a step of 0.05. Simulations show, both in the 220 GW and 440 GW scenario, that larger values of the variance cause a higher number of PV systems to disconnect, resulting in load shedding. In the 440 GW network load with 10% solar penetration scenario, no value of variance was sufficient to trip the frequency below 49.2 Hz (cf. first row of Table 5.3). Clearly, when the network load is higher, the reliability of the network itself is enhanced, having a higher time-to-launch and therefore a larger inertia.

Figure 5.6 shows the frequency drop in a 220 GW network load scenario, with 10% solar penetration (cf. first row of Table 5.2). Different values of the variance (between $0.15 - 0.25$ Hz^2) result in different numbers of PV systems to disconnect, potentially resulting in load shedding. Similar outcomes are shown in Fig. 5.7, where a network load of 440 GW and solar penetration of 20% are considered (cf. second row of Table 5.3).

Composite Scenario: We slightly modify the settings used above to achieve a more faithful description of the heterogeneity of the grid. We employ the data from the European Dispersed Generation Report [41], that are outlined in Chapter 3. For completeness, we additionally resume the so-called *power at risk* here in Table 5.4. We assume that the PV population is divided in four groups, each with different under- and over-frequency limits. A group characterises the smallest working interval that complies with regulations, and different groups comply with different regulations. A given PV system must have a working frequency interval $\mathcal{I}_f$ that is wider than the characteristic interval of its group: as an example, PV systems installed with the 49.5 Hz underfrequency limit can have a threshold $f_{uf} \leq 49.5$. Similarly, PV systems are allowed to have $f_{of} \geq 50.2$ in the 50.2 Hz overfrequency case. The probability distribution that best describes this scenario is the χ^2 distribution. Note that, in contrast with the previous case study, the χ^2 distribution is not symmetric around its average value. The contribution in energy production from each group is on the third row of Table 5.4, whereas the fourth row shows the partition size in percentage terms. The total amount of energy production from the entire population accrues to approximately 10% of solar penetration in the 220 GW network, and to 5% in the 440 GW network: see the last row of Tables 5.2 and 5.3, respectively.

Similarly to the previous case studies, we test five values of variance, ranging from 1 to 5 Hz^2 (with a step of 1), and how they possibly result to load shedding (cf. low end of values of variance for the 220 GW case on bottom row of Table 5.2). Note that increasing variance leads to distributing the population thresholds away from the nominal frequency. Indeed, simulations show that increasing the variance enhances the reliability of the grid, as depicted in Fig. 5.8: whilst no scenario results in load shedding, a higher heterogeneity smooths the frequency response.

Remark 4 *In practical terms, it is interesting to emphasise that new European Union regulations [120] allow for a broader interval of frequency values for inverters than in the past: newly manufactured solar inverters should work within interval* $\mathcal{I}_f = [47.5, 51.5]$ *Hz. Based on the developed models, as demonstrated above and*

Underfrequency threshold	49.8	49.7	49.5	49.0
Overfrequency threshold	50.2	50.2	50.2	50.2
Energy production [MW]	154	11500	1508	4000

Table 5.4: PV population *at risk* divided according to different frequency thresholds.

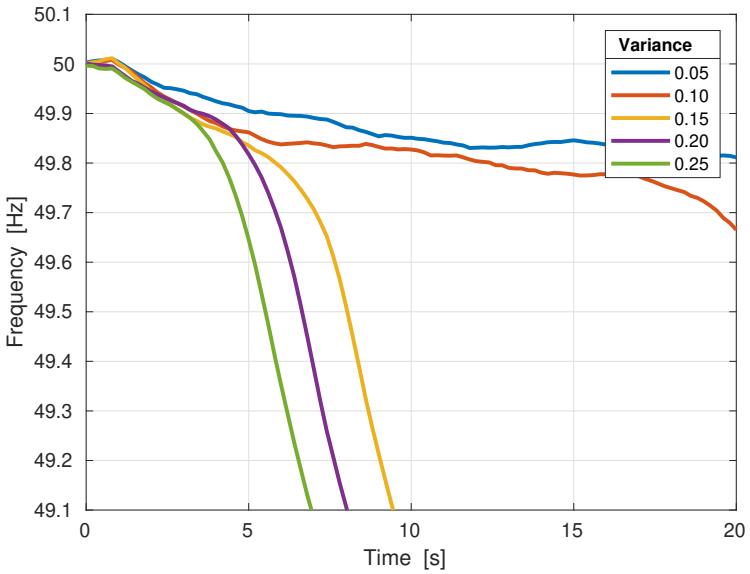

Figure 5.6: Frequency response after a 3 GW infeed loss in a 220 GW network with 10% solar penetration.

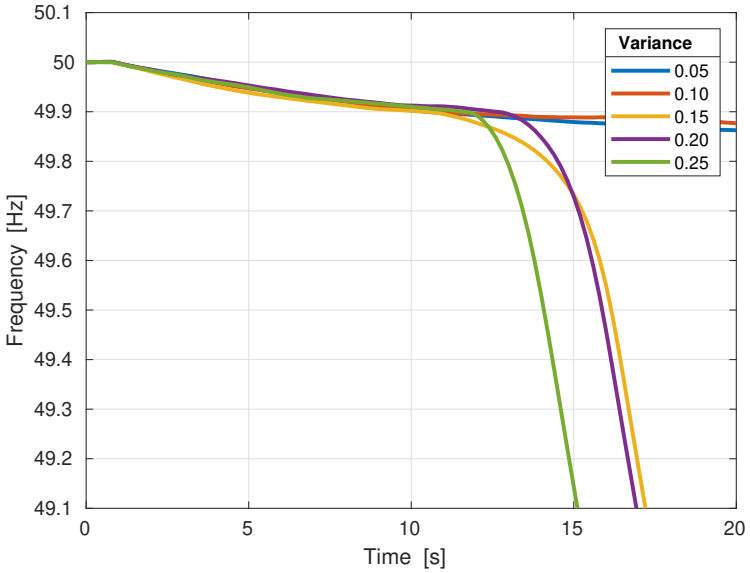

Figure 5.7: Frequency response after a 3 GW infeed loss in a 440 GW network with 20% solar penetration.

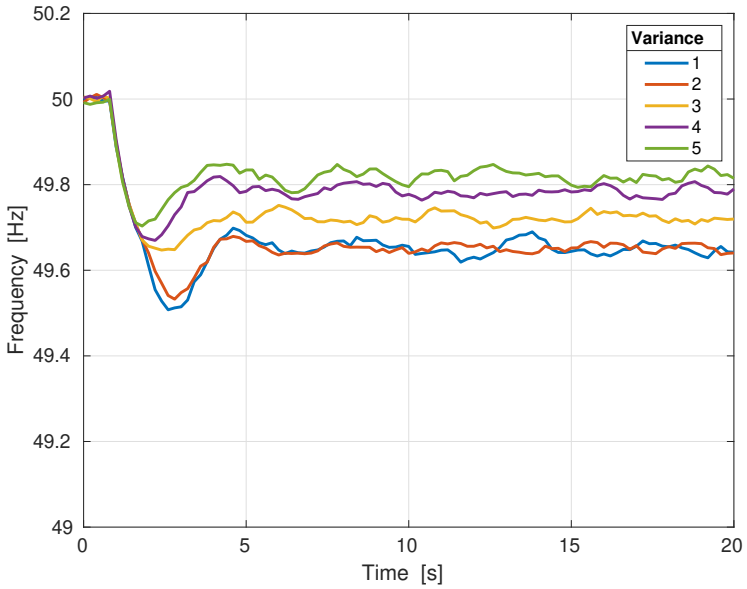

Figure 5.8: Frequency response after a 3 GW infeed loss in a 220 GW network with approximately 10% solar penetration.

further discussed in [15], this larger frequency interval is likely going to contribute to the reduction of the risk of load shedding after normal incidents. On the other hand, old solar inverters or devices with badly implemented frequency-measuring algorithms will likely not abide by these new requirements. □

5.5 Decentralised Control Design

Insofar, PV systems have been described as a power source that may disconnect as a consequence of frequency oscillations. Especially within scenarios of under-frequency deviations, the uncontrolled loss of power generation must be prevented. If the infeed loss cannot be avoided, it must be counterbalanced with control operations that act in the same time frame of the power loss.

A rather natural step, after the presentation of the modelling framework in Chapter 4 and after the initial study about the impact of the heterogeneity of a population of PV systems at the beginning of this Chapter, is the introduction of a control design for PV systems. In view of the distributed nature of the solar population, we imagine to be able to control the power output of each individual

solar device, in a decentralised manner.

We briefly recall here the topic of power output of PV systems, which has been discussed in Section 3.2.1. The power output of PVs is tightly linked to the Maximum Power Point Tracking (MPPT) algorithm. At given conditions of temperature and solar irradiance, the MPPT computes the best (in terms of power) point in the voltage-current plane. We argue that with small effort a different kind of procedure can be implemented: once the MPPT finds the optimal point, it can backtrack towards another point, where the power output represents a pre-defined ratio (e.g. 90%) of the maximum. A schematic representation of the proportional control is shown in Fig. 5.9.

This small and easily implementable modification allows to vary the solar power output to support the grid regulation when needed. This provides fast compensation for losses, yet it cannot be used extensively in time due to the lack of storage. Several works address the renewables as a source of ancillary service, see e.g. [153, 200–202], especially within the smart grids framework. Beyond technical feasibility, the economic aspects must be carefully considered. Fundamental to the use of PV systems in frequency regulation is the presence of storage: whether in remote areas [202], or by using intelligent [153] and industrial [201] loads, or by studying the behaviour of a whole micro-grid [200]. As such, the approach proposed in the following section would benefit from the implementation of storage solutions, along with a long-term study on the economic costs.

5.5.1 Design of a Proportional Control

We have so far considered the power output of connected PV systems to be constant in time. However, in practice the operational mode of a PV system is crucially determined by the so called Maximum Power Point Tracking (MPPT). This mechanism is embedded in the PV system's inverter and employed to maximise its power output, and therefore its efficiency, which depends on several factors, such as solar radiation and external temperature [170, 172]. Within our modelling framework, the MPPT algorithm acts solely when a PV system is active, namely in

state ON, whereas it remains idle when the PV system disconnects. Based on this practical setup, we raise the assumption to be able to control the power output of each device level by means of a proportional gain controller, which is introduced in the following Eq. (5.12). We argue that in practice this control scheme can be implemented within the inverter alongside the MPPT method with little effort. The overall algorithm works as follows: (i.) it first finds the MPP; (ii.) it then computes the desired power output via a proportional control (described below); (iii.) it finally selects a new working point based on the ratio between the desired power and the maximum power for the PV system.

Building on the described control algorithm for the single PV system, we propose a decentralised scheme to control the entire population of PV systems, which can be interleaved with the current primary network control. We assume that each device, during normal operation, injects a power level P_{eq} into the grid that is strictly less than P_{MAX} – the maximum available power to the PV system. The quantity P_{eq} can be tuned according to specifications and requirements by the Transmission System Operator [185].

We utilise the following proportional control law:

$$P_{PV}(k) = P_{eq} + k_p \cdot \Delta f(k), \tag{5.12}$$

where k_p is a constant gain to be determined, and $\Delta f = f(k) - f_0$ represents the frequency deviation from the nominal value. The tuning of k_p is crucial: a small value generates a slow control action, whereas a high value determines an oscillatory response.

Disturbance Rejection Method The gain k_p is calculated via a method called *disturbance rejection at the steady state*. This method tunes the proportional gain k_p according to the maximum desired output deviation after a predefined step disturbance. Practically speaking, we are aware of the contingency (the maximum normal contingency, i.e. 3 GW power loss), and we define the maximum frequency deviation that the power grid can handle (for example, 0.5 Hz) after such an

incident. We thus set a value of k_p that ensures that the controlled system fulfils the specifications. We define f_{ss}, representing the value of frequency at steady state, in the case of a step disturbance $d(t) = A > 0,\ t \geq 0$, so that

$$f_{ss} = A \cdot \frac{G(0)}{1 + k_p G(0)}, \tag{5.13}$$

where $G(0)$ is the steady-state gain of $G(s)$. Considering a maximum acceptable steady-state disturbance value f_{ss}^{max}, we can characterise a working region for the gain as

$$k_p \geq G(0)^{-1} \left(\frac{A \cdot G(0)}{f_{ss}^{max}} - 1 \right). \tag{5.14}$$

This results in a minimum value for k_p, which is then compared to the root locus analysis in Section 5.2 (cf. Fig. 5.3). The root locus shows the coordinates of the poles of the interconnected system on the complex plane. The imaginary component results in oscillations of the response, which are generally undesirable: with the goal to avoid oscillations of the frequency response, we can select k_p resulting in real poles for the closed loop.

Deadband Controller The intrinsic noisy nature of the frequency signal, due to the imperfect matching between generation and load, suggests the introduction of a (small) deadband for the controller design. Accordingly, regulation aims at maintaining the network frequency within a (small) interval around f_0. The introduction of the deadband implies that a control action is performed solely when the frequency is outside this predefined interval, whereas if the frequency is close enough to the nominal value f_0 no control action is needed. In the case of underfrequency, we denote the interval over which the control acts as $[f_u^{min}, f_u^{max}]$; similarly in case of overfrequency, $[f_o^{min}, f_o^{max}]$. Let us remark that these intervals are not related to the working thresholds of the devices.

Finally, the power output of the solar population is limited between P_{MAX} and 0, as the PV systems cannot go above the their maximum power or have negative output. We show the saturated control law, with the addition of the deadband around f_0, in Fig. 5.9.

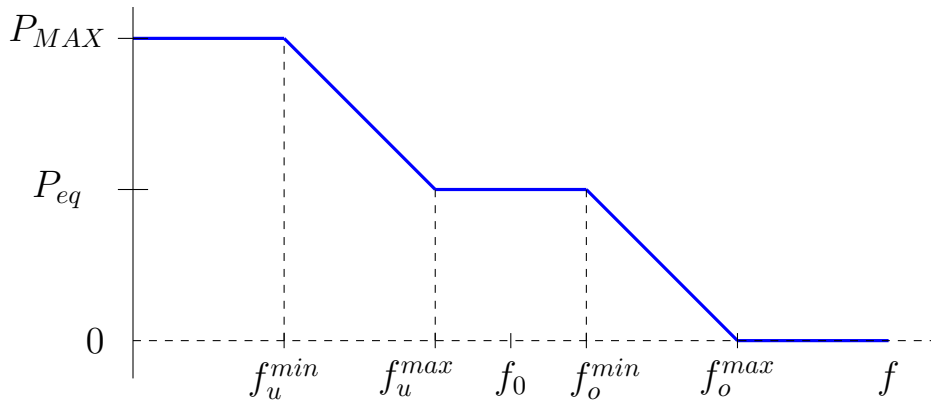

Figure 5.9: Power output in relation to the network frequency.

5.6 Experiments with a Controlled Power Output

We now implement the control architecture discussed in the previous section. In particular, we test the use of the proportional gain in those scenarios from above where the generation loss would result in the activation of the load shedding procedure. We aim at highlighting that, with the implementation of a simple proportional control architecture, household devices can enhance the resilience of the grid and mitigate load shedding scenarios.

In the following simulations, controllers utilise a frequency deadband of $[49.95, 50.05]$ Hz. We assume that every PV system implements the same proportional control design. The value of the control gain is set according to the requirements of the disturbance rejection of Eq. (5.14) and accounting for the root locus plot of Section 5.2: we thus set it to $k_p = 4$. P_{eq} is set to 90% of the maximum available power output, following the dead band approach illustrated in Fig. 5.9. As an example, in a 220 GW network with 10% solar penetration, the solar power output accounts for a total of 22 GW. P_{eq} is thus 19.8 GW, i.e. the 90% of 22 GW, whereas the available control power for frequency regulation amounts to 2.2 GW, i.e. the remaining 10%.

We revisit the scenarios outlined in Section 5.4 and compare against the new controlled settings. Figure 5.10 depicts the frequency response after the infeed loss without any controlled device (solid lines), and with controlled PVs (dashed lines). PV systems contribute to 10% of the 220 GW network. It can be seen that the

controlled signals are able to sustain the network after the incident, and to prevent PV systems from disconnecting, which is the cause of load shedding. Similarly, Fig. 5.11 considers a 440 GW network load scenario with the solar penetration set to 20%.

Finally, on a 220 GW network load scenario Fig. 5.12 shows the frequency signals with uncontrolled and controlled power output, obtained with the highest and lowest values of variance tested (the uncontrolled signals correspond to those in Fig. 5.8). In both cases, the controlled output keeps the frequency very close to the nominal value (within tens of mHz). As evidenced in Fig. 5.8, no load shedding procedure is activated even in the absence of the control design. Nevertheless, a control architecture enhances the electricity grid, providing an immediate ancillary service and, in this scenario, avoids the disconnection of a portion of PV systems – the portion of the population with disconnection thresholds at 49.8 Hz and 49.7 Hz. Further, a frequency decrease of hundreds of mHz represents a stressful situation for the grid: with an active control the frequency decrease remains within tens of mHz, thus significantly improving the aftermath of an incident.

As our simulations clearly show, the implementation of a simple proportional control architecture is sufficient to avoid the activation of the load shedding procedure in all scenarios. The frequency signals nevertheless carry a steady-state error due to the proportional nature of the control. As discussed in section 2.1.3, the steady-state error will be addressed by the secondary control, which acts after few seconds after a significant frequency deviation, hence beyond the analysed time horizon.

Our studies underline the relevance and the vulnerability associated to solar power generation. Photovoltaic systems are a useful resource of clean energy and of control power. Our initial modelling framework considers PV systems only as ON/OFF devices (from the perspective of their power production) that might turn off under stress conditions in the electricity grid. If they remain unable to actively participate to the grid regulation, they can have a destabilising impact on the overall network stability.

On the other hand, it is easy to understand why PV systems cannot be the only source of regulation: our control approach is intended to show the potential

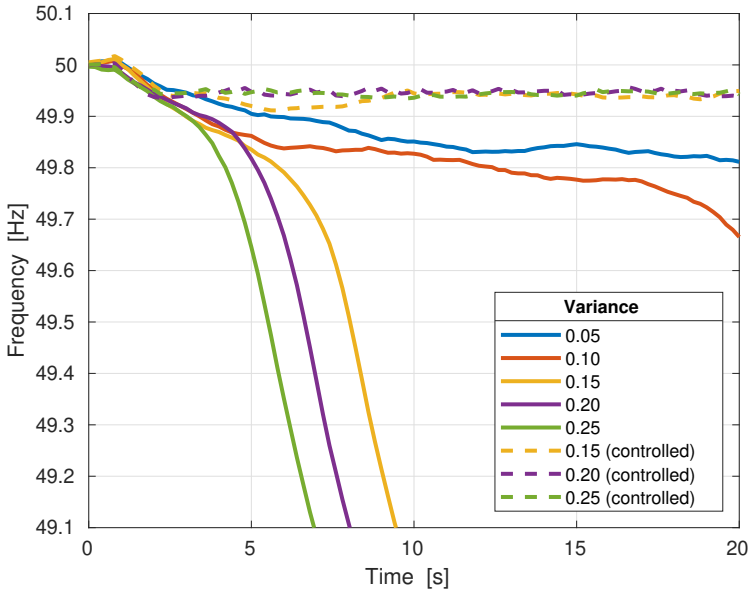

Figure 5.10: Comparison between frequency response after a 3 GW loss of production incident, with 10% of renewable power on a 220 GW network load.

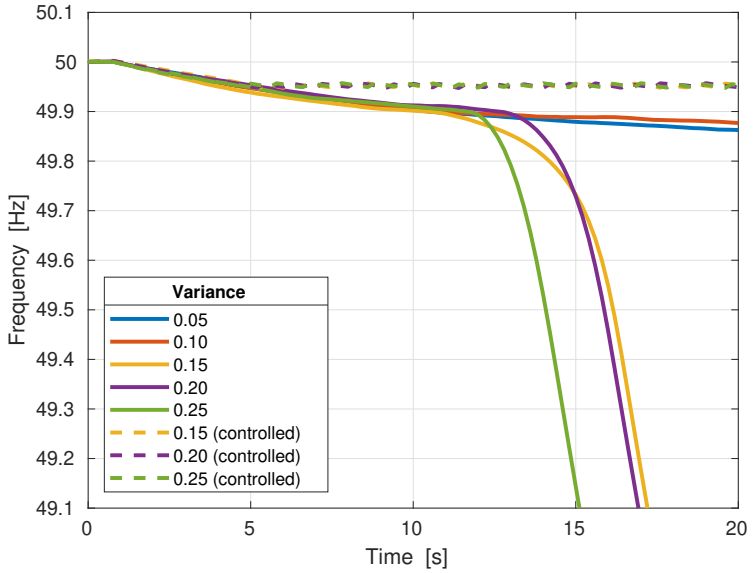

Figure 5.11: Comparison between frequency response after a 3 GW loss of production incident, with 20% of renewable power on a 440 GW network load.

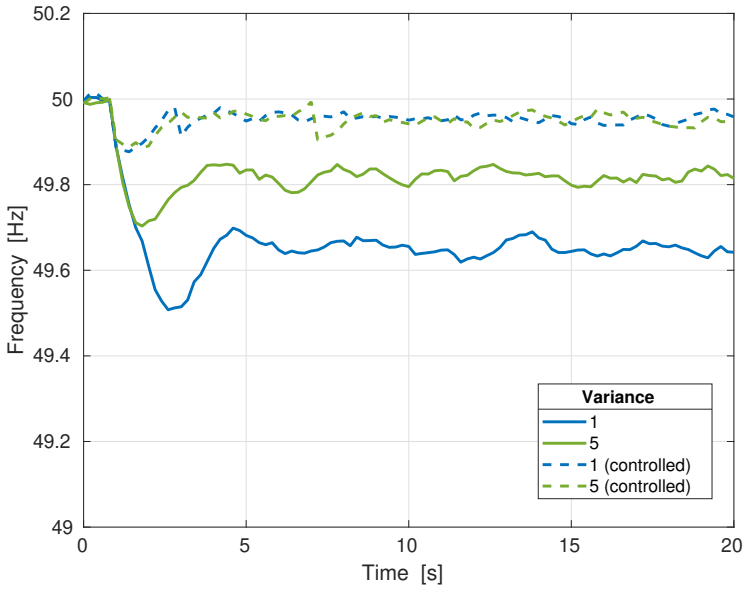

Figure 5.12: Power generation loss of 3 GW on a 220 GW network load with χ^2 distribution for the thresholds. Solar penetration is around 10%.

of such a resource, however this has clear practical shortcomings. First, from the economic point of view, the constant 10% loss of production – by setting P_{eq} to 90% of the available power – is a non-negligible amount over an extended period of time. Further, renewables need a storage system to provide a long-term reliable regulation: weather unpredictability and occlusions are among the factors that need to be accounted for. Simply put, P_{MAX} takes different values at around midday and at around 6 pm. A reliable power grid control architecture should account also for this shortcoming. As such, the discussed control approach is useful in extreme situations, e.g. after an incident, rather than in daily regulation scenarios.

5.7 Concluding Remarks

Chapters 4 and 5 have presented a modelling framework for the aggregation of a heterogeneous population of PV systems connected to the electric grid via Markov models. The frequency dynamics are simulated with a discrete-time equation encompassing the primary control with modified inertia. Specifically, a linear dependency between the amount of conventional power and the time to launch of

the network is used. As a result, we show that the introduction of solar power increases the oscillatory response of the grid after an incident. Wider overshoots are a dangerous inconvenience, as they can stress the physical network and lead to a load shedding procedure.

Unfortunately, the model utilised in this chapter lacks a detailed explanation of all its components or sub-loops. In particular, we highlight that the the sub-loop defining P_m, on the right side of Fig. 5.1, lacks a tunable gain and we could not justify its absence. Notice instead that all the other sub-loops feature a tunable gain (e.g. the load self-regulation k_a).

Regarding other network models, the ENTSO-E has recognised the importance of shared power grid models across the various TSOs: it has recently ruled a standard for the exchange of dynamical models, named Common Grid Model Exchange Standard (CGMES) [203]. This standard ought to be used for the exchange of power system models between TSOs for the purpose of performing bilateral, regional or pan-European studies within TSOs projects. Notably, the standard considers the exchange of network models encoded in the XML format.

Further, we have tested the stability of the network in several scenarios of generation loss with different distributions of population thresholds. Experiments show a correlation between load shedding and the variance of the chosen distributions, or more precisely, to the number of PV systems that disconnect at values of frequency close to the nominal value. A Gaussian distribution has a symmetric shape around its average value: when its variance increases, the tails on both sides spread out. If frequency thresholds are distributed according to a Gaussian, increasing the variance of the thresholds distribution causes a higher number of PV systems (represented by the tails of the distribution) to have thresholds closer to f_0. Consequently, we observe more PV systems with a narrow working interval around the nominal frequency. As a result, the network is more likely to fail, because it becomes more susceptible to small frequency deviation: as highlighted by Tables 5.2 and 5.3, this issue is more relevant when a larger population is connected to the grid. In the case of the χ^2 distribution,

increasing their variance leads to larger thresholds. As expected, the experiments show that an increased heterogeneity guarantees a more reliable network.

Finally, we have introduced a proportional control scheme to control the aggregate power coming from PV systems. Photovoltaic systems individually regulate their power output to restore nominal frequency conditions after an incident. The control framework supports a decentralised implementation. The control design is simple, made to resemble the already existing primary frequency regulation. However, traditional primary network regulation does not act as fast as the PV systems do, being related to the inertia coefficient. Despite its sheer simplicity, the control implementation presented in this work clearly increases the resilience of the network in case of sudden losses for brief time instants. Simulations have proved that these devices are a useful resource for frequency regulation in extreme situations. We have verified that such a simple control could sustain the network and avoid load shedding. This approach can be implemented by flexible loads, such as in the case of cooler devices, as outlined in [16, 55, 56]. Extensions of this work might include modelling storage systems, such as batteries or electric vehicles.

Concluding, we have used simulation-based approaches to highlight and detect the presence of potential issues as a chain of disconnections leading to a load shedding event. However, these case studies are just *one* representation of the power grid response to a contingency: the intrinsic stochastic nature of the system at play may undermine the validity of our approach. One way of dealing with stochastic systems is the use of Montecarlo-based approaches to validate and to study the variability of our results. Notably these approaches are computationally heavy and may not be sufficiently accurate to find insidious corner cases. To overcome these impediments, in the next Chapters we will use formal methods to provide guarantees of the correctness of our results and to offer quantitative answers to the possibility of a load shedding after a contingency.

6

Certification of the Electric Network's Safety

Contents

In the previous chapters we have presented the modelling framework (see Chapter 4), simulation-based scenarios on the impact of a large population of PV systems after an incident and a control approach in line with the grid regulations (see Chapter 5). Simulations help visualise the grid behaviour in the instants following an incident, and may be used to prove the presence (or the risk) of load shedding. However, especially in a stochastic setting, a simulation is just *one* instance of much more complex dynamics. As an example, every result in Chapter 5 assumes an initial frequency value $f(0) = 50$ Hz. One may ask what would be the result if we run tests using $f(0) = 49.95$ Hz, or what if the noise takes different values over the simulation

time. We could set up a series of tests varying the initial frequency value, running (e.g.) 100 simulations per test, and finally averaging and comparing the variance of the results. Following this approach, the number of runs needed to fully characterise a scenario tends to increase with the precision that a researcher wants to achieve.

Formal methods, on the other hand, represent a viable, compact and comprehensive approach for the study of a stochastic scenario. The methodology of formal abstractions in particular incorporates the stochasticity into the verification framework, thus overcoming the ambiguity of a random simulation. In this chapter we present a formal abstraction framework to identify network conditions that lead to load shedding, while varying several network parameters.

The approach translates continuous dynamics – the grid frequency and solar power signals of our setting – into a Markov model via a state-space discretisation, whereas transitions among states are defined by the marginalisation of the additive noises ω_f and ω_P. The procedure comes with a tunable error bound, which guarantees the overall correctness of the results.

As outlined in the previous chapters, we assume that a load shedding procedure is activated as soon as the grid frequency reaches the value of 49.2 Hz. As such, every value of $f(k) \leq 49.2$ Hz represents a critical circumstance which is defined as an unsafe state – cf. Section 2.1.4. Notably, formal methods allow us to compute the *probability* of reaching the unsafe state in order to offer a provably correct, formal certificate of the grid reliability. To provide a complete investigation of the risk of load shedding, we vary network parameters as the renewables penetration, total load, and PV system's working intervals.

6.1 Formal Abstractions of the Grid Dynamics

We now apply the formal abstraction technique discussed in Section 2.2.5 to the closed-loop model introduced in Eq. (5.11). Similarly to the simulation-based approach of Chapter 5, we introduce stochastic kernels t_f and t_P at the level of

frequency and solar power evolution. In practical terms, we add $\omega_f(k)$ and $\omega_P(k)$, i.e. the realisations of t_f and t_P at time k, to the equations of the dynamics, as

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \alpha_2 \Delta f(k-1) + \beta_1 \Delta P_{PV}(k) + \beta_2 \Delta P_{PV}(k-1) + \omega_f(k), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)), \\ P_{PV}(k) = \bar{P}Nx(k) + \omega_P(k), \end{cases} \tag{6.1}$$

where $\Delta f(k) = f(k) - f_0$ and $\Delta P_{PV}(k) = P_{PV}(k) - P_0$. The dynamics of variables $x(k)$ and $y(k)$ represent the portion of PV systems in the population that are in state ON and in state WAIT at time k, respectively (cf. Fig. 4.7). Both $x(k)$ and $y(k)$, as well as $P_{PV}(k)$, are continuous variables: this feature makes their formal verification undecidable [65].

A First-order Modelling Translation The first step towards the application of the formal abstractions technique is the transition from a second-order to a first-order model. The dynamics of the frequency follow a second-order (stochastic) difference equation – cf. the definition of $G(z)$ in Eq. (5.7) – witnessed by the presence of terms $f(k-1)$ and $P_{PV}(k-1)$. We may translate the second-order into first-order dynamics by variable renaming. Formally, we introduce two new state variables

$$\phi(k) = f(k-1), \qquad \xi(k) = P_{PV}(k-1),$$

so that

$$\Delta\phi(k) = f(k-1) - f_0 = \Delta f(k-1),$$

$$\Delta\xi(k) = P_{PV}(k-1) - P_0 = \Delta P_{PV}(k-1).$$

Recall that f_0 and P_0 denote fixed quantities at the equilibrium points: the reference network frequency and the reference solar power production, respectively. We can

now outline an overall first-order model: the system of equations in (6.1) becomes

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \alpha_2 \Delta\phi(k) + \beta_1 \Delta P_{PV}(k) + \beta_2 \Delta\xi(k) + \omega_f(k), \\ \Delta\phi(k+1) = \Delta f(k), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)), \\ P_{PV}(k) = \bar{P}Nx(k) + \omega_P(k), \\ \xi(k+1) = P_{PV}(k). \end{cases} \tag{6.2}$$

Variables' Range We hereby present a model composed of six variables: $x(k)$, $y(k)$, $\Delta f(k)$, $\Delta\phi(k)$, $P_{PV}(k)$, $\xi(k)$. These quantities belong to different domains, hence dwell within different intervals. By construction, $x(k)$ and $y(k)$ belong to the interval $[0, 1]$: indeed they represent the portion of active and waiting PV systems, respectively. On the other hand, $\Delta f(k)$, $\Delta\phi(k)$, $P_{PV}(k)$ and $\xi(k)$ range over $\mathbb{R}$.

However, the system in (6.2) is valid under *normal grid operations*, i.e. when no centralised frequency-control actions are active. As mentioned earlier, whenever $f(k)$ exits its normal operational range, and in particular when $f(k) \leq 49.2$ Hz (for instance, because of a generation loss), primary control mechanisms act to restore the frequency to its nominal value. When the LFC action is active, the current model (6.2) is thus no longer valid. As such, we can limit our framework to values of Δf within the operational range $\mathbb{F} = [f_u, f_o] = [-0.8, +0.8]$ Hz, which corresponds to the frequency interval $[49.2, 50.8]$ Hz. Similarly, the power output of the solar aggregation can never be negative nor exceed the sum of the individual contributions: we restrict the domain of $P_{PV}(k)$ to the interval $\mathbb{P} = [0, \bar{P}N]$ to model the physical limitation of real devices. Finally, we denote the interval $\mathbb{X} = [0, 1]$ for variables x, y.

The state space of the model is thus characterised by the vector variable q and its domain $\mathcal{Q}$

$$q = (\Delta f, \Delta\phi, x, y, P_{PV}, \xi) \in \mathbb{F}^2 \times \mathbb{X}^2 \times \mathbb{P}^2 := \mathcal{Q}, \tag{6.3}$$

with six continuous components. Let us also introduce a noise vector

$$\omega(k) = (\omega_f(k),\ \omega_P(k)), \tag{6.4}$$

where we assume that both ω_f and ω_P belong to a Gaussian kernel; formally, $\omega_f, \omega_P \in \mathcal{N}$.

Stochastic Transitions The noise term $\omega(k)$ plays the key role of defining the stochastic transitions between one (abstract) state and another. As an example, consider the equation $P_{PV}(k+1) = \bar{P}Nx(k)+\omega_P(k)$ from the system of equations in (6.2). We define the stochastic variable ω_P as belonging to a Gaussian distribution with zero average and given variance σ_P^2; formally, $\omega_P \in \mathcal{N}(0, \sigma_P^2)$. On the other hand, the term $P_{PV}(k+1)$ is composed of ω_P and of the deterministic element $\bar{P}Nx(k)$: hence, we may think of $P_{PV}(k+1)$ as a random variable with average value $\bar{P}Nx(k)$ as $P_{PV}(k+1) \in \mathcal{N}(\bar{P}Nx(k), \sigma_P^2)$.

From this perspective, the value of P_{PV} at the next time step is a random variable: as such, we may compute the *probability* of P_{PV} being equal to, or less than, a value of interest. Similarly, we can study the probability distribution of P_{PV} *given that* $x(k)$ is equal to a given value $\underline{x}$. We are, as a matter of fact, interested in this particular question: given that $x(k) = \underline{x}$, what is the probability of $P_{PV}(k+1)$ being equal to, or less than, a given value?

More generally, given the whole six-dimensional state q of Eq. (6.3), we aim at examining the probability distribution of the next state q'. The behaviour of state q depends on the noise ω, composed of two stochastic variables, ω_f and ω_P, whose transition densities are denoted t_f and t_P respectively. Thus, transitions between the current state q and the next state q' are determined [113] via

$$t_\omega(q' \,|\, q) = (\ t_f(q' \,|\, q),\ t_P(q' \,|\, q)\ ), \tag{6.5}$$

where t_ω indicates the one-step transition density function of $\omega(k)$ conditional on point q, and where $t_f(\cdot \,|\, q)$, $t_P(\cdot \,|\, q)$ are the transition density of ω_f and ω_P conditional on point q, respectively. We consider the two densities t_f and t_P decoupled, in view of the assumed independence of the two noise processes.

The stochastic kernel in Eq. (6.5) governs the transitions of the system in Eq. (6.2), which is composed of two stochastic and four deterministic equations.

Also the deterministic elements play a role in the transitions from q to q'. As an example, recall the equation

$$\Delta\phi(k+1) = \Delta f(k).$$

In this case, the next state q' is represented by the value $\Delta\phi$ at time $(k+1)$; we may denote it as $\Delta\phi'$. The transition towards state q' can happen solely if $\Delta\phi'$ is equal to Δf. We rewrite this condition using a Dirac delta [1] $\delta(\Delta\phi' - \Delta f)$, which is centred at $(\Delta\phi' - \Delta f)$ to represent a transition that happens if and only if $\Delta\phi'$ is equal to Δf.

The transition density t_ω can be thus written as a product of four Dirac delta functions, one per deterministic equation in the system (6.2), and the two densities t_f and t_P. Conditional on point $q \in \mathcal{Q}$ it results in

$$\begin{aligned} t_\omega(q' \,|\, q) = t_f(\Delta f' - \alpha_1 \Delta f - \alpha_2 \Delta\phi - \beta_1 \Delta P_{PV} - \beta_2 \Delta\xi) \\ \cdot\ \delta(\Delta\phi' - \Delta f) \cdot\ \delta(x' - (1-a)x - b\varepsilon y) \\ \cdot \delta(y' - b(1 - x - \varepsilon y)) \cdot t_P(P'_{PV} - \bar{P}Nx) \cdot \delta(\xi' - P_{PV}), \end{aligned} \quad (6.6)$$

where primed variables indicate the next value in time.

6.1.1 Finite Abstraction via State-Space Partitioning

Insofar, we have introduced the stochastic model of Eq. (6.2) and have characterised the transitions between continuous states via the conditional kernel of Eq. (6.6). In order to translate the continuous states q into a set of finite, discrete states s we employ a state-space partitioning technique. Roughly speaking, we divide the six-dimensional space $\mathcal{Q}$ into a finite number of hyper-rectangles s, and compute the transition probability from (every) one hyper-rectangle to another.

The formal abstraction technique is based on a state-space partitioning procedure: consider an arbitrary and finite partition of the continuous domains

$$\mathbb{F} = \bigcup_{i=1}^{n} \mathcal{F}_i, \qquad \mathbb{X} = \bigcup_{i=1}^{m} \mathcal{X}_i, \qquad \mathbb{P} = \bigcup_{i=1}^{m} \mathcal{P}_i, \quad (6.7)$$

[1] The Dirac delta $\delta(x)$ is a generalised distribution, which assumes value 1 where $x = 0$ and assumes value 0 otherwise.

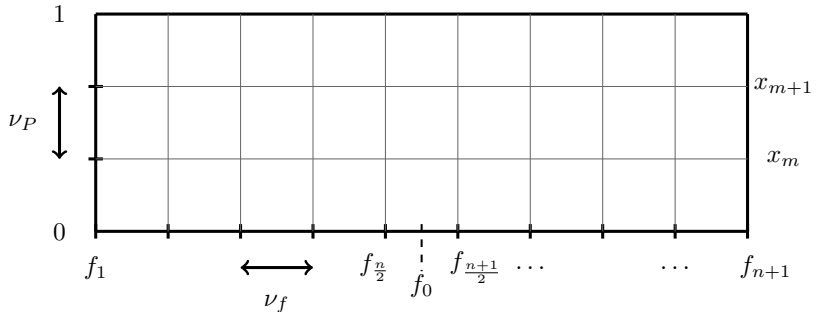

Figure 6.1: Partition intervals for the joint frequency-power state space.

where $\mathcal{F}_i$, $\mathcal{X}_i$, $\mathcal{P}_i$ are non-overlapping intervals. Partition intervals $\mathcal{F}_i$ represent the ranges of values for Δf and $\Delta\phi$, intervals $\mathcal{X}_i$ represent the range of values for x and y, whereas intervals $\mathcal{P}_i$ represent the range of values for ΔP_{PV} and $\Delta\xi$. Figure 6.1 depicts a partitioning procedure in a two-dimensional space which considers solely the network frequency f and the PV systems' active population x.

Within each partition, we choose a representative point that serve as a reference for the whole interval. Let us consider a set of representative points within the partitions

$$\{\bar{f}_i \in \mathcal{F}_i,\ i = 1, \dots n\}, \tag{6.8}$$

which in practice are taken to be their middle points. Similarly, we define representative points

$$\{\bar{x}_j \in \mathcal{X}_j,\ j = 1, \dots m\}, \qquad \{\bar{p}_j \in \mathcal{P}_j,\ j = 1, \dots m\}, \tag{6.9}$$

for variables $\Delta\phi$, x, y, ΔP_{PV} and $\Delta\xi$, respectively. In principle, we could employ different partitioning intervals for different variables, however to ease the notation we use n intervals for frequency-related variables Δf, $\Delta\phi$, and m intervals for $x,\ y,\ P_{PV},\ \xi$.

The partitions $\mathcal{F}_i$ are constructed in the following manner. We first select the partition size ν_f and consequently obtain the number of partitions n by dividing the whole interval $\mathbb{F} = [f_u, f_o]$ as

$$n = \frac{f_o - f_u}{\nu_f}. \tag{6.10}$$

Notably, we have chosen a symmetric interval $\mathbb{F} = [-0.8, +0.8]$ Hz: if n is odd, f_0 becomes the reference point of the $n/2$-th partition, i.e. $\bar{f}_{\frac{n}{2}} = f_0$. Analogously for the power-related variables, we denote ν_P as the partition size and $m = \frac{1}{\nu_P}$ as the number of partitions in the active PV systems domain.

We denote the boundary points of the partitions as

$$f_{i+1} = f_i + \nu_f, \quad i = 1, \dots n, \quad \mathcal{F}_i = [f_i, f_{i+1}), \quad \mathbb{F} = \bigcup_{i=1}^{n} \mathcal{F}_i, \tag{6.11}$$

$$x_{j+1} = x_j + \nu_P, \quad j = 1, \dots m, \quad \mathcal{X}_j = [x_j, x_{j+1}), \quad \mathbb{X} = \bigcup_{j=1}^{m} \mathcal{X}_j, \tag{6.12}$$

and analogously for $\Delta\phi$, y, and P_{PV}, ξ. Let us highlight the interval

$$\mathcal{F}_1 = [f_1, f_2) = [f_u, f_u + \nu_f) = [-0.8, -0.8 + \nu_f), \tag{6.13}$$

that represents the values of frequency just above the load shedding value $f_u = 49.2$ Hz. In order to represent all the values $f < 49.2$ Hz we introduce the unsafe interval $\mathcal{F}_0 = (-\infty, f_1)$. Notably, to compute the load shedding probability, it is thus sufficient to compute the probability of reaching the interval $\mathcal{F}_0$.

The Translation into a Markov Model Introduce now a discrete-time and finite-state Markov chain $\mathscr{M}$, composed by $n^2 \times m^4$ abstract states

$$s = (\bar{f}_{i_1}, \bar{\phi}_{i_2}, \bar{x}_{j_1}, \bar{y}_{j_2}, \bar{p}_{j_3}, \bar{\xi}_{j_4}), \tag{6.14}$$

where $i_1, i_2 \in [1, n]$ and $j_1, j_2, j_3, j_4 \in [1, m]$. We ought to denote an abstract state with six indices, as $s_{i_1,i_2,j_1,j_2,j_3,j_4}$. For simplicity and ease of notation, let us introduce bijective mappings to uniquely and concisely represent the six-indexed states of $\mathscr{M}$ with a single index l. We present two bijections

$$l_n(i_1, i_2) = i_1 + i_2 n, \quad i_1, i_2 \in \mathbb{N}_n \tag{6.15}$$

and

$$l_m(j_1, j_2, j_3, j_4) = j_1 + j_2 m + j_3 m^2 + j_4 m^3, \quad j_1, j_2, j_3, j_4 \in \mathbb{N}_m, \tag{6.16}$$

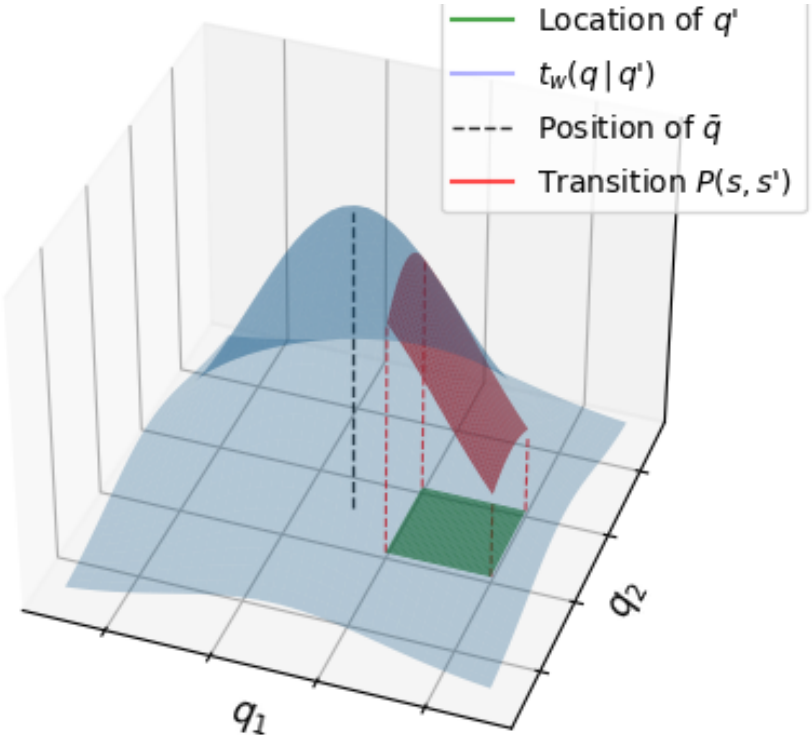

Figure 6.2: Computation of the transition probability $P(s, s')$.

to identify the frequency related and the power related indices, respectively. Notice that l_n and l_m maximum values are $n^2 + n$ and $m^4 + m^3 + m^2 + m$, respectively.

Finally we introduce the mapping

$$l(l_n, l_m) = l_n + l_m(n^2 + n), \tag{6.17}$$

that allows us to unambiguously map the six-indexed state to a single subscript l and indicate a state s_l of the Markov chain $\mathscr{M}$.

Denote by $\mathcal{S}$ the finite state space of $\mathscr{M}$ and by $s_l \in \mathcal{S}$ one of its states. The subscript l is obtained by applying the mappings $l_n(\cdot, \cdot)$ to the i indices, $l_m(\cdot, \cdot, \cdot, \cdot)$ to the j indices, and finally $l(\cdot, \cdot)$ to the resulting values. The abstract state s_l corresponds to a hyper-rectangle centred at $(\bar{f}_{i_1}, \bar{\phi}_{i_2}, \bar{x}_{j_1}, \bar{y}_{j_2}, \bar{p}_{j_3}, \bar{\xi}_{j_4})$ and with bounds defined by the intervals $\mathcal{F}_i$, $\mathcal{P}_j$ and $\mathcal{X}_j$ and corresponding copies, as per Eqs. (6.11), (6.12). Denote $M : \mathcal{S} \to \mathcal{Q}$ the one-to-one mapping between the abstract state s and the corresponding region of the state-space q.

We are now ready to compute the transition probability from the abstract state s to the next abstract state s'. The transition probability matrix of $\mathscr{M}$

comprises the probabilities obtained by marginalising the kernel t_ω over the hyper-rectangular partitions, as

$$P(s, s') = \int_{M(s')} t_\omega(q' \,|\, q) = \int_{M(s')} t_\omega((df', d\phi', dx', dy', dP'_{PV}, d\xi') \,|\, q). \tag{6.18}$$

Roughly speaking, we compute the transition from s to s' by integrating the density *given that* the current state is $q = M(s)$. We consider in practice the transition starting from the reference point of state q, that is, from Eq. (6.14), $\bar{q}$. The integration is performed over the region corresponding to s', i.e. over q', using the map M. Figure 6.2 offers a schematic depiction of the transition probability computation in a two dimensional space, where q_1 and q_2 denote the first and second component of quantity q. The kernel t_w (shown in light blue) is centred at $\bar{q}$, the reference point of q (denoted with a black line). The transition probability (red) is evaluated as the integral over the green region q'.

In view of the presence of deterministic dynamics in the system of equations (6.2), we talk about a degenerate stochastic model. The abstraction of such a model results in a Markov chain with a peculiar structure, as the following example illustrates.

Example 2 *Consider, as an illustrative example, the following model:*

$$\begin{cases} r_1(k+1) = \zeta_1 r_1(k) + \zeta_2 r_2(k) + \omega_z(k), \\ r_2(k+1) = r_1(k), \end{cases} \tag{6.19}$$

where ζ_1, ζ_2 are constants and $\omega_z(k)$ is a Gaussian noise term at time $k \in \mathbb{N}$. These models are typical in control engineering, as they derive from auto-regressive systems, such as

$$r_1(k+1) = \zeta_1 r_1(k) + \zeta_2 r_1(k-1) + \omega_z(k),$$

where the new variable $r_2(k)$ is introduced to replace the delayed variable of interest. Let us assume a 2-set partition of $\mathbb{R} = A \cup B$ with reference points $\bar{r}_1 = r_A, \bar{r}_2 = r_B$. Within the abstract framework, both variables $r_1(\cdot)$ and $r_2(\cdot)$ can either take value r_A or r_B.

The state-space of (6.19) *is $q = (\bar{r}_1, \bar{r}_2) \in \{r_A, r_B\}^2 = \{r_A r_A, r_A r_B, r_B r_A, r_B r_B\}$. The dynamics of r_2 allows only for deterministic transitions, as the next value of r_2*

must be the current value of r_1. *As an example, if the current state is* $q = (r_A, r_A)$, *the next state must be* $q' = (*, r_A)$ *and cannot be* $q' = (*, r_B)$. *Thus, adding an auxiliary variable* r_2 *practically expands the state-space while forbidding several transitions. This results in a particular structure of the transition matrix, as*

$$P = \begin{bmatrix} * & 0 & * & 0 \\ * & 0 & * & 0 \\ 0 & * & 0 & * \\ 0 & * & 0 & * \end{bmatrix},$$

where the elements in $*$ *are obtained by marginalising the density of* ω_z, *offset by the drift term, over the corresponding set* A *or* B, *as per Eq.* (6.18). □

The peculiar structure of the transition matrix allows a computational speed up by using sparse matrices.

The abstraction procedure applied to the model in (6.2) carries a discretisation error: in the following, we formally derive a bound for this error as a function of the discretisation steps ν_f and ν_P. As argued in [113], a finer grid results in a smaller abstraction error, however it generates a larger state space. Further, we will formally outline the calculations that will provide the value of the probability of load shedding.

6.1.2 Quantification of the Load Shedding Probability and of the Abstraction Error

Let us now formally characterise the load shedding probability within a finite time horizon. Consider the model in (6.2) with initial state q_0 and select a discrete time horizon h. We assume that the electric network activates the load shedding procedure whenever $f(k) \leq 49.2$ Hz, namely if $q(k) \in \mathcal{L}$, where $\mathcal{L} := \{\Delta f \leq -0.8\}$ is the continuous-domain equivalent of partition $\mathcal{F}_0$. The probability of load shedding within h time steps can be defined as

$$p_{q_0}(\mathcal{L}) := \mathsf{P}(q(i) \in \mathcal{L}, i \in [1, h] \mid q_0), \tag{6.20}$$

where q_0 is the initial state of the continuous model. On a reverse perspective, the probability of load shedding is the complement to the probability of remaining within a safe set $\mathcal{A}$ for h steps.

Value Functions and Backward Iteration We now illustrate the computation of the probability of reaching a general set $\mathcal{A}$, as outlined originally in [65], to then adapt the procedure to the set $\mathcal{L}$. Recalling Section 2.2.4, we devise a procedure to compute a *safety* probability, i.e. the probability of a system to remain within a safe set $\mathcal{A}$ during the whole time horizon.

To this end, we set up a procedure to evaluate the probability that the system dynamics, associated with the initial condition $q_0 \in \mathcal{A}$, will remain within the set $\mathcal{A}$ during the whole time horizon $[0, H]$. Formally, we compute the probability

$$p_{q_0}(\mathcal{A}) := \mathsf{P}_{q_0}\{q(k) \in \mathcal{A} \text{ for all } k \in [0, h]\}. \tag{6.21}$$

Practically speaking, if $p_{q_0}(\mathcal{A}) \geq \varepsilon$, with $\varepsilon \in (0, 1]$, we state that the system is *safe* with an ε probabilistic guarantee, when initialized at q_0. Notably, $p_{q_0}(\mathcal{A})$ depends on the initial condition, yet its evaluation is all but trivial. The works in [65, 112] present an equivalent rewriting of the problem, which we report here, showing that $p_{s_0}(\mathcal{A})$ is calculated via a backward iterative procedure by representing $p_{q_0}(A)$ as a multiplicative function.

Let $\mathbb{1}_C : \mathcal{Q} \to \{0, 1\}$ denote the indicator function of set $C \subseteq \mathcal{Q}$: formally, $\mathbb{1}_C(q) = 1$ if $q \in C$, and $\mathbb{1}_C(q) = 0$ if $q \notin C$. In view of the discrete-time framework, we may compute the indicator function *at every time step* to gather information about the location of state q. Further, the product

$$\prod_{k=0}^{h} \mathbb{1}_{\mathcal{A}}(q(k)) = \begin{cases} 1 & \text{if } q(k) \in \mathcal{A} \text{ for all } k \in [0, h], \\ 0 & \text{otherwise,} \end{cases} \tag{6.22}$$

where $q(k) \in \mathcal{Q}$, $k \in [0, h]$, offers a binary result: 1 if q has remained within $\mathcal{A}$ for all time steps, and 0 otherwise. Notice that this measurement can only be applied *a posteriori*, i.e. after the state $q(k)$ at time k has "landed" inside or outside $\mathcal{A}$. The system in (6.2) is probabilistic, hence we employ a probabilistic measure to fully characterise the behaviour of such model. We may express the value $p_{q_0}(\mathcal{A})$ in (6.21) as the expectation of the random variable $\prod_{k=0}^{h} \mathbb{1}_{\mathcal{A}}(q(k))$:

$$p_{q_0}(\mathcal{A}) := \mathbb{E}_{q_0}\left[\prod_{k=0}^{h} \mathbb{1}_{\mathcal{A}}(q(k))\right], \tag{6.23}$$

where $\mathbb{E}[\,\cdot\,]$ denotes the expectation of a random variable. The value of $p_{q_0}(\mathcal{A})$ can be computed iteratively, starting from q_0 and multiplying, at each time step, the probability of remaining within $\mathcal{A}$.

We will follow a dynamic programming approach, that evaluates $p_{q_0}(\mathcal{A})$ result through a backward recursive procedure using value functions while guaranteeing computational tractability and ease of implementation. Consider the sequence of functions $V_k : \mathcal{Q} \to [0,1]$, $k = 0, 1, \ldots, h$, defined by:

$$V_k(q) = \mathbb{1}_{\mathcal{A}}(q) \int_{\mathcal{A}} V_{k+1}(q(k+1)) t_\omega(u \,|\, q) du, \quad q \in \mathcal{Q}, \tag{6.24}$$

initialised with $V_h(q) = \mathbb{1}_{\mathcal{A}}(q)$, $q \in \mathcal{Q}$. The method, generally speaking, starts at time h – recall it is a backward operation – where it defines

$$V_h(q) = \mathbb{1}_{\mathcal{A}}(q). \tag{6.25}$$

In other words, V_h is 1, since we aim at computing the probability of remaining within $\mathcal{A}$ for all time steps, including step h. The next (backward) step evaluates V_{h-1} by integrating over $\mathcal{A}$ the product of V_h times the stochastic kernel t_ω. In practice, the integral of t_ω over $\mathcal{A}$ gives the probability of being within $\mathcal{A}$ at the next time step, multiplied by the probability of the previous time steps – in this case, represented only by V_h. As the index k decreases, we accumulate the previous probability value in V_{k+1} until the procedure reaches $k = 1$.

Computation of the Load Shedding Probability The load shedding set $\mathcal{L}$ is formally defined as an unsafe set, i.e. a set that is outside the domain $\mathcal{Q}$. Thus, we compute the probability of reaching the set $\mathcal{L}$ over the time horizon h as its complement to one. Analytically,

$$p_{q_0}(\mathcal{L}) = 1 - p_{q_0}(\mathcal{Q}). \tag{6.26}$$

Computation of the Abstraction Error Let us recall the abstraction error computation outlined in Section 2.2.5. The formal abstraction technique encompasses a safe set $\mathcal{A}$, whose size is $\mathscr{L}(\mathcal{A})$. We partition $\mathcal{A}$ in hyper-rectangles of diameter δ_A and compute the transition probabilities among them. To this end we apply the marginalisation of the noise kernel, whose Lipschitz constant is l_w. The abstraction error can be upper bounded (cf. Eq. (2.6)) by

$$\epsilon_{abs} \leq (h-1)\,\delta_A\,l_w\,\mathscr{L}(\mathcal{A}),$$

where h is the time horizon of the specification we are checking.

In the dtMC model, functions $a(k)$ and $b(k)$ are approximated and assume a finite number of values (one for each of the $\bar{f}_i$). This introduces the error term a_{max}

$$a_{max} = \max_{\substack{i\in[1,n]\\ f_i\in\mathcal{F}_i}} \left\| \int_{f_i}^{\bar{f}_i} p^d(u)du \right\|, \tag{6.27}$$

where p^d represents the probability distribution of the frequency thresholds for disconnection. This quantity defines the maximum approximation error introduced with the discretisation in the computation of $a(k)$.

Note that the presence of $\delta(\cdot)$ functions in Eq. (6.6) introduces discontinuities within the domain of the kernel: continuity regions of the kernel density are parts of the state space where the $\delta(\cdot)$ functions are equal to one. Within such regions (which are formally defined in Appendix B.1) the value functions are continuous, and the following holds over pairs of points $q, \tilde{q}$:

$$|V_k(q) - V_k(\tilde{q})| \leq \frac{\alpha_1}{\sigma_f\sqrt{2\pi}}|\Delta f - \Delta\widetilde{f}| + \frac{a_{max}}{\sigma_P\sqrt{2\pi}}|\Delta P_{PV} - \Delta\widetilde{P_{PV}}|, \tag{6.28}$$

where α_1 is the coefficient of the transfer function $G(z)$ introduced in Eq. (5.8). The full algebraic calculation of the error bound is reported in Appendix B.3.

We now abstract the aggregated population of PV systems as a Markov chain based on the procedure described in Section 2.2.5. Computing the solution of (6.24) over the Markov chain, the overall approximation error can be upper-bounded [65] as follows

$$|p_{q_0}(\mathcal{L}) - p_{s_0}(\mathcal{L}_s)| \leq (h-1)\left[\frac{\alpha_1}{\sigma_f\sqrt{2\pi}}\nu_f + \frac{a_{max}}{\sigma_P\sqrt{2\pi}}\nu_P\right]. \tag{6.29}$$

We highlight, with respect to Eq. (6.28), the substitution of $|\Delta f - \Delta\tilde{f}|$ with ν_f, i.e. the partition size of the frequency variable defined in (6.10); similaraly for ν_P. Further, the single-step error of (6.28) is multiplied by $(h-1)$, i.e. the length of the time horizon. Notice that we multiply by $(h-1)$ instead of by h since we assume the exact knowledge of the location of the initial condition q_0 and of the abstract s_0. This error allows to refine the outcomes of the model checking procedure (obtained from $p_{s_0}(\mathcal{L}_s)$) over the concrete population model (corresponding to the unknown quantity $p_{q_0}(\mathcal{L})$).

6.2 Safety of the Electrical Network

We are now ready to use the abstract Markov chain to compute the load shedding probability after a sudden generation loss, under several scenarios. In practical terms, we quantitatively compute the safety of the grid in terms of the probability of reaching the load shedding.

In line with the ENTSO-E requirements [27] and similarly to the experiments in Chapter 5, we assume an infeed loss of 3 GW in a global network with a demand of $S = 220$ GW and $S = 440$ GW. The network model and PV systems parameters are selected in accordance to Chapter 5: in particular, we recall the variance of ω_P and ω_f are respectively $\sigma_f^2 = 0.025$ and $\sigma_P^2 = 0.01$.

The probabilistic model checking tests are implemented using the MATLAB software. Due to the large state space, ν_f and ν_P are set to 0.02 and 0.05 respectively. Recall that the frequency-related variables range over an interval of 1.6 Hz, whereas the power-related variables range over a unit interval. This implies 80^2 partitions coming from the two frequency-related variables, and 20^4 partitions coming from the four power-related variables, for a total of approximately 10^8 states. The transition matrix is stored efficiently as a sparse matrix to enhance computations. As per previous Chapters, the grid frequency is sampled at a rate of 0.2 seconds, consistently with the requirements introduced in [120]. The discussion is focused on the consequences of an incident after a few seconds, hence we consider a time

interval of 20 seconds: the discrete time horizon is thus composed of 100 steps. After this time interval, we assume the centralised LFC intervenes to stabilise $f(\cdot)$.

Numerical Value of the Abstraction Error Following the definition of several numerical parameters, we may compute the abstraction error in Eq. (6.29). The value of α_1 results $^{0.09}/_{50} \simeq 2 \cdot 10^{-4}$ – recall that we consider normalised, i.e. divided by 50, values of the network frequency f.

Assuming that p^d belongs to a Gaussian distribution, the maximum error arising from the computation of the integral in Eq. (6.27) arises when f is close to the mean of the distribution. Intuitively, values of p^d around the mean have a higher numerical value. Recalling that we set $\nu_f = 0.01$, and use a reference point $\bar{f}$ located at the middle of any partition, the maximum error between $\bar{f}$ and any point within a partition is half partition length, i.e. 0.005. The numerical value of a_{max} then results

$$a_{max} = \max_{\substack{i \in [1,n] \\ f_i \in \mathcal{F}_i}} \left\| \int_{f_i}^{\bar{f}_i} p^d(u) du \right\| = 5 \cdot 10^{-3}. \tag{6.30}$$

Substituting these values into Eq. (6.29)

$$\epsilon_{abs} \leq 100 \left[\frac{2 \cdot 10^{-4}}{\sqrt{0.025}\sqrt{2\pi}} 0.01 + \frac{5 \cdot 10^{-3}}{\sqrt{0.01}\sqrt{2\pi}} 0.05 \right] \simeq 0.1. \tag{6.31}$$

This error ought to be attached to the certificates on safety probability derived in the following Sections.

6.2.1 Setup of the Generation-Loss Incidents

As anticipated above, Transmission Systems Operators are tasked with ensuring the safe operation of the grid, and are thus interested in formal guarantees on its dynamics, and in reliable forecasting of potentially problematic situations, such as issues related to frequency responses after a generation loss incident.

Our study concerns the so-called normal incidents, specifically a loss of up to 3 GW of power generation (cf. Section 3.4.3). We assume the initial condition to be $f(0) = f_0$, with the population of PV systems in active (ON) mode ($x(0) = 1$). The generation-loss incident is modelled as a negative step injected into the dynamics

of Eq. (6.2). Assuming that an incident of magnitude M occurs at time $k = \bar{k}$, the dynamics of $f(\bar{k}+1)$ become

$$f(\bar{k}+1) = \alpha_1 \Delta f(\bar{k}) + \alpha_2 \Delta\phi(\bar{k}) + \beta_1(\Delta P_{PV}(\bar{k}) - M) + \beta_2 \Delta\xi(\bar{k}) + \omega_f(\bar{k}), \quad (6.32)$$

and then evolve from time $(\bar{k}+2)$ on as

$$f(\bar{k}+2) = \alpha_1 \Delta f(\bar{k}+1) + \alpha_2 \Delta\phi(\bar{k}+1) + \beta_1(\Delta P_{PV}(\bar{k}+1) - M) + \beta_2(\Delta\xi(\bar{k}+1) - M) + \omega_f(\bar{k}+1). \quad (6.33)$$

Notice that the incident M does not disappear at time $(\bar{k}+2)$ but remains attached to the dynamics *from* the instant $(\bar{k}+1)$ onwards, ideally until the contingency is repaired.

Equations (6.32) and (6.33) display two different deterministic drifts. We ought to apply the formal abstraction procedure to both cases, thus translating the two models into two different dtMCs. Let us denote $\mathscr{M}_1$ the dtMC that abstracts the system with deterministic drift in (6.32) and $\mathscr{M}_2$ for the drift in (6.33). The overall computation of the probability of load shedding encompasses both dtMCs: we employ $\mathscr{M}_1$ to evaluate the probability of load shedding for the first time step and successively employ $\mathscr{M}_2$ for the remaining $(h-1)$ time steps. For simplicity, we assume that $\bar{k} = 0$, namely the incident occurs at the beginning of the time horizon.

Notably, the scenario where the incident happens at time $\bar{k} > 0$ further complicates the procedure: additionally, we ought to translate the incident-less dynamics of Eq. (6.2) into a Markov model, say $\mathscr{M}_0$. The overall load shedding probability thus should employ $\mathscr{M}_0$, $\mathscr{M}_1$, and $\mathscr{M}_2$.

6.2.2 Computation of Load Shedding Probability

Our tests encompass several scenarios, in which we vary:

(a.) the choice of the disconnection distributions p^d;

(b.) μ and σ^2, the associated mean and variance of p^d;

(c.) the total solar penetration in the network.

Distribution	Load Demand [GW]	Renewable [% S]	μ [Hz]	σ^2 [Hz^2]
$\mathcal{U}$	220, 440	10, 20, 30, 40	$[49.5, 49.8]$	$[1, 7.5] \cdot 10^{-3}$
$\mathcal{N}$	220, 440	10, 20, 30, 40	$[49.5, 49.8]$	$[1, 20] \cdot 10^{-3}$
χ^2	220, 440	10, 20, 30, 40	$[49.5, 49.8]$	$[1, 8]$

Table 6.1: Parameters range with a Uniform (first row), Gaussian (second row), χ^2 distribution (third row).

We also ought to recall that the solar penetration modifies the total inertia of the electric network, hence modifies the network transfer function, as per Chapter 5.

We ask our procedure to check the property outlined in (6.20), i.e. we ask *what is the load shedding probability?* Once the model checking procedure has answered this question, we report the parameters' values (mean and variance of the distributions, solar penetration) and consider a combination *satisfactory* if the load shedding probability is smaller or equal to 0.01. Conversely, a combination of parameters is unsatisfactory if the load shedding probability is greater than 0.01.

We test the safety of the electric network varying the average values μ and variance values σ^2 of several uniform thresholds distributions, along with the variation of solar penetration that ranges from 10% to 40% of the network load. We specifically refer to the under-frequency – denoted by the subscript u – disconnection – denoted by the superscript d – values μ_u^d and σ_u^d; for clarity, we will simply indicate them as μ and σ. The ranges for these parameters are summarised in Table 6.1.

The mean value μ ideally represents specific regulations on disconnection and reconnection, whereas the variance σ^2 represents small manufacturing deviations from the desired value. Finally, increasing variance of p^d reflects a more heterogeneous population (in terms of $\mathcal{I}_f$ thresholds).

The threshold distributions for $\mathcal{I}_f$ (the PV systems' operational values of frequency) are either represented by a Uniform distribution, a Gaussian distribution or χ^2 distribution: these are notably dissimilar and are here used to model different PV system's dynamics. The use of different distributions denotes different *modelling* choices: similarly to the simulation-based experiments in Chapter 5, a Gaussian distribution models the inverter measurement noise, whereas a χ^2 can be used

to define a minimum performance setting, namely a minimum working interval. Whilst the Gaussian distribution results in a more realistic choice, the χ^2 offers interesting outlooks on how distributions affect the safety property. The Uniform distribution, on the other hand, represents the classical modelling choice of no prior information about the real devices. Further, the uniform distribution will be a crucial choice in the next Chapter 7 about parameter synthesis. In practice, the selection of a distribution must depend on data measurements coming from real devices, when available.

Our study concerns the effects of the shape of the chosen distribution as a measure for the heterogeneity of the population, and especially the variation of its variance. Ultimately, increasing variance of p^d reflects a more heterogeneous population (more diverse thresholds characterising $\mathcal{I}_f$). As discussed shortly, a larger variance can have opposite consequences on stability, depending on which threshold distribution is used.

$\mathcal{I}_f$ thresholds distributed as a Uniform Table 6.2 shows the "boundary points" of a satisfactory combination, namely the largest values of σ^2 for each μ that offers a load shedding probability smaller than 0.01. Entries ✓ and ☒ indicate that all values and no values, respectively, are satisfactory. We first notice that in a low demand setting ($S = 220$ GW), the load shedding probability is essentially binary (values close to zero or to one), whereas in a high demand scenario ($S = 440$ GW) the load shedding probability grows smoothly with μ and σ^2. A high-demand network is thus more resilient to incidents, since it has a greater inertia, and an incident is proportionally less significant: the maximum load shedding probability value, in this scenario, is around 1.5%. The tests with a wider penetration of solar contributions (10%, 30%, 40% of the total) indicate that the load shedding probability increases when a larger population is connected to the grid: this is intuitive, as more renewable energy renders the frequency response more oscillatory, and thus more likely to deviate from f_0.

220 GW	μ [Hz]				440 GW	μ [Hz]			
[% S]	49.5	49.6	49.7	49.8	[% S]	49.5	49.6	49.7	49.8
10	✓	6	☒	☒	10	✓	✓	☒	☒
20	✓	4.5	☒	☒	20	✓	6	☒	☒
30	✓	2.5	☒	☒	30	✓	5.5	☒	☒
40	✓	☒	☒	☒	40	✓	☒	☒	☒

Table 6.2: Satisfactory combinations with a Uniform distribution.

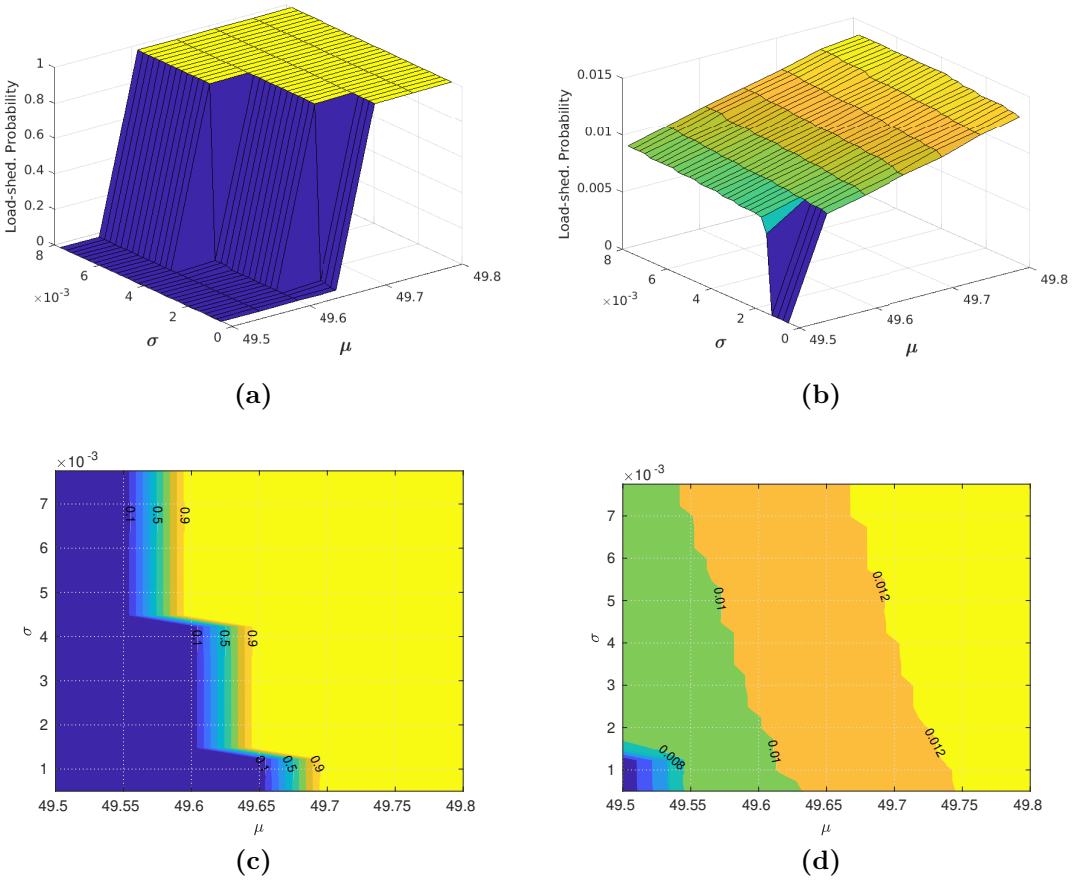

Figure 6.3: Load shedding probability with Uniform thresholds: $S = 220$ GW (left) and $S = 440$ GW (right). The lower figures show the contour plots: notice that the probability reaches approximately 1 with $S = 220$ GW, whereas it reaches 0.015 with $S = 440$ GW.

220 GW	μ [Hz]				440 GW	μ [Hz]			
[% S]	49.5	49.6	49.7	49.8	[% S]	49.5	49.6	49.7	49.8
10	0.05	☒	☒	☒	10	✓	✓	✓	☒
20	☒	☒	☒	☒	20	✓	0.08	☒	☒
30	☒	☒	☒	☒	30	✓	0.06	☒	☒
40	☒	☒	☒	☒	40	✓	0.05	☒	☒

Table 6.3: Satisfactory combinations with a Gaussian distribution.

Figure 6.3 depicts the load shedding probability in the presence of 20% solar penetration, with a network load $S = 220$ GW (left) and $S = 440$ GW (right), where values of μ range within $[49.5, 49.8]$ Hz and those of σ^2 are within $[1, 7.5]$ Hz^2. A higher σ^2 causes a higher number of PV systems to have disconnection thresholds closer to f_0. Consequently, a greater portion of the population is likely to disconnect under frequency deviations, causing the network frequency to decrease. Intuitively, parameter S plays a major role in the verification outcome.

$\mathcal{I}_f$ thresholds distributed as a Gaussian As per Table 6.2, Table 6.3 reports the results in terms of "boundary points" of a satisfactory combination of parameters testing a Gaussian distribution. Similarly to the Uniform distribution, a Gaussian distribution has a symmetric shape around its mean: when its variance increases, the tails on both sides spread out. As such, increasing their variance causes a higher number of PV systems (represented by the tails of the distribution) to have thresholds closer to f_0. When the variance is high, we observe more PV systems with a narrow working interval around the nominal frequency. Therefore, a greater portion of the population is likely to disconnect under frequency deviations, causing the network frequency to decrease. Figure 6.4 depicts the load shedding probability in presence of 20% solar penetration, $S = 220$ GW (left) and $S = 440$ GW (right), varying values of the mean and variance of p^d.

$\mathcal{I}_f$ thresholds distributed as a χ^2 Table 6.4 reports the results in terms of "boundary points" of a satisfactory combination of parameters testing a Gaussian distribution. In contrast to the previous tables, we here report the *minimum*

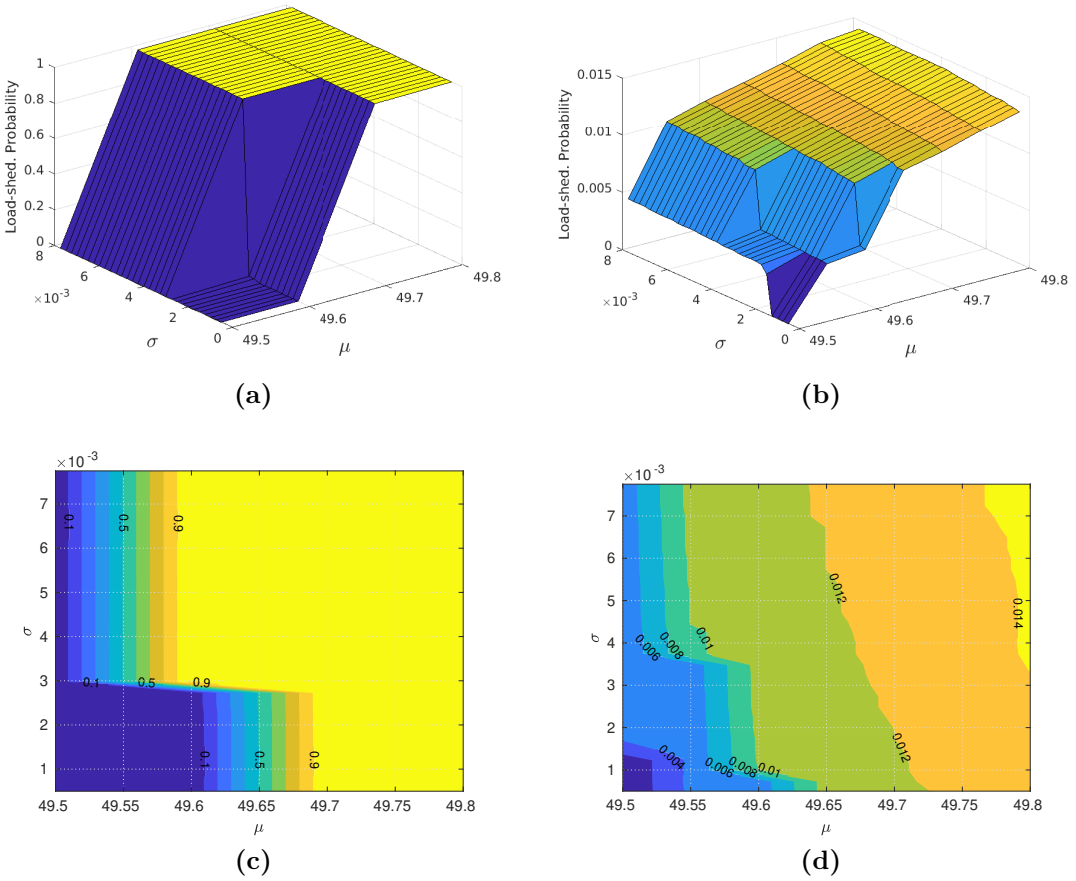

Figure 6.4: Load shedding probability with Gaussian thresholds: $S = 220$ GW (left) and $S = 440$ GW (right). The lower figures show the contour plots: notice that the probability reaches approximately 1 with $S = 220$ GW, whereas it reaches 0.015 with $S = 440$ GW.

220 GW	μ [Hz]				440 GW	μ [Hz]			
[% S]	49.5	49.6	49.7	49.8	[% S]	49.5	49.6	49.7	49.8
10	✓	2	6	7	10	✓	2	5	6
20	✓	4	7	8	20	✓	4	7	8
30	✓	3	5	8	30	✓	2	4	8
40	✓	5	5	8	40	✓	✓	4	8

Table 6.4: Satisfactory combinations with a χ^2 distribution.

value of variance that represents a satisfactory combination of parameters. As the variance of a χ^2 distribution increases, the probability mass – the PV systems' thresholds – move further away from f_0.

Figure 6.5 depicts the load shedding probability under 20% solar penetration, with varying values of the initial point of the support and of the variance of p^d. Note

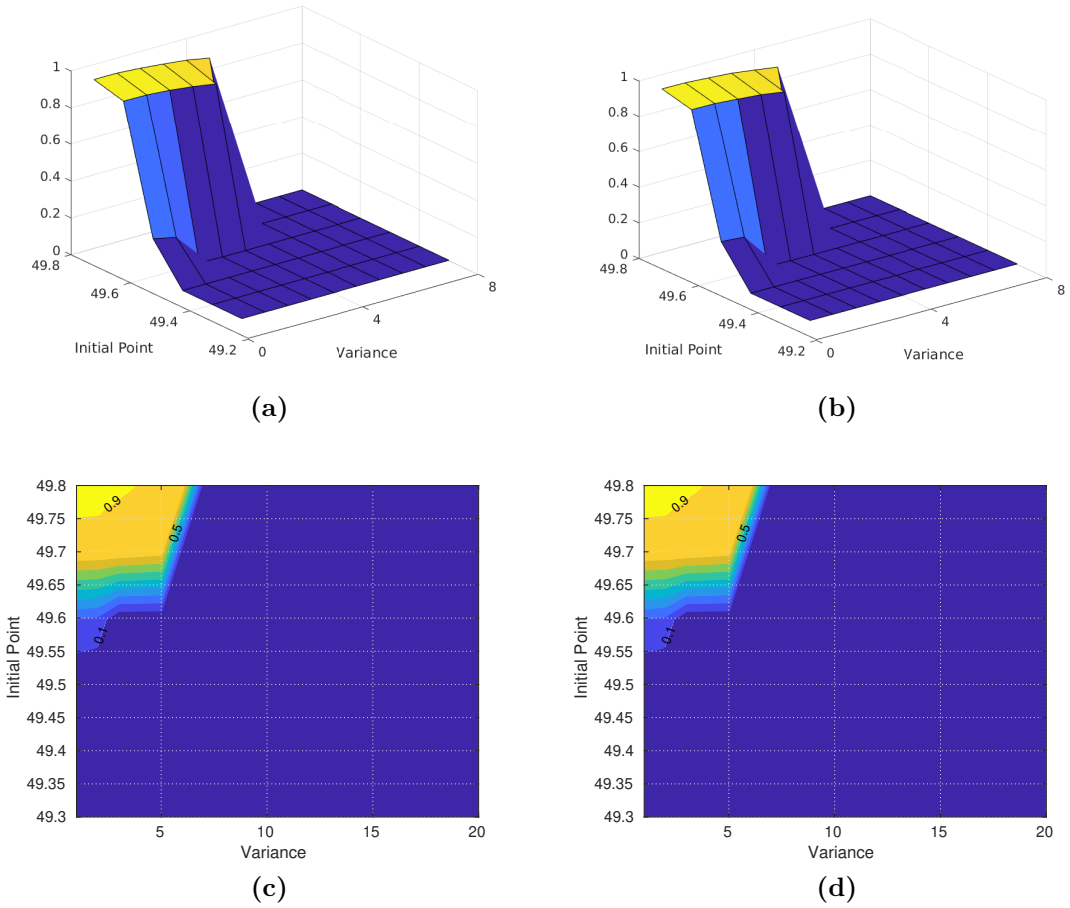

Figure 6.5: Load shedding probability with χ^2 thresholds: $S = 220$ GW (left) and $S = 440$ GW (right). The lower figures show the contour plots.

that, due to the nature of the χ^2 distribution, instead of the average value, we denote an initial point of the support. Unlike the Gaussian case, increasing the variance of a χ^2 distribution results in larger thresholds. As expected, the experiments show that, in this scenario, an increased heterogeneity guarantees a more reliable network.

6.3 Concluding Remarks

In this chapter, we have introduced a formal procedure to abstract the dynamics of a heterogeneous population of PV systems, embedded within the frequency dynamics of the grid. The computation of error bounds on the abstraction guarantees the correctness of the outcomes of a formal verification procedure run on the obtained abstract model. We have focussed our analysis on the probability of load

shedding, encoded as a formal safety property (cf. Section 2.2.4), under significant energy generation from renewables via formal abstractions: we have assessed the load shedding probability of the network, under several scenarios of population heterogeneity. Operators can use these certificates to monitor the distribution of PV systems over the grid and to assess its reliability in case of incidents. In fact, the formal abstraction framework presented in this Chapter may be used as a synthesis and validation tool to define grid connection codes and regulations. We are able to investigate the effect of reconnection and disconnection patterns attaching a provable guarantee of correctness. Further, we may offer guidance on the disconnection and reconnection rules: our results highlight the benefits of a χ^2 over a Gaussian distribution, but we may test other setting or even design one. The next Chapter follows this path by synthesising the parameters that guarantee the networks' safety, rather than just checking them. On the other hand, this approach suffers from the high dimensionality of the network model: the number of states of the Markov chain grows exponentially with the number of partitions and the number of variables under consideration. Hence a trade-off between an accurate underlying model and computational resources ought to be found, according to the desired performance.

7

Synthesis via Parametric Abstractions

Contents

The formal abstraction technique outlined in Section 2.2.5 and applied in Chapter 6 tackles the translation of a stochastic, uncountable state-space model into a probabilistic model as a Markov chain. We have used such a procedure to verify safety requirements on the electric network considering a collection of parameters: the total load of the grid, the kind of threshold distributions and their variance. In practice, we have dealt with a *parametric* (cf. Section 2.2.3) stochastic system: instead of embracing its structure, we have instantiated the parameters, generating a collection of probabilistic models. Every modification of a parameter in the original, continuous model entails a *new* abstraction procedure and a consequent

verification step. This process is clearly time- and resource-consuming.

In this view, it is interesting to analyse how the modifications in the continuous system produce variations in the abstract Markov model. The two models are linked by the formal abstraction procedure, which has at its core the marginalisation of the stochastic kernel over a partitioned state-space. In this Chapter we extend the formal abstraction procedure to parametric models. In practice, we consider a parametric stochastic difference equation, devise a formal abstraction to obtain a parametric Markov chain. Ultimately, we link the SDE parameters with the dtMC parameters: by synthesising the dtMC parameters we are able to evaluate the corresponding parameters on the SDE.

7.1 Formal Abstractions of Parametric SDEs

We offer an approach to formally verify properties on parametric, stochastic, uncountable-state systems by combining formal abstraction techniques and adapting them to a parametric framework. Unlike other approaches in the literature, we consider parameters that can range over unbounded intervals. Given a parametric SDE and a formal property, our procedure entails two steps:

1. a parametric formal abstraction is computed, yielding a (parametric) abstract model;

2. parameters are synthesised on the abstract model in order to satisfy the property given for the SDE.

The abstraction can be iteratively refined to increase its precision. For illustration purposes we analyse three kinds of parametrisation:

(i) additive;

(ii) multiplicative on the deterministic term;

(iii) multiplicative on the noise.

We are interested in the application of formal abstraction techniques [113] to stochastic models, and in particular to stochastic difference equations. Formal abstractions reduce SDEs to finite-state Markov models, performing a finite partitioning of the continuous state space of the SDE and computing probabilities between discrete states by marginalisation of the SDE transition kernel. Such a procedure generates an abstraction error ϵ_{abs}, which in general depends on the continuity of the dynamics of the SDE, and on the diameter of the introduced partitions [65], as outlined in the previous Chapter.

Parameter synthesis problems consider a set of models that share a common structure. Their aim is to find evaluations (instantiations) of the parameters, such that the corresponding model satisfies a given property ϕ (a logical specification). One can be interested in finding single instantiations of the parameters, or rather the maximal region $\mathcal{V}$ (known as *feasible set*) within the parameters space, corresponding to models satisfying a property ϕ. By and large, parameter synthesis problems deal with finite-state models, which in the probabilistic context can be parametric Markov chains, denoted as $\mathscr{M}(\mathbf{v})$, with parameters $\mathbf{v} \in \mathbb{R}^l, l \in \mathbb{N}$, and where $P(\mathbf{v})$ represents a parametric transition matrix for the chain dynamics.

Translation of a Parametric SDE into a parametric dtMC In this Chapter, instead, we consider parametric SDEs, denoted as $\mathscr{F}(\mathbf{t})$, where $\mathbf{t}$ is a parametric vector $\mathbf{t} \in \mathbb{R}^p$, $p \in \mathbb{N}$. We aim at synthesising regions (to be denoted as $\mathcal{T}$) within the parameters space of $\mathbf{t}$, such that $\mathscr{F}(\mathbf{t})$ satisfies a given property ψ [1]. We shall abstract the SDE into a finite-state Markov model with parameters $\mathbf{v}$, and automatically compute set $\mathcal{V}$, such that the corresponding abstract models satisfy ϕ, as outlined in Fig. 7.1. Properties ψ and ϕ for the SDE and the abstract Markov model, respectively, are related via the abstraction error ϵ_{abs}, as detailed in the following section. Finally, we translate the synthesised set $\mathcal{V}$ into a region $\mathcal{T}$ for the vector $\mathbf{t}$ of the parametric SDE. Notice that, in general, $\mathbf{v} \neq \mathbf{t}$: $\mathbf{t}$ represents the parameters vector for the original SDE, whereas $\mathbf{v}$ represents the

[1] We use ψ to denote a requirement specified on the SDE, whereas we use ϕ to denote a requirement specified on the Markov model.

parameters entering the abstract transition probability matrix: we expect $\mathbf{v}$ to be a function of $\mathbf{t}$, say $\mathbf{v} = g(\mathbf{t})$. A key contribution of this Chapter is in finding g automatically, for classes of SDE.

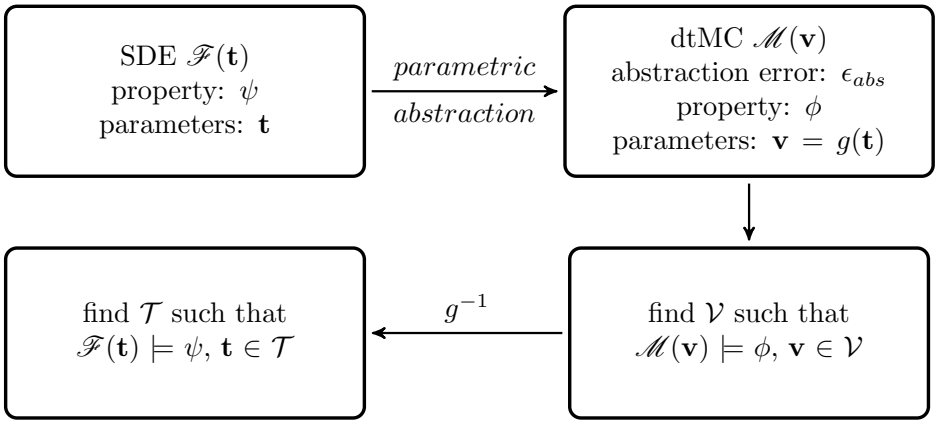

Figure 7.1: Depiction of the procedure for abstraction and parameter synthesis.

In this work, we focus on a discrete-time, parametric SDE of the following form:

$$x(k+1) = f(x(k), \mathbf{t}) + \omega_{\mathbf{t}}(k), \tag{7.1}$$

where $k \in \mathbb{Z}$ indexes the discrete time, $x \in \mathbb{R}^n$ represents the state space variable, f is a continuous vector field $f : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$, and ω denotes an additive noise term, also a function of the parameter vector $\mathbf{t} \in \mathbb{R}^p$.

Let us introduce the unsafe set $\mathcal{Q}_u$, such that when $x \in \mathcal{Q}_u$ the system in Eq. (7.1) is in a *bad* condition. We consider probabilistic safety properties [69] (see Section 2.2.4), which can be encoded formally as

$$\psi = \mathsf{P}_{\leq \eta}[\Diamond_{\leq h}\ \mathcal{Q}_u],\ h \in \mathbb{N}, \tag{7.2}$$

that is satisfied if the probability of reaching[2] the unsafe set $\mathcal{Q}_u$, within h time steps, is smaller than a given constant η.

Parametric Abstraction Error Let us now introduce a special class of the parametric SDEs in Eq. (7.1), namely

$$x(k+1) = \theta + \rho f(x(k)) + \sigma\omega(k), \tag{7.3}$$

[2]We recall that the symbol $\Diamond_{\leq h}$ denotes the reachability within h steps.

where we have further specialised how the parameters in the vector $\mathbf{t} = [\theta, \rho, \sigma]$ enter the dynamics. This parametric SDE, as anticipated in above, deals with:

- θ, an additive parameter;
- ρ, a parameter that is multiplicative on the deterministic drift;
- σ, a parameter that is multiplicative on the noise term.

In a parametric setting, term l_w (see Eq. (2.6)), the Lipschitz constant of the noise kernel, becomes a function of σ: formally we may denote it as $l_w(\sigma)$, whose analytical expression can be computed during a first pre-processing step. As a consequence, the abstraction error ϵ_{abs} depends on parameter σ and is a function of δ_q. We can rewrite Eq. (2.8) as

$$\delta_q \cdot l_w(\sigma) \leq \frac{\epsilon^{max}}{h\mathscr{L}(\mathcal{Q}_u)}, \tag{7.4}$$

where δ_q is the diameter of the partitions of $\mathcal{Q}_u$, ϵ^{max} is the maximum desired abstraction error, h is the time horizon and $\mathscr{L}(\mathcal{Q}_u)$ represents the diameter of set $\mathcal{Q}_u$. Recall that the canonical formal abstraction procedure has one degree of freedom – the partition size δ_q – to finely tune the desired abstraction error. Equation (7.4) instead has both δ_q and $l_w(\sigma)$ on the left-hand side: the abstraction error depends on the parameter σ, which must be synthesised together with the partition size δ_q.

Once we set a value for δ_q and perform the abstraction, we implicitly define a range of values where σ can dwell. Furthermore, the abstraction procedure induces the map g, which translates parameters $\mathbf{v}$ into the abstract parameters $\mathbf{t}$. The actual form of function g depends on the kind of parameter – whether additive, multiplicative – and, above all, on the probability density function (pdf) of the noise. As we discuss in the following sections, a distribution function that can be written as a combination of elementary functions (cf. Section 7.2) is a key feature for the construction of function g.

Abstraction Error within the Probability Bound We now show how to account for the error ϵ_{abs} carried by the abstraction procedure. It can be shown [65] in quite some generality that if the abstract system satisfies a given specification with probability η, the probability of verifying a corresponding property for the concrete SDE lies within $\eta \pm \epsilon_{abs}$. In particular for safety properties, where we are interested in upper bounds for the probability to reach unsafe sets, we can ensure that

$$\mathscr{M} \models \phi : \mathsf{P}_{\leq \eta - \epsilon_{abs}}[\Diamond_{\leq h}\ \mathcal{S}_u] \Longrightarrow \mathscr{F} \models \mathsf{P}_{\leq \eta}[\Diamond_{\leq h}\ \mathcal{Q}_u],\ h \in \mathbb{N},$$

where $\mathcal{S}_u \subset \mathcal{S}$ represents the set of abstract unsafe states corresponding to the continuous unsafe region $\mathcal{Q}_u \subset \mathcal{Q}$. Thus, the satisfaction of ψ for $\mathscr{F}$ (cf. Eq. (7.2)) is guaranteed by the stricter bound on the Markov chain $\mathscr{M}$ that accounts for the abstraction error ϵ_{abs}.

The parameter synthesis on $\mathscr{M}$ defines a satisfiability region $\mathcal{V}$ that is then mapped to parameters $\mathbf{t}$ of the continuous system via g^{-1}, resulting in a region $\mathcal{T}$: if mapping g is analytical, then this can be done explicitly. The final step of the parametric abstraction procedure checks the intersection of region $\mathcal{T}$ and conditions (7.4). If this is non-empty, it defines the satisfiability region of ψ; if, conversely, the intersection is empty, the abstraction must be refined and the procedure repeated until a stopping criterion – e.g. a minimum value δ_q – is met.

7.2 Parametric Abstractions with Elementary Distributions

The formal abstraction procedure described above translates a continuous SDE into a Markov chain $(\mathcal{S}, P)$ via kernel marginalisation, as per Eq. (2.5). As a result, entries of P are evaluations of the noise kernel; if we consider a parametric SDEs, the entries of P become parametric functions: the shape of the cumulative distribution function (cdf) of ω defines function $g(\cdot)$, namely the relation between parameters $\mathbf{v}$ and $\mathbf{t}$. As an example, a polynomial cdf of ω results in polynomial expressions for P. The mapping g is crucial, as its shape contributes to the complexity of the subsequent parameter synthesis procedure.

A relevant class of functions for this work is the class of elementary functions. It consists of polynomials, exponential, logarithmic and trigonometric functions, together with their inverse, equipped with operations of sum, product and composition applied finitely many times. Example of distributions that use these functions are the uniform and triangular distributions.

Unfortunately, many well-known cumulative distribution functions cannot be expressed in terms of elementary functions, as the Gaussian or χ^2 distributions. For these, an analytical inverse cdf function does not exist, thus we have no symbolic way of extracting the parameters from these.

Example 3 (Computing Transition Probabilities) *Consider a stochastic parametric SDE as*

$$x(k+1) = \sigma\omega(k),$$

where $\sigma \in \mathbb{R}$ is a parameter and ω is a zero-mean noise. Assume we aim at checking the probability of $x(1) \leq 1$ and synthesise the values of σ such that this probability is less than 0.75. The formal property of interest is $\psi := \mathsf{P}[x \leq 1] \leq 0.75$. The probability $x(1) \leq 1$ can be formally written

$$\mathsf{P}[x \leq 1] = \int_{-\infty}^{1} t_\omega(u)\, du = \mathscr{C}_\omega(1) - \mathscr{C}_\omega(-\infty), \tag{7.5}$$

where t_ω is the probability density of ω and $\mathscr{C}_\omega$ represents its cumulative function.

Let us now consider the two scenarios where ω belongs to the uniform or to the Gaussian distribution. These two are rather different: a uniform distribution has a finite domain and a linear pdf, whereas the Gaussian has infinite support and exponential pdf. If ω belongs to a uniform, by definition [183] we gather

$$\mathscr{C}_{\mathcal{U}}(x) = \frac{x - a(\sigma)}{b(\sigma) - a(\sigma)}, \tag{7.6}$$

where a and b are the lower and upper bound of the domain of ω. They are a function of the distribution standard deviation σ, hence the notation $a(\sigma)$ and $b(\sigma)$. In particular, we obtain [3] $a(\sigma) = -\sqrt{3}\sigma$ and $b = \sqrt{3}\sigma$. On the other hand,

$$\mathscr{C}_G(x) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x}{\sigma\sqrt{2}}\right)\right], \tag{7.7}$$

[3] We show the complete derivation of $a(\sigma)$ and $b(\sigma)$ in the following Section.

if ω belongs to a Gaussian, where erf *represents the so-called error function and σ is the standard deviation of ω. We may safely assume $\mathscr{C}(-\infty) = 0$ in both cases. In the uniform case we have*

$$\mathsf{P}[x \leq 1] = \frac{1 + \sqrt{3}\sigma}{2\sqrt{3}\sigma} \leq 0.75 \Longrightarrow \sigma \geq \frac{2}{\sqrt{3}} \simeq 1.15,$$

whereas in the Gaussian case

$$\mathsf{P}[x \leq 1] = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{1}{\sigma\sqrt{2}}\right)\right] \leq 0.75 \Longrightarrow \sigma \geq \frac{1}{\sqrt{2}\operatorname{erf}^{-1}(0.5)} \simeq 1.48.$$

Note that this simple case offers an elegant solution for both scenarios. However, when considering slightly more complex expressions, we may not offer equally elegant results. As an example, assume the derivation of $\mathsf{P}[x \leq 1]$ is composed of a sum of similar terms, instead of a single one – as it is the case using the formal abstraction technique. If ω belongs to a uniform distribution, we deal with the sum

$$\mathscr{C}_{\mathcal{U}}(x_1) + \mathscr{C}_{\mathcal{U}}(x_2) = \frac{x_1 + \sqrt{3}\sigma}{2\sqrt{3}\sigma} + \frac{x_2 + \sqrt{3}\sigma}{2\sqrt{3}\sigma}, \tag{7.8}$$

that offers a rational expression of σ. We are hence able to extract the value of σ such that the sum is less than 0.75, as the specification suggests.

On the other hand, the erf *function does not have associative properties. Hence, we are left with an expression*

$$\mathscr{C}_G(x_1) + \mathscr{C}_G(x_2) = 1 + \frac{1}{2}\left[\operatorname{erf}\left(\frac{x_1}{\sigma\sqrt{2}}\right) + \operatorname{erf}\left(\frac{x_2}{\sigma\sqrt{2}}\right)\right], \tag{7.9}$$

that cannot be simplified further. We are thus unable to recover the value of σ that satisfies any property without using numerical methods. □

A naïve workaround is to replace every non-elementary parametric element of the stochastic kernel – in the previous example, every $\operatorname{erf}(\cdot)$ term – with a new parameter, reason with the new parameters and compute their numerical values. However, beyond the introduced numerical approximations, this approach requires a new parameter for every non-elementary term: the number of such substitutions would quickly grow with the order of n^2, i.e. one (or more) parameter per entry of P, making the synthesis practically impossible to scale.

Another approach, which is explored in the following Section, approximates the non-elementary cdf with elementary functions, such as polynomials. This introduces an approximation error to be accounted for, while allowing to control the number of parameters and to apply standard parameter synthesis procedures.

If, on the other hand, the cdf can be represented in terms of elementary functions, the parameter synthesis step can follow standard procedures, within the limits of available theories and software tools: polynomial or rational cdf's are tractable problems, whereas trigonometric or exponential distributions are beyond the reach of existing tools, being either theoretically undecidable or practically not scalable. From a computational point of view, stochastic noises that have polynomial distributions – as the uniform distribution or the more general Irwin-Hall distribution [204] (the sum of independent uniform distributions) – offer both elegant results and ease of computation. In the following, we illustrate the special case of a uniform noise, deriving the abstract model and discussing the computation of a probabilistic safety property.

7.2.1 Piecewise Linear Distribution Function: Uniform Noise

In the following, we consider noise ω belonging to a uniform distribution; for simplicity and ease of understanding, we assume a scalar system: referring to Eq. (7.1), $x \in \mathbb{R}$. The uniform distribution, with its rectangular shape, presents useful features, as an elegant connection between its probability value and the size of the partition intervals.

Parametric Uniform Distribution As mentioned in the previous Section, the uniform distribution has a finite domain: let us define it as $D = [-\lambda, \lambda]$ so that the total length results in $|D| = 2\lambda$. By definition of uniform distribution [183], we may write the expression of the variance v

$$v = \frac{\lambda^2}{3} = \frac{|D|^2}{12} \Longrightarrow \lambda = \sqrt{3v}. \tag{7.10}$$

By definition of probability distribution, we know that the integral of any probability density over the domain is equal to one. Since the domain size is $|D| = 2\lambda$, let us define $\alpha = 0.5\lambda^{-1}$ – i.e. the point-wise value of probability – so that $|D| \cdot \alpha = 1$.

Considering Eq. (7.3), we argue that term $\sigma\omega$ has variance equal to σ^2. We may write the relationship between α and σ, using Eq. (7.10), as

$$\lambda = \sqrt{3}\sigma, \text{ so that } \alpha = 0.5\lambda^{-1} = (2\sqrt{3}\,\sigma)^{-1}. \tag{7.11}$$

In practical terms, we find that the domain size $|D|$ is proportional to the standard deviation as $|D| = 2\lambda = 2\sqrt{3}\sigma$.

After a few observations on the features offered by a uniform noise, we are ready to tackle the whole SDE reported in Eq. (7.3),

$$x(k+1) = \theta + \rho f(x(k)) + \sigma \cdot \omega(k).$$

We notice that, whenever ω belongs to a standard [4] distribution $\mathcal{P}(0,1)$, we may represent the value of $x(k+1)$ as a new random variable, whose mean is $\theta + \rho f(x(k))$ and variance is σ^2; formally, $x(k+1) \in \mathcal{P}(\theta + \rho f(x(k)), \sigma^2)$.

Since we consider a uniform noise ω, we conclude that $x(k+1)$ also belongs to a uniform distribution – with the proper average value and variance. Let us introduce a parametric uniform pdf t_w:

$$t_w(u) = \begin{cases} 0 & \text{if } u < -\lambda + \theta + \rho f(x(k)) \\ & \text{or } u > \lambda + \theta + \rho f(x(k)), \\ \alpha & \text{otherwise,} \end{cases} \tag{7.12}$$

where α is the point-wise probability density value and λ is the domain half-width, as defined in (7.11). The piecewise constant shape of t_w will grant significant simplifications for the formal abstraction technique.

Formal Abstractions: State-space Partitions We may now begin our formal abstraction procedure. The first step is the construction of state-space partitions: since we consider $x \in \mathbb{R}$, each partition is simply an interval of the real line. We

[4] Standard distributions have zero mean and variance equal to one.

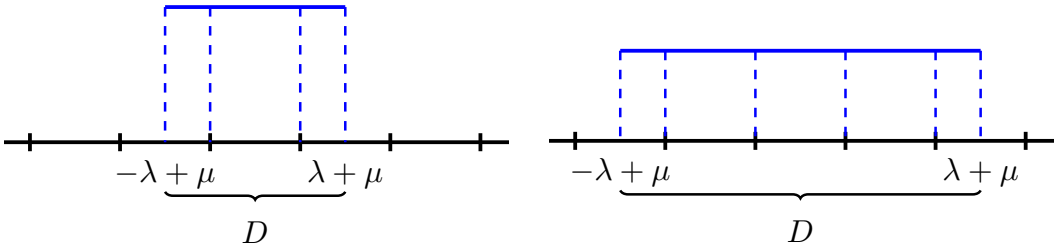

Figure 7.2: Uniform noise probability distribution domain with parametric variance.

thus focus on the transitions between partitions, i.e. intervals of $\mathbb{R}$. Transitions from a partition can only reach partitions where the density t_w is positive – recall that the transition probabilities are computed as the integral of the noise distribution over the partitions. In view of the finite length of domain D, only a limited number of partitions have a positive transition probability, whereas many others have transition probability equal to zero: intuitively, partitions that are far from the current one cannot be reached. On the other hand, the domain D is proportional to σ, the noise standard deviation, as per Eq. (7.10); therefore an increase in variance leads to an increase of $|D|$ (and viceversa). The wider $|D|$, the more partitions are reachable, resulting in more transitions with non-zero probability in the obtained abstract model. Figure 7.2 shows the uniform probability distribution domain when ω has a small variance (left) and a larger variance (right). It depicts a uniform probability density function with average $\mu = \theta + \rho f(x)$: the domain is then $D = [-\lambda + \mu, \lambda + \mu]$. In this example, a small variance confines the noise domain into only 3 partitions (left), whereas a larger variance extends the noise domain into more 5 partitions (right). Starting from the central interval, we compute transitions towards 3 partitions (left) or towards 5 partitions (right).

Following the formal abstraction procedure, each partition q_i is translated into an abstract state s_i. The abstract states compose the dtMC, which will be then employed to formally verify any property of interest.

Transitions between Abstract States We finally ought to evaluate the transition between abstract states. For each partition q_i we select a representative

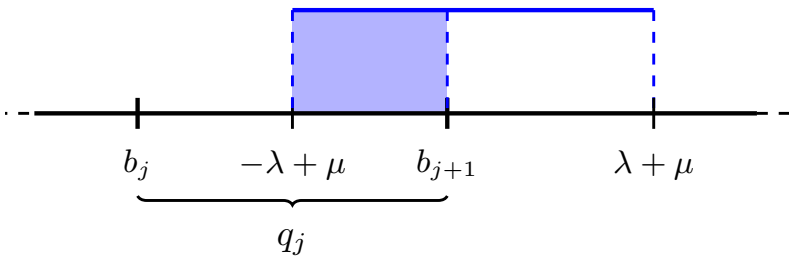

Figure 7.3: The light blue area represents the (s_i, s_j) transition probability.

point r_i that numerically represents the "starting point" of the transitions. Roughly speaking, in order to evaluate the transition from q_i to q_j, we compute the integral of the noise, centred at r_i, over q_j. We compute the transition probability $P(s_i, s_j)$ for each pair of abstract states, to ultimately compose the transition matrix P. Let us now consider a general transition from abstract state s_i, corresponding to partition q_i with representative point r_i, to state s_j, corresponding to partition q_j whose boundary points we denote as b_j and b_{j+1}. In view of the constant probability density of Eq. (7.12), the (s_i, s_j) transition probability can be written as $\alpha \cdot \nu_{i,j}$, where $\nu_{i,j}$ is the portion of partition q_j covered by the probability distribution, as depicted in Fig. 7.3.

The lower boundary of $\nu_{i,j}$ evaluates to

$$\max\{b_j, -\lambda+\mu\},$$

the maximum between b_j, i.e. the lower bound of partition q_j, and $\mu-\lambda$, i.e. the lower bound of the probability distribution. The upper bound of $\nu_{i,j}$ is conversely

$$\min\{b_{j+1}, \lambda+\mu\},$$

the minimum between b_{j+1}, i.e. the upper bound of partition q_j, and $\mu+\lambda$, the upper bound of the distribution. Finally, $\nu_{i,j}$ must be non-negative. Summarising, we recall that the mean μ_i evaluates to $\theta+\rho f(r_i)$, thus the transition values can be written as

$$\begin{aligned} P(s_i, s_j) &= \alpha \cdot \nu_{i,j} \\ &= \alpha \cdot \max\{0, \min\{b_{j+1}, \lambda+\theta+\rho f(r_i)\} - \max\{b_j, -\lambda+\theta+\rho f(r_i)\}\}. \end{aligned} \quad (7.13)$$

Every entry of the transition matrix $P(\mathbf{v})$ is a piecewise linear function of parameters λ, θ and ρ. The abstract system thus depends on four parameters $\mathbf{v} = [\alpha, \lambda, \theta, \rho]$, where the function $g(\mathbf{t}) = \mathbf{v}$ maps

$$g(\alpha, \lambda, \theta, \rho): \quad \alpha = (2\sqrt{3}\sigma)^{-1}, \quad \lambda = \sqrt{3}\sigma, \quad \theta = \theta, \quad \rho = \rho, \tag{7.14}$$

notice that the special case of uniform noise allows writing the transition matrix $P(\mathbf{t})$ as a rational function of the parameters in $\mathbf{v}$.

We have showed the formal abstraction procedure in the scalar case. For higher-dimensional SDE models (where $x \in \mathbb{R}^n$), similar results can be obtained applying the max and min operations to every component independently.

Solving Parameter Synthesis Problems A parameter synthesis problem for piecewise linear functions of parameters can be solved via "case splitting". This entails binary branching for every max or min clause, until only linear functions remain. Standard techniques from the synthesis literature can then be applied, e.g. Gaussian elimination [205], transitions or state elimination [71] and ETR encoding [76]. However, the number of binary cases introduced is, in the worst case, exponential in the number of clauses.

Alternatively, we follow an approach drawing from the reachability analysis literature, e.g. [206]. Assume we aim at computing an h-step probabilistic reachability property. We compute the satisfiability region $\mathcal{V}_i$ for every step i and use the induced constraints to simplify the entries in the transition matrix P. We then iteratively repeat the procedure over the h-step horizon, obtaining at each step stricter constraints, in view of the chain of inclusions $\mathcal{V}_i \supseteq \mathcal{V}_{i+1}, \forall i$.

7.3 Parametric Abstractions with Non-Elementary Distributions

As discussed at the beginning of section 7.2, the shape of the noise's cdf is crucial for the synthesis of the parameters, both of the abstract and of the original system. In practice, in order to symbolically perform parameter synthesis, the

noise's distribution function should be invertible, or at least show some simplifying properties. The most common random variables, from the Gaussian distribution to the χ^2-distribution, have cumulative distribution functions that cannot be evaluated in closed form. Therefore, the parameter synthesis cannot be performed analytically and we must resort to approximations.

With focus on the Gaussian noise, many polynomial approximations of its distribution function F are available in the literature: see, among others, [207]. However, polynomial or rational approximations are reliable only in a finite domain: still, as the natural setting of the abstraction-partitioning procedure is a finite set $\mathcal{A}$, a polynomial approximation of F over $\mathcal{A}$ can be satisfactory. By its definition, the Gaussian cdf is a linear combination of the so-called erf functions – that was first invented specifically to denote the Gaussian cumulative distribution. It can be approximated using several approaches; we select here a polynomial expansion via the (truncated) MacLaurin series as

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \left(x - \frac{x^3}{3} + \frac{x^5}{10} \dots \right), \tag{7.15}$$

or via rational expansions [208], e.g.

$$\mathrm{erf}(x) \simeq 1 - \frac{1}{(1 + a_1 + a_2 x^2 + a_3 x^3 + a_4 x^4)^4}, \tag{7.16}$$

where coefficients a_i, $i = 1, \dots, 4$ are given.

Recalling Eq. (7.9) of Example 3, we may employ these approximations to substitute the expressions of erf within the computation of the transition probabilities. This will allows us to extract the values of the parameters in a symbolic manner, hence to evaluate the parameters that satisfy a given (safety) property. Naturally, a high-order polynomial approximation holds high accuracy but, on the other hand, limits the computational tractability of parameter synthesis. Related to the order of the polynomial, the error introduced by these approximations can be quantified: typically, rational functions have an approximation error that can be upper bounded by some finite constant ξ.

Whenever we apply the polynomial approximation to the transition values, we attach a ξ error to every entry of the transition matrix P. Hence, dealing with a n-dimensional transition matrix carries an approximation error upper bounded by ξn^2 – the number of entries of a transition matrix is indeed n^2. A total of h operations with P accounts for an overall error that can be upper bounded by $h \cdot \xi \cdot n^2$.

We ought to attach this additional approximation error to the formal abstraction procedure. Practically speaking, we consider a safety property with bound $\eta + h\xi n^2$.

7.4 Experimental Evaluation

In this section we outline case studies including polynomially-distributed (uniform) and non-elementary-distributed (Gaussian) noise. Remarkably, we offer a parameter synthesis case study considering the (simplified) power grid model. Despite the high complexity and non-trivial dynamics, we synthesise the heterogeneity of the solar population under a uniform-noise scenario. Drawing inspiration from literature on probabilistic model checking of SDE [65], we focus on the h-step probabilistic reachability probability property. At each step, the set of *feasible parameters* is evaluated and used to simplify entries of matrix P.

7.4.1 Formal Abstractions with Parametric Uniform Noise

We now illustrate a complete parametric formal abstraction procedure. We consider three parameters θ, ρ, σ, and aim at checking a probabilistic safety property. To this end, we introduce the parametric SDE

$$x(k+1) = \theta + \rho \cdot x(k) + \sigma \cdot \omega(k),$$

where θ is an additive parameter, ρ is a multiplicative parameter, and σ affects the variance of the noise $\omega \in \mathcal{U}(0,1)$. Consider a safe set $\mathcal{A} = [0,1)$, or dually the unsafe set $\mathcal{Q}_u = (-\infty, 0) \cup [1, \infty)$, and the probabilistic safety specification

$$\psi = \mathsf{P}_{\leq 0.5}\,[\, x(k+1) \in \mathcal{Q}_u \,].$$

First of all, let us compute the (parametric) abstraction error. Recalling Eq. (7.4), we consider $h = 1$ and $\mathscr{L}(\mathcal{A}) = 1$, hence an upper bound for the abstraction error ϵ_{abs} results in

$$\epsilon_{abs} \leq \delta_A \cdot l_w(\sigma) = \delta_A \cdot \frac{1}{2\sqrt{3}\sigma} = \delta_A \cdot \alpha, \tag{7.17}$$

where we have used the relationship in Eq. (7.11). Note that this bound is conservative, as it accounts for the probability mass over the whole partition. Considering a reference point placed in the centre of mass of the partition, the abstraction error can be halved, hence becoming $\epsilon_{abs} \leq 0.5\,\alpha\,\delta_A$.

Let us select a (desired) maximum abstraction error $\epsilon^{max} \leq 0.25$. From Eq. (7.17), we ought to "split" this value between δ_A and l_w: we choose $\delta_A = \sqrt{\epsilon^{max}} = 0.5$ leaving $l_w \leq \sqrt{\epsilon^{max}}$, resulting in $\sigma \geq \sqrt{3}^{-1}$.

Let us now partition $\mathcal{A}$: we have selected size $\delta_A = 0.5$ hence we construct two partitions, $\mathcal{A}_1 = [0, 0.5)$ and $\mathcal{A}_2 = [0.5, 1)$, with boundary points $b_1 = 0$, $b_2 = 0.5$, $b_3 = 1$ and representative points $r_1 = 0.25$ and $r_2 = 0.75$. These are represented by the abstract states s_1, s_2, respectively, whereas the unsafe set $\mathcal{Q}_u$ is represented by s_3 – we do not need representative points for the unsafe set. The transition probability matrix P may be written as

$$P = \begin{bmatrix} p_{11} & p_{12} & 1 - p_{11} - p_{12} \\ p_{21} & p_{22} & 1 - p_{21} - p_{22} \\ 0 & 0 & 1 \end{bmatrix},$$

where $p_{i,j} = \alpha \max\{0, y_{i,j}\}$, $y_{i,j} = \min\{b_{j+1}, \lambda + \theta + \rho r_i\} - \max\{b_j, \theta + \rho r_i - \lambda\}$ in view of Eq. (7.13).

Assume the initial state is $x_0 = 0$, which is mapped to the abstract state is s_1. To account for the abstraction, we aim at verifying that the probability of reaching $\mathcal{Q}_3$ is smaller than $\eta = 0.5$; we ought to subtract the abstraction error $\epsilon^{max} = 0.25$, hence the final bound results $\eta - \epsilon^{max} = 0.25$. Dually, we might evaluate the probability of remaining within s_1 or s_2 greater or equal than 0.75. We are to solve the inequality

$$p_{11} + p_{12} = \alpha[\max\{0, y_{11}\} + \max\{0, y_{12}\}] \geq 0.75 \Leftrightarrow$$
$$\max\{0, y_{11}\} + \max\{0, y_{12}\} - 1.5\lambda = c(\theta, \rho, \lambda) \geq 0, \tag{7.18}$$

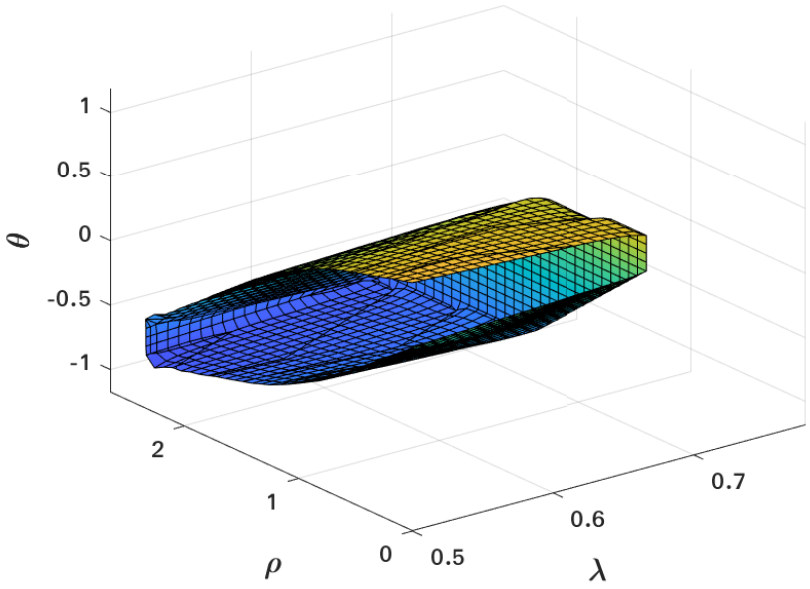

Figure 7.4: Depiction of the surface $c(\theta, \rho, \lambda) = 0$ of the satisfiability region.

where we substitute α as per Eq. (7.11). In particular, the set $c(\theta, \rho, \lambda)$ offers solutions only for values of $\lambda \leq 2/3$, that is $\sigma \leq 0.34$, which is disjoint from the initial setting $\sigma \geq 1/\sqrt{3}$. This instance of the parametric verification offers no solution.

We then consider a smaller partition step $\delta_A = 0.25$, which results in a wider initial constraint $\sigma \geq (2 \cdot \sqrt{3})^{-1}$, that is $\lambda \geq 0.5$. The reachability property is now satisfiable with the upper probability bound $0.5 + \epsilon^{max} = 0.75$, which results in the set depicted in Fig. 7.4.

7.4.2 Formal Abstractions with Parametric Gaussian Noise

We now discuss an abstraction procedure of an SDE in presence of a Gaussian noise. Differently from the uniform noise, a Gaussian distribution has exponential probability density with infinite support.

Let us consider a safe set $\mathcal{A} = [0, 1)$, the unsafe set $\mathcal{Q}_u = (-\infty, 0) \cup [1, \infty)$ and the SDE model in Eq. (7.3)

$$x(k+1) = \theta + \rho x(k) + \sigma w(k),$$

where $w \in \mathcal{N}(0,1)$ and initial state $x_0 = 0$. Once again, we consider the safety property

$$\psi = \mathsf{P}_{\leq 0.5}\,[\,x(k+1) \in \mathcal{Q}_u\,]. \tag{7.19}$$

We set a maximum abstraction error $\epsilon^{max} = 0.25$. Noting that $h = 1$ and $\mathscr{L}(\mathcal{A}) = 1$, we may upper bound the abstraction error as

$$\epsilon_{abs} \leq \delta_A\, l_w(\sigma) = \delta_A\, \frac{1}{\sqrt{2\pi}\sigma^3}.$$

We choose to "split" equally the upper bound between δ_A and l_w, as

$$\delta_A = \sqrt{\epsilon^{max}} = 0.5, \quad \sigma^3 \geq \frac{1}{\sqrt{2\pi\epsilon^{max}}} \Longrightarrow \sigma \geq (0.5\pi)^{-1/6},$$

defining a first feasibility region for σ. In view of $\delta_A = 0.5$, we divide $\mathcal{A}$ into two partitions $\mathcal{A}_1 = [0, 0.5)$, $\mathcal{A}_2 = [0.5, 1)$ that are represented by the abstract states s_1, s_2, respectively. The unsafe set $\mathcal{Q}_u$ is represented by the abstract state s_3.

Partitions $\mathcal{A}_1$ and $\mathcal{A}_2$ have boundary points $b_1 = 0$, $b_2 = 0.5$, $b_3 = 1$ (b_2 is shared) and representative points $r_1 = 0.25$ and $r_2 = 0.75$. The transition probability matrix results in

$$P = \begin{bmatrix} e_{11} & e_{12} & e_{13} \\ e_{21} & e_{22} & e_{23} \\ 0 & 0 & 1 \end{bmatrix}, \tag{7.20}$$

where the elements $e_{i,j}$ can be written as

$$e_{i,j} = \frac{1}{2}\left[\operatorname{erf}\left(\frac{b_j - (\rho f(r_i) - \theta)}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{b_i - (\rho f(r_i) - \theta)}{\sqrt{2}\sigma}\right)\right],$$

for $i, j = 1, 2$, whereas $e_{i\,3} = 1 - e_{i\,1} - e_{i\,2}$.

As discussed in Section 7.3, we approximate the erf function with the polynomial expansion in Eq. (7.15): the approximation error is bounded from above by 10^{-4} [208]. We note 4 terms ($e_{i,j}$ for $i, j = 1, 2$) introducing such approximation from the transition matrix in (7.20). We account for this error by increasing the safety bound. The original specification requires a probability bound of 0.5, with abstraction error 0.25 and approximation $4 \cdot 10^{-4}$: the abstract probability bound η results in 0.2496. We aim at synthesising parameters on the dtMC that verify

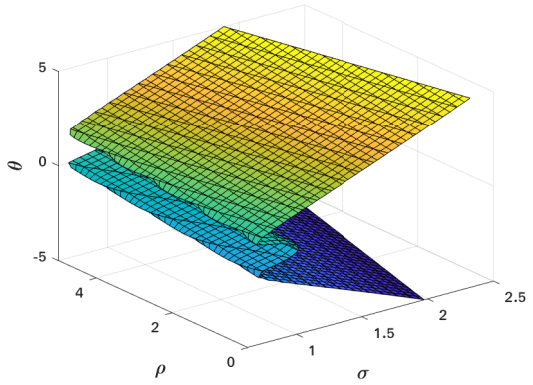

Figure 7.5: Plot of the surface $c(\theta, \rho, \sigma) = 0.2496$ delimiting the satisfiability region.

$\mathsf{P}_{\leq 0.2496}[\,\text{next state} = s_3\,]$: this can be written as an inequality $c(\theta, \rho, \sigma) \geq 0.2496$, where $c(\theta, \rho, \sigma)$ is the polynomial function

$$\begin{aligned} c(\theta,\rho,\sigma) = \sqrt{2}[(\sigma^4(30\rho - 120\theta) + \sigma^4(-30\rho + 120\theta + 120) \\ - 20\sigma^2(0.25\rho - \theta)^3 - 20\sigma^2(-0.25\rho + \theta + 1)^3 \\ + 3(0.25\rho - \theta)^5 + 3(-0.25\rho + \theta + 1)^5)] \cdot (120\sqrt{\pi}\sigma^5)^{-1}. \end{aligned}$$

The surface $c(\theta, \rho, \sigma) = 0.2496$ is depicted in Fig. 7.5. This region of parameter space, combined with set $\sigma \geq (0.5\,\pi)^{-1/6} \simeq 0.9275$, defines the parameters that guarantee that $\mathsf{P}_{\leq 0.5}[x(k+1) \in \mathcal{Q}_u]$ is satisfied.

7.4.3 Formal Abstractions of the Power Grid Dynamics and a Parametric Aggregation of PV Systems

Finally, we outline a formal abstraction procedure for a parametric version of Eq.(5.11). As mentioned in the previous chapters, the model checking techniques suffers from the abundance of parameters in every scenario. In order to succeed, however, we leverage a few simplifying assumptions on the underlying model in (5.11): we will consider uniform reconnection and disconnection thresholds, uniform noise and we will ease the frequency dynamics.

Our analysis focusses on the synthesis of the variance of a given threshold distribution. Practically speaking, we assume a parametric disconnection and reconnection distribution, i.e. parametric functions $a(k)$ and $b(k)$ of the system of

equations (5.11). For simplicity, we consider the disconnection and reconnection thresholds (and correspondingly $a(k)$ and $b(k)$) belonging to a uniform distribution. We recall that $a(k)$ and $b(k)$ are the cumulative distribution function of a uniform distribution, which may be written as

$$\begin{aligned} a(k) &= \begin{cases} 0 & \text{if } f(k) < -\lambda + \mu, \\ \dfrac{f(k) + \lambda - \mu}{2\lambda} & \text{if } |f(k) - \mu| \leq \lambda, \\ 1 & \text{if } f(k) > \lambda + \mu, \end{cases} \\ &= \min\left\{0, \max\left\{1, \frac{f(k) + \lambda - \mu}{2\lambda}\right\}\right\}, \end{aligned} \tag{7.21}$$

where μ represents the average value and λ the half-length of the distribution domain; similarly we may rewrite $b(k)$. This alternative formulation allows an easier handling of the the resulting parametric transitions.

Notice that this setting is more complex than the ones analysed earlier: the parameters λ and μ are argument of a nonlinear function.

For clarity and ease of understanding, we outline a reduced version of the closed-loop dynamics, as

$$\begin{cases} \Delta f(k+1) = \bar{\alpha}\Delta f(k) + \bar{\beta}\Delta P_{PV}(k) + \omega_f(k), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)), \\ P_{PV}(k) = \bar{P}Nx(k) + \omega_P(k), \end{cases} \tag{7.22}$$

where we have simplified the dynamics of the network frequency. The frequency now evolves according to a first-order difference equation, with coefficients $\overline{\alpha}$, $\overline{\beta}$. These coefficients have been chosen in order to make the reduced transfer function

$$\overline{G}(z) = \frac{\overline{\beta}}{z - \overline{\alpha}}, \tag{7.23}$$

dynamically similar to the original $G(z)$ of Eq. (5.7). Practically speaking, we are interested in the response of $G(z)$ and $\overline{G}(z)$ in terms of settling time (cf. Section 5.2) and final value [5]. To this end, we recall [198] that the settling time t_s for

[5] The final value is the value of a function as time approaches infinity.

a first-order transfer function is

$$\overline{G}(z) = \frac{\overline{\beta}}{z - \overline{\alpha}} \quad \Longrightarrow \quad t_s \simeq \frac{-4h}{\ln|\overline{\alpha}|}, \tag{7.24}$$

where h is the sampling time of the discrete-time framework: in our setting $h = 0.2$ seconds. Using the equivalent method for a second-order transfer function, $G(z)$ provides a settling time of approximately 7.6 seconds. Thus, we equate

$$\frac{-4h}{\ln|\overline{\alpha}|} = 7.6 \quad \Longrightarrow \quad \overline{\alpha} = \pm 0.90. \tag{7.25}$$

We choose the positive value $\overline{\alpha} = 0.90$ as it guarantees less oscillations [198] with respect to the choice $\overline{\alpha} = -0.90$. In order to have the same final value, we set

$$G(1) = \overline{G}(1) \quad \Longrightarrow \quad \overline{\beta} = 0.0315. \tag{7.26}$$

To further help the synthesising procedure, we consider ω_f and ω_P belonging to a uniform distribution. Notice that these stochastic signals are not parametric, and the only parameters are within $a(k)$ and $b(k)$.

Let us define the state variable $q = [\Delta f, \Delta P_{PV}, y]$, where we can omit x as it is strictly related to ΔP_{PV}. We recall from Section 6.1 the definitions of $\mathbb{F} = [-0.8, 0.8]$, $\mathbb{P} = [0, \bar{P}N]$, $\mathbb{X} = [0, 1]$ and we define the three-dimensional safe set $\mathcal{A} = \mathbb{F} \times \mathbb{P} \times \mathbb{X}$. Dually, we define the unsafe set $\mathcal{Q}_3$ that represents a load shedding event. We produce a partitioning grid with $\delta_A = 0.01$ that ensures an abstraction error $\epsilon_{abs} = 0.1$, as per Eq. (7.17).

The network settings of this test are as follows: we assume a generation incident of 3 GW in a network with load demand of 220 GW; the solar penetration is set to 20% of the load; we set the average value of the disconnection equal to $\mu = -0.2$, which corresponds to a disconnection at 49.8 Hz. We aim at computing the variance of the uniform noise σ^2 that guarantees a load shedding probability smaller than 0.1 after 2 time steps.

Our procedure returns a nonlinear inequality in view of the presence of the min and max functions. The actual expression is rather complex and does not

offer much insight, hence we omit it from the present dissertation. The inequality solver offers the symbolic solution

$$\sigma \leq 0.080, \tag{7.27}$$

which is then corroborated by the numerical evaluation depicted in Fig. 7.6a. We read this result as a requirement for a low degree of population heterogeneity. In concert with the previous chapter's results, a uniform distribution of thresholds with high variance causes a greater number of disconnections: the upper bound of the distribution is in fact closer to the nominal frequency f_0. Smaller deviations are sufficient to start the disconnection process.

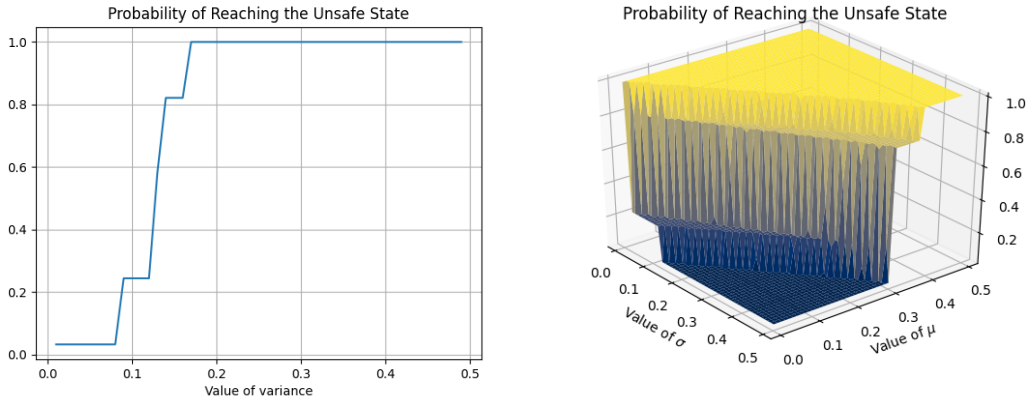

(a) **(b)**

Figure 7.6: Probability of reaching the unsafe state (load shedding) varying solely the variance (left) or the mean and variance (right) of the distribution.

Similarly, we repeat the test considering μ, the average value of the disconnection, as a parameter. The procedure returned an even more complex inequality expression, which provided the solution

$$\sigma \leq 0.81\mu - 0.08, \tag{7.28}$$

which is numerically confirmed by the depiction of Fig. 7.6b. Note that Fig. 7.6b presents very similar results the the ones presented in the previous chapter (compared to Fig. 6.3), witnessing the correctness of the parametric formal abstraction results.

We read Eq. (7.28) as follows. Increasing μ allows a higher σ: in practical terms, if the thresholds are far from f_0 (a high value of μ), the heterogeneity can grow (represented by the standard deviation σ), since the disconnection of PV systems is unlikely. On the other hand, if the thresholds are close to f_0 (low value of μ) we must decrease the heterogeneity to ensure a low-probability load shedding.

Unfortunately, the solution of the nonlinear inequality is computationally hard for a symbolical solver. In order to provide a symbolical solution, we passed the resulting inequality to a powerful commercial tool as Mathematica [209]. Our method is indeed limited by the power of the inequality (or equality) solver used. We may compute the probability of reaching the unsafe state over more than 2 steps, yet the computational time needed to solve the resulting inequality will grow indefinitely. Further, also the number of states complicates the resulting inequality: (potentially) every cell of the transition matrix contains a parametric item, hence the more states, the more cells, the more parametric terms within the inequality to solve.

Finally, the use of uniform distribution is a significant limitation of this approach. To overcome this first-order approximation while maintaining analytical tractability we might use a weighted summation (i.e. the arithmetic mean) of uniform distributions to mimic a Gaussian distribution. Such probability is known as *Bates* distribution [183]. Interestingly, considering the mean of just 3 uniform distributions, we have a probability profile very close to the Gaussian one, as Fig. 7.7 depicts. The mean squared error between the Gaussian distribution (in blue) and the mean of 3 uniform distributions (in red) is below $4 \cdot 10^{-3}$. This expedient may allow us to overcome the analytical constraints of the hardly-invertible Gaussian function and is currently matter of additional research.

7.5 Concluding Remarks

In this Chapter we have presented a method to perform parameter synthesis over parametric SDEs using formal abstraction techniques. Formal abstractions are employed to transform the parametric, continuous-space SDE into a parametric, discrete-space Markov chain; we then perform a parameter synthesis technique in

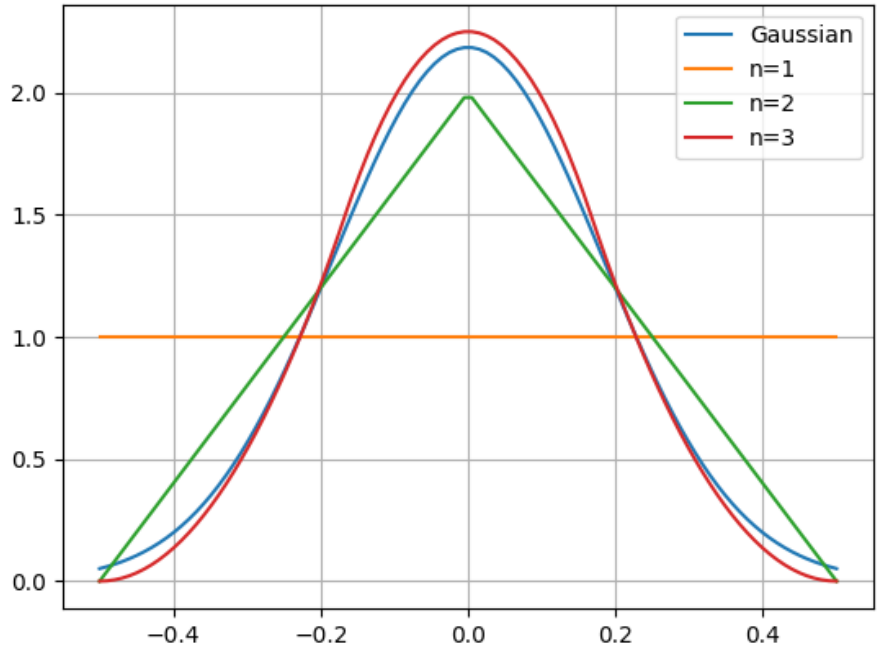

Figure 7.7: Gaussian distribution profile (blue) next to a single uniform distribution ($n = 1$, orange), and the mean of two ($n = 2$, green) and three ($n = 3$, red) uniform distributions.

order to verify a formal safety property. Such new parametric abstractions carry an error that certifies the verification technique, and that is used to obtain sets of parameters for the concrete SDE satisfying the given specification.

We have considered both finite- and infinite-support transition kernels, expressed either as elementary functions or not: the first instance offers an elegant dependence between the entries of the transition matrix and the parameters, which however is limited to a few kernels with special structure; more generally, the latter exploits a polynomial approximation of the kernel. Several experiments show the application of this methodology, from small parametric SDEs to a closed-loop model of the interconnection of the power grid with a population of heterogeneous PV systems.

8

Conclusions and Future Work

Contents

This chapter summarises the results of this dissertation, discusses the key contributions of the present work and outlines some directions of research that are currently pursued, or sought, by the author.

8.1 Conclusions

The rapid growth of renewable energy sources and the current sustainability and environmental outlook poses challenges to the power grid operators, customers and to the electric network's development. As we have shown, high penetration of renewable-based power sources impacts practically everything, from the very basis of modelling to the concept of stability and control design; from the emphasis on different time-scales to the need of new analytical tools with novel computational and practical insights. Finally, we highlight that the need for methods that provide certificates and guarantees on the reliability of the network's operations.

The crucial issue in the path towards the achievement of this goal remains the lack of suitable modelling methods and frameworks, which must be coupled with verification approaches capable of providing formal certifications. The electric network is also naturally exposed to a high degree of uncertainty while being affected by external, unpredictable dynamics – think of weather conditions – that inevitably increase the complexity of the modelling phase. The modelling stage ought to consider the new forms of renewable-based power generation. Not only we must analyse large, geographically confined and power-controlled wind or solar farms; we also study the geographically scattered, control-less populations of household PV systems. A single unit of such population will not affect the power grid, but the aggregation of a large number of them may bring deleterious effects to the grid's stability. When considering a large population of household solar panels, the modelling issue becomes evident: (i) they are geographically widespread, hence the weather may be different from one location to another; (ii) they have different ages, follow different regulations, and are built by different manufacturers, hence the performance both in terms of power output and dynamic behaviour may differ significantly; (iii) they have no form of power output control, neither centralised or decentralised.

In this work, we have presented a modelling framework for a large, heterogeneous population of solar panels. We have further translated the stochastic models into Markov models that enable us to verify properties concerning the stability of the power grid with formal methods. In Chapter 2 we present the two main areas involved in this thesis. We formally define some basic notions from the power system's field that have been useful for our approach, we show that the frequency might be used as a measure for the network's status and how the conventional control architectures handle frequency deviations. We then present the stochastic hybrid systems framework, introduce the formal abstractions approach that translates a stochastic system into a Markov model to be formally verified and provide guarantees on the safety and stability of the network.

Chapter 3 opens with the definition of the main object of study of this dissertation, the solar device and its noteworthy characteristics. We analyse particularly the rules for the connection to the grid: a solar panel can operate if and only if the network's frequency remains within a narrow interval of its nominal value. The width of this interval varies based upon countries' regulations – which may change over the years – and ultimately by the single device's measurements, generating a certain degree of heterogeneity in a large population of units. We thus outline reconnection and disconnection tests of real devices to understand the underlying heterogeneity and suggest a probabilistic description. Issues arise when a significant population of solar panels has a narrow operational interval, because this may lead to a chain disconnection of devices. The size of the population of solar panels with a narrow operational interval is a matter of concern, as highlighted by several ENTSO-E reports [41, 188, 190].

In Chapter 4 we propose a (time-varying) Markov model of the aggregated population. We describe the heterogeneity with probability distributions and represent the delayed reconnection to the grid with an augmented model. Using a time-varying coefficient, we are able to lump the population model and offer a succinct version of our framework. Finally, we test the reliability of our modelling choices against an explicit model of the heterogeneous population.

The solar device's behaviour depends on the network's frequency value and reciprocally the frequency is influenced by the population of panels' power output. Chapter 5 introduces the dynamics of the power grid, based on the swing equation, with an inertia coefficient that is inversely proportional to the amount of renewables connected to the network. The addition of renewable energy sources decreases the total inertia, causing the network to have wider and more rapid frequency oscillations. To counterbalance this instability, we imagine to control the power output of the panels' population with a proportional controller, in accordance to the traditional frequency control architecture. We simulate the scenario of a significant generation loss and value the possibility of load shedding, while varying the heterogeneity level of the population.

The stochastic nature of both the power grid (the intrinsic imbalance between generation and demand) and the output of solar panels (weather conditions, occlusions) coupled with the significant number of parameters (total network load, penetration of renewables, degree of heterogeneity) impairs the effectiveness of simulation-based case studies. The number of simulations to fully and thoroughly characterise every possible scenario is extremely high, hence the safety analysis of the network's dynamics becomes significantly time and resources consuming and may omit some corner cases.

To mitigate this impediment, Chapter 6 introduces formal methods. We characterise the stability of the power grid with a safety property: we require the network's frequency never to leave a given interval. Via a formal abstraction approach, we translate the closed-loop dynamics of the power grid and solar panels into a Markov model. As a result, we are able to provide a mathematical proof, a formal certificate of the reliability of the power grid, in terms of probability of load shedding. We formally test how the different kinds of heterogeneity and other parameters influence the frequency response. Operators can use these certificates to monitor the distribution of solar panels over the grid and to assess its reliability in case of incidents.

The formal abstraction approach provides mathematical guarantees of correctness, yet our modelling framework still suffers from the considerable amount of parameters. For each formal verification test, we need to evaluate every parameter, check the probability of load shedding, and repeat the operation with a different set of parameters. Chapter 7 bridges the gap by proposing a parameter synthesis framework. We first devise a parametric formal abstraction procedure, that translates a parametric SDE into a parametric Markov model. The SDE and the dtMC are equipped with a safety specification and the abstract version of it. The formal procedure comes with a parametric abstraction error that certifies the correctness of the approach. Interestingly, we also provide a mapping that links the SDE parameters to the dtMC parameters. From the parametric Markov model, we apply a canonical parameter synthesis technique to evaluate for which ranges

of parameters the dtMC satisfies the abstract specification. Finally, exploiting to the (inverse) mapping, we are able to reconstruct the parameters of the original SDE that satisfy the original specification.

We have considered both polynomial and non-polynomial kernels: the first instance offers an elegant dependence between the entries of the transition matrix and the parameters, which, however, is limited to a few kernels; more generally, the latter can be approximated with a polynomial expansion of the kernel. We have tested this novel method with the power grid model: fixing an upper bound on the probability of load shedding, we synthesise the degree of heterogeneity of a solar population that fulfils the requirement.

8.2 Recommendation for Future Development

From a general perspective, the power grid is undergoing a transition that requests new ways of thinking and new methodologies. A few of the main actions and new characteristics that are needed may be summarised as follows:

- new models, both for novel devices and to reshape the power grid, which must find a trade-off between level of details and description of crucial features, able to provide efficient analytical insight and computational power;
- revised stability criteria, with special attention to the new devices as renewable energy sources, new type of loads, the use of energy storage, and the corresponding time-scales;
- computational work aided by data-driven models and approaches, to share important data and apply methodologies from the big-data analysis community;
- new control architecture that reflect the much faster time-scales brought by power electronic devices and counteract the related frequency oscillations.

This work has offered models for a heterogeneous population of solar panels, but we have further provided a modelling framework that may be applied to all inverter-connected devices. Similarly to other scientific fields, formal methods bring

mathematical guarantees that are simply not achievable with simulation-based, numerical approaches. A few more research directions, closely related to the topics presented in this dissertation, are outlined in the following.

Modelling Framework Extensions on the modelling include estimations of solar power output, connecting with existing literature and tools [191]. This outlook modifies the considered time horizon: from the actual few seconds to hours or days. We aim to predict and monitor the network with several degrees of penetration of renewables into play, integrating uncertainties due to weather or occlusions. Estimations can be accurate and will enhance network reliability with the substantial presence of renewables. As an example, the Irish grid [121] reaches more than half power production from wind generation. They have developed a tool to precisely estimate the total eolic potential and they plan to extend and to exploit even more this power source.

Formal Abstractions In the current implementation of the formal abstraction method, we consider the closed-loop model where the solar panels do not have a controlled output. The extension of formal abstractions to controlled panels would introduce Markov decision processes instead of Markov chains, and would allow us to verify a given control architecture. Further, they allow the synthesis of a control policy that is not related to the canonical control theoretical approach (proportional, integral, derivative) while providing a formal certificate of correctness. On the other hand, the formal abstractions technique may be used to generate interval Markov chains (or interval Markov decision processes) [95] which embed the stochastic nature of a system via interval probability distributions. Both approaches however suffer from the well-known state-space explosion problem. The work in this thesis has utilised a sparse matrix representation for the probability transition matrix, which offers faster computational results, yet the extremely large state-space issue persists. The extension towards adaptive and sequential state-space partitioning may alleviate this impediment.

Parameter Synthesis Future work will further push the computational scalability of this methodology, focus on the relevant and general case of general distribution with non-elementary kernels, and develop results for advanced applications.

Software Tool The efforts in this dissertation may be assembled towards a software tool that analyses and formally verifies the electric grid behaviour. Currently, the only tool that performs formal abstractions of linear and non-linear models is FAUST^2. The computational requirement for its use are quite significant, especially when scaling up the dimensionality of the model or if we aim for a high-precision certificate. As an example, each test shown in Chapter 6 takes approximately 20 minutes, on an average machine[1], to be completed. Whilst the single test time may be reasonable, a test is to be intended as a single run of the algorithm, once parameters as solar penetration, global demand, disconnection thresholds, their mean and variance are set: each entry of Table 6.2 summarises 13 tests (each with a different value of variance). The table alone has 32 entries, making a total of 416 test taking up roughly 140 hours. Each entry of Tables 6.3 and 6.4 summarises a similar quantity of tests (20 and 8, respectively) taking approximately 220 hours and 85 hours, respectively. Rather clearly, this is the main rock to be climbed. Further, FAUST^2 could not be used in a *plug-and-play* fashion, as the transition matrix of the model changes during the experiment (see discussion about the two different transition matrices in Section 6.2.1). This had a significant impact in the overall procedure time consumption.

As such, a development of the FAUST^2 architecture is needed in order to simplify the computations with time-varying models. A possible alternative, started with the work in Chapter 7, is writing the code in Python to accelerate the computation times and do not lose its user-friendly features. Further, approximations of the Gaussian kernel (as sum of Uniform distributions) may lead the way towards a much faster verification procedure. Each entry of the transition matrix would then become a polynomial function rather than the integral of a squared exponential

[1]CPU Intel i5 and 8 GB RAM, running Ubuntu 18.04.

function. Naturally, delegating a powerful server to hold all the computations helps: a part of the tests has indeed been moved to the research group's server and the computational time shrunk by a factor of 4.

On the other hand, dimensionality remains an issue even in a powerful machine. A key feature for the feasibility of this approach still resides in a low number of variables: the number of states (10^8) of the Markov chain tests in Chapter 6 hits the ceiling of the known capabilities of formal verification engines. The computation with such number of states is possible solely thanks to the sparse nature of the transition matrix: only approximately 10^5 transition are non-zero. Thus, simple underlying models are crucial to use formal analysis. A breakthrough may come from newer formal verification approaches that make use of interval Markov models, as [95]: it computes transitions in a different fashion and needs fewer states to provide the same kind of certification.

Appendices

A

Proofs Related to the Model Formulation

Contents

A.1 Observability and Observer Design

Whilst it is in principle possible to have real-time information about every photovoltaic system in the network, this would be highly impractical. We thus consider a more realistic scenario, where the only accessible (i.e., observable) output is the frequency of the network. Note that if the value of $f(k)$ is known, also $a(k)$ and $b(k)$ can be computed, as we assume to know the population distributions.

Observability of the Models

Let us show that both the 3-state Markov Chain model and the n-state one (cf. Eqs. (5.9), (5.10)) are observable when the only output signal is $f(k)$. This means, in practical terms, that we can derive the value of the other variables of the system, reconstructing them solely from the value of $f(k)$.

Let us focus on a simplified version of the three-state Markov model, namely

$$\begin{cases} f(k+1) = \alpha_1 f(k) + \beta_1 x(k), \\ x(k+1) = (1 - a(k))x(k) + b(k)\varepsilon(k)y(k), \\ y(k+1) = b(k)(1 - x(k) - \varepsilon(k)y(k)). \end{cases}$$

This version, featuring first-order frequency dynamics, is simpler to analyse whilst retaining the significant features of the model in Eq. (5.10). Observability analysis with the second-order frequency dynamics can be attained following a similar procedure. Limiting the access to the frequency, the output matrix is the constant quantity

$$C(k) \equiv C = [1\ 0\ 0],$$

and let us define $A(k)$, the Jacobian matrix of the linearised system at time k

$$A(k) = \begin{bmatrix} \alpha_1 & \beta_1 & 0 \\ 0 & \bar{a}(k) & b(k)\varepsilon(k) \\ 0 & -b(k) & -b(k)\varepsilon(k) \end{bmatrix},$$

where $\bar{a}(k) = 1 - a(k)$. The observability matrix results in

$$\begin{aligned} \mathcal{O}(k) &= \begin{bmatrix} C \\ CA(k+1) \\ CA(k+1)A(k+2) \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & 0 \\ \alpha_1 & \beta_1 & 0 \\ \alpha_1^2 & \alpha_1\beta_1 + \beta_1\bar{a}(k+2) & \beta_1 b(k+2)\varepsilon(k+2) \end{bmatrix}. \end{aligned}$$

The rank is full as long as $b(k)\varepsilon(k) \neq 0$, namely when both $b(k) \neq 0$ and $\varepsilon(k) \neq 0$. Let us now analyse what the conditions $b(k) = 0$ and $\varepsilon(k) = 0$ elicit.

1) $\varepsilon(k) = 0$ – The system reduces to

$$\begin{cases} f(k+1) = \alpha_1 f(k) + \beta_1 x(k), \\ x(k+1) = \bar{a}(k)x(k), \\ y(k+1) = b(k)(1 - x(k)). \end{cases}$$

In this case $y(k)$ is clearly not observable, as $x(k+1)$ loses its dependency on $y(k)$ and at the same time $y(k)$ has no effect on the output. However, if we restrict

our attention to the reduced system composed of the first two equations, this is observable: the observability matrix becomes

$$\mathcal{O}(k) = \begin{bmatrix} 1 & 0 \\ \alpha_1 & \beta_1 \end{bmatrix},$$

which is always full rank. Note that $y(k)$ depends only on $x(k)$: in the next section we will be able to compute its estimated value $\hat{y}(k)$ from $\hat{x}(k)$, the estimated value of $x(k)$.

2) $b(k) = 0$ – Similarly, the system reduces to

$$\begin{cases} f(k+1) = \alpha_1 f(k) + \beta_1 x(k), \\ x(k+1) = \bar{a}(k)x(k), \\ y(k+1) = 0, \end{cases}$$

and analogous conclusions can be drawn.

Regarding the observability of the $(n+2)$-state model, the analysis may be carried out similarly. Equation (5.9) suggests a chain-dependency of $w_i(k)$ from $w_{i-1}(k-1)$, which keeps the rank of the observability matrix full. This leads to a chain of substitutions that guarantees the observability in n steps, under the condition $b(k)\varepsilon(k) \neq 0$. If $b(k)\ \varepsilon(k) = 0$, then we can repeat a similar argument. Finally, the analysis is easily generalisable to the second-order frequency dynamics.

Observer Design

An observer can be built for both models in Eqs (5.9) and (5.10). We start with the observer for the smaller model. Let us define $\hat{x}(k)$ and $\hat{y}(k)$ as the estimated values of $x(k)$ and of $y(k)$, respectively. Assume the unique output is the frequency of the network, namely $\hat{f}(k) = f(k)$. Our aim is the computation of $\hat{x}(k)$ and $\hat{y}(k)$. As above, considering the three-state model, we assume $\varepsilon(k)$ known as we carry on the analysis. Later, for the $(n+2)$-state model estimator computation, we build $\hat{\varepsilon}(k)$, the estimation of $\varepsilon(k)$.

Let us now focus on the three-state model and show how to build $\hat{x}(k)$ and $\hat{y}(k)$. After some algebra, from Eq. (5.10) we obtain:

$$\begin{bmatrix} \hat{f}(k-2) \\ \hat{x}(k-2) \\ \hat{y}(k-2) \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1/\beta_1 & -\alpha_1/\beta_1 \\ l_1 & l_2 & l_3 \end{bmatrix} \cdot \begin{bmatrix} f(k) \\ f(k-1) \\ f(k-2) \end{bmatrix},$$

where $[l_1,\ l_2,\ l_3]$ is equal to

$$\left[\frac{1}{\beta_1 b(k-2)\varepsilon(k-2)},\ -\frac{\bar{a}(k-2)+\alpha_1}{\beta_1},\ -\frac{\alpha_1\bar{a}(k-2)}{\beta_1}\right].$$

Our estimation is two steps behind the current time instant. We then set up a two-steps predictor following the dynamical equations of the system, after some algebra, as

$$\begin{bmatrix}\hat{f}(k)\\ \hat{x}(k)\\ \hat{y}(k)\end{bmatrix}=\begin{bmatrix}1 & 0 & 0\\ obs_{2,1} & obs_{2,2} & obs_{2,3}\\ obs_{3,1} & obs_{3,2} & obs_{3,3}\end{bmatrix}\begin{bmatrix}f(k)\\ f(k-1)\\ f(k-2)\end{bmatrix}+\begin{bmatrix}0\\ 0\\ b(k-1)\end{bmatrix},$$

where

$$obs_{2,1}=\frac{\bar{a}(k-1)}{\beta_1},\qquad obs_{2,2}=-\frac{\bar{a}(k-1)\alpha_1}{\beta_1},\quad obs_{2,3}=0,$$
$$obs_{3,1}=-\frac{1+c_1}{\beta_1}-c_1,\quad obs_{3,2}=-\frac{\alpha_1}{\beta_1}(1+2c_1),\quad obs_{3,3}=0,$$

and where

$$c_1=\frac{\varepsilon(k)\bar{a}(k-1)}{b(k-1)\varepsilon(k-1)}.$$

Regarding the $(n+2)$-state model, a similar situation is expected. In fact, for the dynamical equation we need an $(n+2)$-step predictor. The value of $\hat{x}(k)$ can be attained using two temporal values of $f(\cdot)$, as it is easy to note from Eq. (5.9). The key part is how to connect the estimation of $x(k)$ to $w_1(k)$. We note that, again from Eq. (5.9)

$$x(k+1)+\sum_i w_i(k+1)=\bar{a}(k)x(k)+b(k)(1-x(k)),$$

so that

$$\sum_i w_i(k)=\bar{a}(k-1)x(k-1)+b(k-1)(1-x(k-1))-x(k).$$

We can substitute this into the estimator of $w_1(k)$, giving

$$\hat{w}_1(k)=b(k-1)[1-b(k-1)(1-\hat{x}(k-1))-\bar{a}(k-1)\hat{x}(k-1)].$$

The observer and predictor can be built as follows:

$$
\begin{aligned}
&\hat{f}(k) = f(k),\\
&\hat{x}(k-i) = \frac{f(k-i+1) - \alpha_1 f(k-i)}{\beta_1},\ i = 1, 2,\\
&\hat{x}(k)|\hat{x}(k-1) = \bar{a}(k-1)\hat{x}(k-1) + b(k-1)\sum_i \tau_i \hat{w}_i(k-1),\\
&\hat{w}_1(k) = b(k)[1 - b(k-2)(1 - \hat{x}(k-2)) - \bar{a}(k-2)\hat{x}(k-2)],\\
&\hat{w}_i(k) = b(k-1)(1-\tau_{i-1})\hat{w}_{i-1}(k-1), \qquad i = 2, \dots n-1,\\
&\hat{w}_n(k) = b(k-1)[(1-\tau_{n-1})\hat{w}_{n-1}(k-1)\\
&\qquad\qquad + (1-\tau_n)\hat{w}_n(k-1)].
\end{aligned}
$$

With this technique we prove that we are able to build an observer for the system, estimate the $w_i(k)$ and compute the $\hat{\varepsilon}(k)$ value at each time step as

$$\hat{\varepsilon}(k) = \frac{\sum_i \tau_i \hat{w}_i(k)}{\sum_i \hat{w}_i(k)}.$$

Note that the observer has a transient of $(n+2)$ steps, which is necessary to initially compute $\hat{w}_i$, $\forall i$. This value is then used in the estimation and prediction of the three-state model.

A.2 Stability Analysis of the Closed-Loop Model

We illustrate the stability analysis for the closed-loop models in Eqs. (5.9) and (5.10). In order to ease the notation and to make the discussion clearer, we utilise a reduced first-order transfer function, as done already in Appendix A.1. The analysis with the second-order transfer function introduced in the paper can be carried out similarly.

We start with the closed-loop three-state model: the analysis of the $(n+2)$-state model is analogous to what we present in the following. Whilst for simplicity we manipulate quantities $a(k)$ and $b(k)$, recall that they are functions of the frequency $f(k)$, namely we have $a(f(k))$ and $b(f(k))$. More specifically, they are cumulative distribution functions of the PV systems thresholds, where $f(k)$ enters as an extremum of the integral.

Let us re-write the second equality of Eq. (5.10): we need the dependency on $\Delta x(k)$ to be explicit. Recall from Chapter 5 that $\Delta x(k) = x(k) - x_0$, where x_0 is the portion of PV systems in state ON when $f(\cdot) = f_0$. In our setting, we assume $x_0 = 1$ (all PV systems ON), i.e. all PV systems are active when the frequency is at its nominal value. Notice that in this case the equilibrium point of this system is the origin, however if $x(0) \neq 0$, the following analysis must be modified to accommodate for this.

Substituting $x(k) = \Delta x(k) + x(0)$, we obtain an alternative formulation for the three-state model as

$$\begin{cases} \Delta f(k+1) = \alpha_1 \Delta f(k) + \beta_1 \Delta x(k), \\ \Delta x(k+1) = \bar{a}(k)\Delta x(k) + b(k)\varepsilon(k)y(k) + a(k)x(0), \\ y(k+1) = -b\Delta x(k) - b(k)\varepsilon(k)y(k) + b(k)(1 - x(0)). \end{cases}$$

We compute the Jacobian as in the previous Appendix and evaluate it at the equilibrium point: this holds an eigenvalue in 1: the eigenvalue analysis is not sufficient in this case (the stability depends on the non-linear components of the vector field), and analysis via Lyapunov function is thus necessary to assess the stability of the equilibrium.

Recall that at $f(\cdot) = f_0$ the functions $\bar{a}(k) = 1$, $b(k) = 1$, $a(k) = 0$, $x(0) = 1$. Therefore we have $\varepsilon(k) = 0$, $y(k) = 0$, $\Delta x(k) = 0$, $\Delta f(k) = 0$. Let us define the vector $\xi(k) = [\Delta f(k), \Delta x(k), y(k)]$. Notice that, thanks to the stability of the first-order transfer function, we ensure the stability of $f(k)$ when $\Delta x(k) \to 0$. We then select the following Lyapunov function:

$$V(\xi(k)) = (\Delta x(k) + y(k))^2,$$

which is positive everywhere but at the origin. In order to assess the stability of the interconnection, we compute

$$\begin{aligned} V(\xi(k+1)) - V(\xi(k)) = {} & [(\bar{a}(k) - b(k))^2 - 1]\Delta x(k)^2 \\ & + 2[a(k)(\bar{a}(k) - b(k)) - y(k)]\Delta x(k) + a(k)^2 - y(k)^2. \end{aligned}$$

From the definitions of $a(\cdot)$ and $b(\cdot)$, we have

$$a(f_0) = 0, \quad \bar{a}(f_0) = 1, \quad b(f_0) = 1. \tag{A.1}$$

Substituting these values we obtain:

$$V(\xi(k+1)) - V(\xi(k)) = -(\Delta x(k) + y(k))^2 \leq 0.$$

A similar reasoning can be followed in presence of estimated values. Let us define $\hat{f}$, $\Delta\hat{x}$ and $\hat{w}_i$ as the estimated frequency, power deviation and i-th waiting state, respectively. Recall that we assume to measure the frequency, hence $\hat{f} = f$ and, as such, the frequency-dependent quantities $a(k)$ and $b(k)$ are known. Let us define the Lyapunov function:

$$V(\hat{\xi}(k)) = \left(\Delta\hat{x}(k) + \sum_{i=1}^{n} \hat{w}_i(k)\right)^2.$$

Let us recall the time evolution of the two quantities

$$\Delta\hat{x}(k+1) = \bar{a}(k)\Delta\hat{x}(k) + b(k)\sum_{i=1}^{n} \tau_i\hat{w}_i(k) + \bar{a}(k)x(0)$$
$$\sum_i \hat{w}_i(k+1) = b(k)\left[1 - \Delta\hat{x}(k) - x(0) - \sum_{i=1}^{n} \tau_i\hat{w}_i(k)\right]$$

We similarly compute

$$\begin{aligned} V(\xi(k+1)) - V(\xi(k)) &= [(\bar{a}(k) - b(k))^2 - 1]\Delta x(k)^2 \\ &+ 2(\bar{a}(k) - b(k))^2\Delta x(k)[(\bar{a}(k) - b(k) - 1)x(0) + b(k)] - \left(\sum_i \hat{w}_i(k)\right)^2 \\ &- 2\Delta x(k) \cdot \sum_i \hat{w}_i(k) + b(k)(1 - x(0))^2. \end{aligned}$$

We ensure the stability of a neighbourhood around f_0, namely I_{f_0}, where we may apply the conditions in (A.1) to get

$$V(\hat{\xi}(k+1)) - V(\hat{\xi}(k)) = -\left(\Delta\hat{x}(k) + \sum_{i=1}^{n} \hat{w}_i(k)\right)^2 \leq 0.$$

The negativity of the Lyapunov function is guaranteed as long as the conditions in (A.1) hold. We argue that these conditions also hold in a non-trivial neighbourhood of f_0 (call it I_{f_0}): in practice, PV systems disconnect and reconnect at values $f_0 \pm \delta$,

$\delta > 0$, which means that the network is stable as long as no PV system disconnects (given $x_0 = 1$). As soon as $f(\cdot)$ exits I_{f_0}, the stability is not guaranteed. The transition from stability to instability depends on the threshold distributions for the population of PV systems, and in particular on their averages and variances. The stability of the interconnection depends heavily on the interval I_{f_0}: the larger is the interval, the more resilient is the system to oscillations. Let us focus on $a(k)$ for simplicity. Assume also that $a(k)$ is the integral of truncated normal distribution, as

$$a(k) = \begin{cases} \int_{-\infty}^{f(k)} p_{t,o}^{d}(u)du & \text{if } \ f(k) > f_0, \\ \int_{f(k)}^{+\infty} p_{t,u}^{d}(u)du & \text{otherwise,} \end{cases}$$

where $p_{t,o}^{(\cdot)}$ and $p_{t,u}^{(\cdot)}$ are the truncated versions of the probability density functions considered above in over-frequency and under-frequency, respectively. Let us focus on the condition $f(k) > f_0$ (the $f(k) \leq f_0$ side is handled symmetrically). Define $\lambda_o = [\lambda_1, \lambda_2]$ to be the support of $p_{t,o}^{d}(\cdot)$, and suppose that λ_0 supports a symmetric truncated normal $p_{t,o}^{d}$ (i.e. its average value is $\mu = (\lambda_1 + \lambda_2)/2$) for any choice of λ_1 and of λ_2. By definition of the truncated normal distribution $p_{t,o}^{d}$, the quantity $a(k)$ becomes

$$a(k) = \begin{cases} 0 & \text{if } f_0 \leq f(k) < \lambda_1, \\ \int_{\lambda_1}^{f(k)} p_{t,o}^{d}(u)du & \text{if } \lambda_1 \leq f(k) \leq \lambda_2, \\ 1 & \text{otherwise.} \end{cases}$$

In the case under consideration, the upper limit of I_{f_0} is λ_1. As such, the length of I_{f_0} increases as the support of $p_{t,o}^{d}$ is far from f_0, namely as λ_1 increases. With an increase of I_{f_0}, the stability of the interconnection is preserved for a larger set of network frequency values.

On the other hand, also λ_2 plays an important role in the stability analysis. Increasing λ_2 enlarges the support of $p_{t,o}^{(\cdot)}$. This reduces the rate of increase of $a(k)$ when $f(k)$ increases: the distribution $p_{t,o}^{(\cdot)}$, over a larger support, has lower values. Note that the sign of $V(\xi(k-1)) - V(\xi(k))$ outside I_{f_0} depends heavily on the slope of $a(k)$.

Let us define F, a value of frequency so that $|F - f_0| >> 0$. When $f(\cdot) = F$, then $a(F) = 1$, $\bar{a}(F) = 0$, $b(F) = 0$. Let us analyse the Lyapunov function under these conditions:

$$V(\xi(k+1)) - V(\xi(k)) = -(\Delta x(k) + y(k))^2 + a(F)^2,$$

which is positive as long as $(\Delta x(k) + y(k))^2 < a(F)^2$. Note that $\Delta x(\cdot), y(\cdot) \in [0, 1]$, and are such that $\Delta x(k) + y(k) \leq 1 \; \forall k$, by definition. The slope of $a(\cdot)$ here plays a major role: if the growth rate of $a(k)$ is faster than the one of $(\Delta x(k) + y(k))$, the system becomes unstable.

Let us finally study the system for values of $f(k)$ so that $a(f),\ \bar{a}(f),\ b(f) \in (0, 1)$. The Lyapunov function becomes

$$V(\xi(k+1)) - V(\xi(k)) = -(\Delta x(k) + y(k))^2 + (a(k) + d(k)\Delta x(k))^2,$$

where $d(k) = \bar{a}(k) - b(k)$. The analysis is similar to the one above: the Lyapunov function can become positive if $a(\cdot)$ and $d(\cdot)$ grow faster than $\Delta x(k)$ and $y(k)$.

In summary, under all the three conditions analysed, the increase of $a(\cdot)$ and $b(\cdot)$ is dictated by the variance of the disconnection and reconnection distributions, respectively. A small value of the variance renders the system less robust and more prone to instability. This situation occurrs with a population having individual $\mathcal{I}_f$ with small variations. A higher value of variance corresponds to a higher degree of population heterogeneity, which we have shown to increase the stability region of the Lyapunov function. Heterogeneity, in this sense, can be exploited to enhance the network reliability and resilience against oscillations and incidents. On the other hand, as the analysis on the first condition suggests, instability can be mitigated by moving the average value away from f_0.

This stability analysis can be generalised to the second-order frequency dynamics using the same Lyapunov function. Furthermore, considering the $(n+2)$-state model, the Lyapunov function

$$V(\xi(k)) = \left(\Delta x(k) + \sum_{i=1}^{n} w_i(k)\right)^2$$

leads to a similar analysis.

B

Proof Related to Formal Abstractions

B.1 Definition of Kernel Continuous Regions

We want to underline the discontinuity of the kernel density $t_\omega(\cdot \,|\, q)$ caused by the presence of the $\delta(\cdot)$ functions. Let us define

$$g(\Delta f) = -\alpha_1 \Delta f - \alpha_2 \Delta\phi - \beta_1 \Delta x - \beta_2 \Delta\xi, \quad l(P_{PV}) = -\bar{P} N x,$$

and

$$h_1(x) = -(1-a)x - b\varepsilon y, \quad h_2(y) = -b(1 - x - \varepsilon y).$$

The transition kernel density can be written as

$$t_\omega(q' \,|\, q) = \begin{cases} t_f(\Delta f' - g(\Delta f)) \cdot t_P(P'_{PV} - l(P_{PV})) & \text{if } \Delta\phi' = \Delta f \wedge x' = h_1(x) \wedge \\ & \qquad \wedge\, y' = h_2(y) \wedge \xi' = P_{PV}, \\ 0 & \text{otherwise,} \end{cases}$$

defining the continuous regions $\mathcal{C} = \{\Delta\phi' = \Delta f \wedge x = h_1(x) \wedge y' = h_2(y) \wedge \xi' = P_{PV}\}$. Note that in the abstraction framework, regions $\mathcal{C}$ assume the discretised form $\mathcal{C}_d = \{\Delta\bar{\phi}' = \Delta\bar{f}_i \wedge \bar{x} = h_1(\bar{x}) \wedge \bar{y}' = h_2(\bar{y}) \wedge \bar{\xi}' = \bar{P}_j\}$.

B.2 Probabilistic Safety for Partially Degenerate Models

Let us show that for a partially degenerate stochastic model the safety probability computation depends only on the stochastic state. Consider the model

$$\begin{cases} x(k+1) = f(z(k)) + \omega(k), \\ y(k+1) = x(k), \end{cases}$$

where $\omega(k) \sim \mathcal{N}(0, \sigma^2)$ and where $z = (x,\ y)^T$ denotes the complete state vector. Let us denote with $t_\omega(\cdot)$ the density of the Gaussian kernel. The one-step transition probability kernel can be split as follows:

$$P(x(k+1) \,|\, z(k)) = t_\omega(f(z(k))),$$
$$P(y(k+1) \,|\, z(k)) = P(y(k+1)|x(k)) = \delta(y(k+1) - x(k)),$$

where $\delta(z-p)$ represents the Dirac delta function of variable z, centred at point p. Let us consider a safe set $A = A_x \times A_y$, where A_x and A_y denote its projections on variables x and y, respectively. Define the value function at time step H as $V_H(z) = \mathbf{1}_A(z)$ and compute the one-step backward recursion:

$$\begin{aligned} V_{H-1}(z) &= \int_A V_H(z')P(z'|z)dz' = \int_A P(z'|z)dz' \\ &= \int_{A_y}\int_{A_x} t_\omega(dx'|f(z))\delta(dy'-x) = \mathbf{1}_{A_y}(z)\int_{A_x} t_\omega(dx'|f(z)) \\ &= \int_{A_x} t_\omega(dx'|f(z)), \end{aligned}$$

showing that the computation of the safety probability depends solely on the stochastic kernel affecting the dynamics of variable x.

B.3 Value Function Continuity for Probabilistic Safety

In the following, we consider a generation-loss incident scenario; the load-loss case can be derived analogously. Recall the value function definition from Eq. (6.24) and compute the backward Bellman equation as

$$V_k(q) = \mathbf{1}_{\mathcal{L}}(q)\int_{\mathcal{Q}} V_{k+1}(\tilde{q})t_s(\tilde{q} \,|\, q)\, d\tilde{q}, \quad \text{with } V_H(q) = \mathbf{1}_{\mathcal{L}}(q).$$

We show that the value functions are continuous within the continuity regions of the state space, thus there must exist a constant γ so that

$$|V_k(q) - V_k(\tilde{q})| \leq \gamma \, \|q - \tilde{q}\|. \tag{B.1}$$

To enhance the readability, let us define $g(\Delta f) = -\alpha_1 \Delta f - \alpha_2 \Delta\phi - \beta_1 \Delta x - \beta_2 \Delta\xi$, $h(P_{PV}) = -\bar{P} N x$ and $\Delta f = \rho$, $P_{PV} = \psi$. We now show the validity of Eq. (B.1) by finding a value for γ. From the definition of $V_k(q)$, we obtain:

$$\begin{aligned}
&\Big| \int_{\mathcal{Q}} V_{k+1}(q) t_f(\underline{\rho} - g(\rho)) t_P(\underline{\psi} - h(\psi)) d(\underline{\rho}) d(\underline{\psi}) \\
&\qquad - \int_{\mathcal{Q}} V_{k+1}(\tilde{q}) t_f(\underline{\rho} - g(\tilde{\rho})) t_P(\underline{\psi} - h(\tilde{\psi})) d(\underline{\rho}) d(\underline{\psi}) \Big| \\
&\leq \Big| \int_{\mathcal{F}} V_{k+1}(q) t_f(\underline{\rho} - g(\rho)) d(\underline{\rho}) \cdot \int_{\mathcal{P}} V_{k+1}(q) t_P(\underline{\psi} - h(\psi)) d(\underline{\psi}) \\
&\qquad - \int_{\mathcal{F}} V_{k+1}(\tilde{q}) t_f(\underline{\rho} - g(\tilde{\rho})) d(\underline{\rho}) \cdot \int_{\mathcal{P}} V_{k+1}(\tilde{q}) t_P(\underline{\psi} - h(\tilde{\psi})) d(\underline{\psi}) \Big|,
\end{aligned}$$

where $\mathcal{F}$ and $\mathcal{P}$ denote the domain of frequency and power, respectively. Let us now introduce a useful lemma.

Lemma 1 *Assume $A, B, C, D \in [0, 1]$, then $|AB - CD| \leq |A - C| + |B - D|$.*

Proof 1 *Assume $A > C$, then*
if $AB - CD > 0$,
$|AB - CD| \leq |CB - CD| = C|B - D| \leq |B - D| \leq |B - D| + |A - C|$.
if $AB - CD < 0$,
$|AB - CD| \leq |AB - AD| = A|B - D| \leq |B - D| \leq |B - D| + |A - C|$.
Analogously for $A \leq C$. □

Thanks to this Lemma, we can write

$$\begin{aligned}
&\Big| \int_{\mathcal{F}} V_{k+1}(q) t_f(\underline{\rho} - g(\rho)) d(\underline{\rho}) \cdot \int_{\mathcal{P}} V_{k+1}(q) t_P(\underline{\psi} - h(\psi)) d(\underline{\psi}) \\
&\qquad - \int_{\mathcal{F}} V_{k+1}(\tilde{q}) t_f(\underline{\rho} - g(\tilde{\rho})) d(\underline{\rho}) \cdot \int_{\mathcal{P}} V_{k+1}(\tilde{q}) t_P(\underline{\psi} - h(\tilde{\psi})) d(\underline{\psi}) \Big| \\
&\qquad \leq \Big| \int_{\mathcal{F}} t_f(\underline{\rho} - g(\rho)) d(\underline{\rho}) \cdot \int_{\mathcal{P}} t_P(\underline{\psi} - h(\psi)) d(\underline{\psi}) \\
&\qquad\qquad - \int_{\mathcal{F}} t_f(\underline{\rho} - g(\tilde{\rho})) d(\underline{\rho}) \cdot \int_{\mathcal{P}} t_P(\underline{\psi} - h(\tilde{\psi})) d(\underline{\psi}) \Big| \\
&\qquad\qquad \leq \int_{\mathcal{F}} \Big| t_f(\underline{\rho} - g(\rho)) - t_f(\underline{\rho} - g(\tilde{\rho})) \Big| \, d(\underline{\rho}) \\
&\qquad\qquad\qquad + \int_{\mathcal{P}} \Big| t_P(\underline{\psi} - h(\psi)) - t_P(\underline{\psi} - h(\tilde{\psi})) \Big| d(\underline{\psi}).
\end{aligned}$$

Let us focus on the first integral:

$$\int_{\mathcal{F}} |t_f(\underline{\rho} - g(\rho)) - t_f(\underline{\rho} - g(\tilde{\rho}))| d(\underline{\rho}) = \frac{1}{\sigma_f} \int_{\mathcal{F}} \left| \Phi\left(\frac{\underline{\rho} - g(\rho)}{\sigma_f}\right) - \Phi\left(\frac{\underline{\rho} - g(\tilde{\rho})}{\sigma_f}\right) \right| d\underline{\rho}$$
$$= \int_{\mathcal{F}} \left| \Phi\left(u - \frac{\alpha_1(\rho - \tilde{\rho})}{2\sigma_f}\right) - \Phi\left(u + \frac{\alpha_1(\rho - \tilde{\rho})}{2\sigma_f}\right) \right| d\underline{\rho} \leq \frac{2\alpha_1}{\sqrt{2\pi}\sigma_f} |\rho - \tilde{\rho}|,$$

and similarly for the second integral. Therefore,

$$|V_k(q) - V_k(\tilde{q})| \leq \frac{2\alpha_1}{\sqrt{2\pi}\sigma_f} |\rho - \tilde{\rho}| + \frac{2a_{max}}{\sqrt{2\pi}\sigma_P} |\psi - \tilde{\psi}|.$$

References

[1] *The Paris Agreement*. https://unfccc.int/process-and-meetings/the-paris-agreement/the-paris-agreement. Accessed: 09/12/2018.

[2] ENTSO-E. *Statistical Factsheet*. https://www.entsoe.eu/publications/statistics-and-data/. 2018.

[3] H. Wirth. *Recent Facts about Photovoltaics in Germany*. https://www.ise.fraunhofer.de/en/publications/studies/recent-facts-about-pv-in-germany. 2018.

[4] Solar Power Europe. *Global Market Outlook for Solar Power 2016-2020*. https://www.solarpowereurope.org/global-market-outlook-2018-2022/. 2016.

[5] ENTSO-E. *Power Facts – Europe 2019*. https://www.entsoe.eu/data/. 2019.

[6] Federico Milano et al. "Foundations and Challenges of Low-inertia Systems". In: *2018 Power Systems Computation Conference (PSCC)*. IEEE. 2018, pp. 1–25.

[7] Annual Market Update. *Global Wind Report*. Tech. rep. 2019.

[8] Arnulf Jager-Waldau. "Snapshot of Photovoltaics - February 2020". In: *Energies* 13 (Feb. 2020), p. 930.

[9] International Energy Agency. *Electricity Information: Overview*. Tech. rep. 2020.

[10] Göran Andersson. "Power System Analysis". In: *ETH Zurich* (2011).

[11] Yashar Ghiassi-Farrokhfal et al. "Optimal Design of Solar PV Farms with Storage". In: *IEEE Transactions on Sustainable Energy* 6.4 (Oct. 2015), pp. 1586–1593.

[12] B Shiva Kumar and K Sudhakar. "Performance Evaluation of 10 MW Grid Connected Solar Photovoltaic Power Plant in India". In: *Energy reports* 1 (2015), pp. 184–192.

[13] Chagitha Ranhotigamage and Subhas Chandra Mukhopadhyay. "Field Trials and Performance Monitoring of Distributed Solar Panels Using a Low-cost Wireless Sensors Network for Domestic Applications". In: *IEEE Sensors Journal* 11.10 (2011), pp. 2583–2590.

[14] A. Peruffo et al. "Aggregated Markov Models of a Heterogeneous Population of Photovoltaic Panels". In: *International Conference on Quantitative Evaluation of Systems* (2017), pp. 72–87.

[15] A. Peruffo et al. "Synchronous Frequency Grid Dynamics in the Presence of a Large-Scale Population of Photovoltaic Panels". In: *2018 Power Systems Computation Conference (PSCC)*. June 2018, pp. 1–7.

[16] A. Peruffo et al. “Impact of Solar Panels and Cooling Devices on Frequency Control after a Generation Loss Incident”. In: *Decision and Control (CDC) 2018 IEEE 57th Annual Conference Proceedings*. IEEE. 2018.

[17] A. Peruffo et al. “Aggregation and Control of a Heterogeneous Population of Solar Panels Over the Grid Frequency”. In: *IEEE Transactions on Control Systems Technology* (2020), pp. 1–17.

[18] Andrea Peruffo et al. “Safety Guarantees for the Electricity Grid with Significant Renewables Generation”. In: *International Conference on Quantitative Evaluation of Systems*. Springer. 2019, pp. 332–349.

[19] Andrea Peruffo et al. “Model-based Formal Reliability Analysis of Grid Dynamics with Solar Energy Sources”. In: 15th European Workshop on Advanced Control and Diagnosis. Springer, 2019.

[20] Alessandro Abate Andrea Peruffo. “Formal Abstraction and Synthesis of Parametric Stochastic Processes”. In: *under revision* (2020).

[21] International Energy Agency. *Annual Energy Storage Deployment by Country, 2013-2019*. `https://www.iea.org/data-and-statistics/charts/annual-energy-storage-deployment-by-country-2013-2019`. Accessed: 2020-06-20.

[22] Swiss Grid. *Grid Stability*. `https://www.swissgrid.ch/en/home/operation/regulation/grid-stability.html`. Accessed: 2020-07-06.

[23] ENTSO-E. *Continuing frequency deviation in the Continental European Power System originating in Serbia/Kosovo: Political solution urgently needed in addition to technical*. `https://www.entsoe.eu/news/2018/03/06/press-release-continuing-frequency-deviation-in-the-continental-european-power-system-originating-in-serbia-kosovo-political-solution-urgently-needed-in-addition-to-technical`. Accessed: 2020-06-20.

[24] Stan Mark Kaplan. *Smart Grid: Modernizing Electric Power Transmission and Distribution*. The Capitol Net Inc, 2009.

[25] *Energy industry in Germany - on the trail of electricity*. `https://www.divia.de/blog/2013/10/24/energiewirtschaft-in-deutschland-auf-den-spuren-des-stroms`. Accessed on 10/12/2018.

[26] *ENTSO-E Website*. `https://www.entsoe.eu/`, note = Accessed: 2020-06-20.

[27] ENTSO-E. *Policy 1: Load-frequency Control and Performance*. `https://www.entsoe.eu/publications/statistics-and-data/`. 2009.

[28] A. Ulbig, T. S. Borsche, and G. Andersson. “Impact of Low Rotational Inertia on Power System Stability and Operation”. In: *World Congress Conference* (Aug. 2014).

[29] Wen Tan. “Unified Tuning of PID Load Frequency Controller for Power Systems via IMC”. In: *IEEE Transactions on power systems* 25.1 (2009), pp. 341–350.

[30] Seyed Abbas Taher, Masoud Hajiakbari Fini, and Saber Falahati Aliabadi. “Fractional Order PID Controller Design for LFC in Electric Power Systems using Imperialist Competitive Algorithm”. In: *Ain Shams Engineering Journal* 5.1 (2014), pp. 121–135.

[31] S Ali Pourmousavi and M Hashem Nehrir. "Introducing Dynamic Demand Response in the LFC Model". In: *IEEE Transactions on Power Systems* 29.4 (2014), pp. 1562–1572.

[32] Vandy Ganesh, K Vasu, and P Bhavana. "LQR based Load Frequency Controller for Two Area Power System". In: *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering* 1.4 (2012).

[33] H Shayeghi, HA Shayanfar, and OP Malik. "Robust Decentralized Neural Networks based LFC in a Deregulated Power System". In: *Electric Power Systems Research* 77.3-4 (2007), pp. 241–251.

[34] Mohamed I Mosaad and Fawzan Salem. "LFC based Adaptive PID controller using ANN and ANFIS Techniques". In: *Journal of Electrical Systems and Information Technology* 1.3 (2014), pp. 212–222.

[35] H Shayeghi, A Jalili, and HA Shayanfar. "Robust Modified GA based Multi-stage Fuzzy LFC". In: *Energy Conversion and Management* 48.5 (2007), pp. 1656–1670.

[36] Pukar Mahat, Zhe Chen, and Birgitte Bak-Jensen. "Review of Islanding Detection Methods for Distributed Generation". In: *2008 third international conference on electric utility deregulation and restructuring and power technologies*. IEEE. 2008, pp. 2743–2748.

[37] Zhihong Ye et al. "Evaluation of Anti-islanding Schemes based on Nondetection Zone Concept". In: *IEEE transactions on power electronics* 19.5 (2004), pp. 1171–1176.

[38] Faridaddin Katiraei, Mohammad Reza Iravani, and Peter W Lehn. "Micro-grid Autonomous Operation During and Subsequent to Islanding Process". In: *IEEE Transactions on power delivery* 20.1 (2005), pp. 248–257.

[39] Charles Concordia, Lester H Fink, and Geroge Poullikkas. "Load Shedding on an Isolated System". In: *IEEE Transactions on Power Systems* 10.3 (1995), pp. 1467–1472.

[40] ENTSO-E. *Report on Deterministic Frequency Deviations*. `https://consultations.entsoe.eu/system-development/deterministic_frequency_deviations_report/`. 2020.

[41] ENTSO-E. *Dispersed Generation Impact on CE Region, Dynamic study*. `https://www.entsoe.eu/data/data-portal/`. 2014.

[42] Uros Markovic et al. "Understanding Stability of Low-inertia Systems". In: (2019).

[43] Mohammad Dreidy, H Mokhlis, and Saad Mekhilef. "Inertia Response and Frequency Control Techniques for Renewable Energy Sources: A Review". In: *Renewable and sustainable energy reviews* 69 (2017), pp. 144–155.

[44] David L King et al. "Performance Model for Grid-connected Photovoltaic Inverters". In: *Sandia National Laboratories SAND2007-5036* (2007).

[45] Farzaneh Mirzapour et al. "A New Prediction Model of Battery and Wind-solar Output in Hybrid Power System". In: *Journal of Ambient Intelligence and Humanized Computing* 10.1 (2019), pp. 77–87.

[46] Emad M Natsheh, Alhussein Albarbar, and Javad Yazdani. "Modeling and Control for Smart Grid Integration of Solar/wind Rnergy Conversion System". In: *2011 2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies*. IEEE. 2011, pp. 1–8.

[47] Lewis M Fraas and Larry D Partain. *Solar Cells and their Applications*. Vol. 236. John Wiley & Sons, 2010.

[48] A Marti et al. "Novel Semiconductor Solar Cell Structures: The Quantum Dot Intermediate Band Solar Cell". In: *Thin solid films* 511 (2006), pp. 638–644.

[49] Miro Zeman. *Solar Cells*. Accessed: 2020-06-20. `https://ocw.tudelft.nl/courses/solar-cells/`, 2020.

[50] Hongmei Tian et al. *Detailed Performance Model for Photovoltaic Systems*. Tech. rep. NREL/JA-5500-54601. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2012.

[51] J. A. Kratochvil D. L. King W. E. Boyson. *Photovoltaic Array Performance Model*. Tech. rep. SAND2004-3535. Sandia, 2007.

[52] G. M. Galbraith D. L. King S. Gonzalez. *Performance Model for Grid-Connected Photovoltaic Inverters*. Tech. rep. SAND2007-5036. Sandia, 2007.

[53] *Collegato alla rete pugliese il più grande parco fotovoltaico d'Italia*. `https://www.rinnovabili.it/energia/fotovoltaico/parco-fotovoltaico-grande-italia/`. In Italian, Accessed on 20/06/2020.

[54] Xi Fang et al. "Smart Grid – The New and Improved Power Grid: A Survey". In: *IEEE communications surveys & tutorials* 14.4 (2011), pp. 944–980.

[55] Simon H. Tindemans, Vincenzo Trovato, and Goran Strbac. "Frequency Control Using Thermal Loads Under the Proposed ENTSO-E Demand Connection Code". In: *2015 IEEE Eindhoven PowerTech*. June 2015, pp. 1–6.

[56] Simon H. Tindemans, Vincenzo Trovato, and Goran Strbac. "Decentralized Control of Thermostatic Loads for Flexible Demand Response". In: *IEEE Transactions on Control Systems Technology* 23.5 (Sept. 2015), pp. 1685–1700.

[57] George C Konstantopoulos, Antonio T Alexandridis, and Panos C Papageorgiou. "Towards the Integration of Modern Power Systems into a Cyber - Physical Framework". In: *Energies* 13.9 (2020), p. 2169.

[58] Prashant V Kamat. "Quantum Dot Solar Cells. Semiconductor Nanocrystals as Light Harvesters". In: *The Journal of Physical Chemistry C* 112.48 (2008), pp. 18737–18753.

[59] Jianghai Hu, John Lygeros, and Shankar Sastry. "Towards a Theory of Stochastic Hybrid Systems". In: *International Workshop on Hybrid Systems: Computation and Control*. Springer. 2000, pp. 160–173.

[60] Henk AP Blom et al. *Stochastic Hybrid Systems: Theory and Safety Critical Applications*. Vol. 337. Springer, 2006.

[61] Fedor Shmarov et al. "SMT-based Synthesis of Safe and Robust PID Controllers for Stochastic Hybrid Systems". In: *Haifa Verification Conference*. Springer. 2017, pp. 131–146.

[62] Alessandro Abate et al. “ARCH-COMP19 Category Report: Stochastic Modelling”. In: *ARCH@ CPSIoTWeek*. 2019, pp. 62–102.

[63] Kendra Lesser and Meeko Oishi. “Finite State Approximation for Verification of Partially Observable Stochastic Hybrid Systems”. In: *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*. 2015, pp. 159–168.

[64] John Lygeros and Maria Prandini. “Stochastic Hybrid Systems: a Powerful Framework for Complex, Large Scale Applications”. In: *European Journal of Control* 16.6 (2010), pp. 583–594.

[65] Alessandro Abate et al. “Approximate Model Checking of Stochastic Hybrid Systems”. In: *European Journal of Control* 16.6 (2010), pp. 624–641.

[66] Kai Lai Chung. “Markov Chains”. In: *Springer-Verlag, New York* (1967).

[67] John G Kemeny and J Laurie Snell. *Markov Chains*. Springer-Verlag, New York, 1976.

[68] Olle Häggström et al. *Finite Markov Chains and Algorithmic Applications*. Vol. 52. Cambridge University Press, 2002.

[69] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT press, 2008.

[70] D. Knuth and A. Yao. “Algorithms and Complexity: New Directions and Recent Results”. In: Academic Press, 1976. Chap. The complexity of nonuniform random number generation.

[71] Conrado Daws. “Symbolic and Parametric Model Checking of Discrete-time Markov Chains”. In: *International Colloquium on Theoretical Aspects of Computing*. Springer. 2004, pp. 280–294.

[72] Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. “Probabilistic Reachability for Parametric Markov Models”. In: *International Journal on Software Tools for Technology Transfer* 13.1 (2011), pp. 3–19.

[73] Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. “Approximate Parameter Synthesis for Probabilistic Time-bounded Reachability”. In: *2008 Real-Time Systems Symposium*. IEEE. 2008, pp. 173–182.

[74] Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina. “Parametric Probabilistic Transition Systems for System Design and Analysis”. In: *Formal Aspects of Computing* 19.1 (2007), pp. 93–109.

[75] Tobias Winkler et al. “On the Complexity of Reachability in Parametric Markov Decision Processes”. In: *arXiv preprint arXiv:1904.01503* (2019).

[76] Sebastian Junges et al. *Parameter Synthesis for Markov Models*. 2019. arXiv: `1903.07993 [cs.LO]`.

[77] Luca Cardelli et al. “Syntax-guided Optimal Synthesis for Chemical Reaction Networks”. In: *International Conference on Computer Aided Verification*. Springer. 2017, pp. 375–395.

[78] Aurélien Rizk et al. “On a Continuous Degree of Satisfaction of Temporal Logic Formulae with Applications to Systems Biology”. In: *International Conference on Computational Methods in Systems Biology*. Springer. 2008, pp. 251–268.

[79] Grégory Batt et al. “Robustness Analysis and Tuning of Synthetic Gene Networks”. In: *Bioinformatics* 23.18 (2007), pp. 2415–2422.

[80] Linda Herrmann et al. “Formal Parameter Synthesis for Energy-Utility-Optimal Fault Tolerance”. In: *European Workshop on Performance Engineering*. Springer. 2018, pp. 78–93.

[81] E Allen Emerson and Edmund M Clarke. “Characterizing Correctness Properties of Parallel Programs using Fixpoints”. In: *International Colloquium on Automata, Languages, and Programming*. Springer. 1980, pp. 169–181.

[82] Edmund M Clarke, E Allen Emerson, and Joseph Sifakis. “Model Checking: Algorithmic Verification and Debugging”. In: *Communications of the ACM* 52.11 (2009), pp. 74–84.

[83] M. Kwiatkowska, G. Norman, and D. Parker. “PRISM 4.0: Verification of Probabilistic Real-time Systems”. In: *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*. Ed. by G. Gopalakrishnan and S. Qadeer. Vol. 6806. LNCS. Springer, 2011, pp. 585–591.

[84] Claude Girault and Rüdiger Valk. *Petri Nets for Systems Engineering: a Guide to Modeling, Verification, and Applications*. Springer Science & Business Media, 2013.

[85] Rajeev Alur. “Timed automata”. In: *International Conference on Computer Aided Verification*. Springer. 1999, pp. 8–22.

[86] Rajeev Alur et al. “Hybrid Automata: An Algorithmic Approach to the Specification and Verification of Hybrid Systems”. In: *Hybrid systems*. Springer, 1992, pp. 209–229.

[87] Graham Clark, Stephen Gilmore, and Jane Hillston. “Specifying Performance Measures for PEPA”. In: *International AMAST Workshop on Aspects of Real-Time Systems and Concurrent and Distributed Software*. Springer. 1999, pp. 211–227.

[88] Costas Courcoubetis and Mihalis Yannakakis. “Verifying Temporal Properties of Finite-state Probabilistic Programs”. In: *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society. 1988, pp. 338–345.

[89] Joost-Pieter Katoen. “The Probabilistic Model Checking Landscape”. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. 2016, pp. 31–45.

[90] Christel Baier et al. “Model Checking Probabilistic Systems”. In: *Handbook of Model Checking*. Springer, 2018, pp. 963–999.

[91] Christian Dehnert et al. “A Storm is Coming: A Modern Probabilistic Model Checker”. In: *International Conference on Computer Aided Verification*. Springer. 2017, pp. 592–600.

[92] Ernst Moritz Hahn et al. “The 2019 Comparison of Tools for the Analysis of Quantitative Formal Models”. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer. 2019, pp. 69–92.

[93] Sadegh Esmaeil Zadeh Soudjani, Caspar Gevaerts, and Alessandro Abate. "FAUST^2: Formal Abstractions of Uncountable-STate STochastic Processes". In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer. 2015, pp. 272–286.

[94] Arnd Hartmanns and Holger Hermanns. "The Modest Toolset: An Integrated Environment for Quantitative Modelling and Verification". In: *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014*. Ed. by Erika Ábrahám and Klaus Havelund. Vol. 8413. Lecture Notes in Computer Science. Springer, 2014, pp. 593–598.

[95] Nathalie Cauchi and Alessandro Abate. "StocHy: Automated Verification and Synthesis of Stochastic Processes". In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer. 2019, pp. 247–264.

[96] Hans Hansson and Bengt Jonsson. "A Logic for Reasoning about Time and Reliability". In: *Formal aspects of computing* 6.5 (1994), pp. 512–535.

[97] Adnan Aziz et al. "It Usually Works: The Temporal Logic of Stochastic Systems". In: *International Conference on Computer Aided Verification*. Springer. 1995, pp. 155–165.

[98] Moshe Y Vardi. "Probabilistic Linear-time Model Checking: An Overview of the Automata-theoretic Approach". In: *International AMAST Workshop on Aspects of Real-Time Systems and Concurrent and Distributed Software*. Springer. 1999, pp. 265–276.

[99] Christos G Cassandras and John Lygeros. *Stochastic Hybrid Systems*. CRC Press, 2018.

[100] Costas Courcoubetis and Mihalis Yannakakis. "The Complexity of Probabilistic Verification". In: *Journal of the ACM (JACM)* 42.4 (1995), pp. 857–907.

[101] Rajeev Alur et al. "Discrete Abstractions of Hybrid Systems". In: *Proceedings of the IEEE* 88.7 (2000), pp. 971–984.

[102] Antoine Girard and George J Pappas. "Approximation Metrics for Discrete and Continuous Systems". In: *IEEE Transactions on Automatic Control* 52.5 (2007), pp. 782–798.

[103] Antoine Girard, Giordano Pola, and Paulo Tabuada. "Approximately Bisimilar Symbolic Models for Incrementally Stable Switched Systems". In: *IEEE Transactions on Automatic Control* 55.1 (2009), pp. 116–126.

[104] A Agung Julius and George J Pappas. "Approximations of Stochastic Hybrid Systems". In: *IEEE Transactions on Automatic Control* 54.6 (2009), pp. 1193–1203.

[105] Christel Baier et al. "Comparative Branching-time Semantics for Markov Chains". In: *Information and computation* 200.2 (2005), pp. 149–214.

[106] Kim G Larsen and Arne Skou. "Bisimulation through Probabilistic Testing". In: *Information and computation* 94.1 (1991), pp. 1–28.

[107] Josée Desharnais, Abbas Edalat, and Prakash Panangaden. "Bisimulation for Labelled Markov Processes". In: *Information and Computation* 179.2 (2002), pp. 163–193.

[108] Holger Hermanns. "Interactive Markov Chains". In: *Interactive Markov Chains*. Springer, 2002, pp. 57–88.

[109] Harold Joseph Kushner. *Approximation and Weak Convergence Methods for Random Processes, with Applications to Stochastic Systems Theory*. Vol. 6. MIT press, 1984.

[110] Xenofon D Koutsoukos. "Optimal Control of Stochastic Hybrid Systems Based on Locally Consistent Markov Decision Processes". In: *Proceedings of the 2005 IEEE International Symposium on, Mediterrean Conference on Control and Automation Intelligent Control, 2005.* IEEE. 2005, pp. 435–440.

[111] Maria Prandini and Jianghai Hu. "A Numerical Approximation Scheme for Reachability Analysis of Stochastic Hybrid Systems with State-dependent Switchings". In: *2007 46th IEEE Conference on Decision and Control*. IEEE. 2007, pp. 4662–4667.

[112] A. Abate et al. "Probabilistic Reachability and Safety for Controlled Discrete Time Stochastic Hybrid Systems". In: *Automatica* 44.11 (2008), pp. 2724–2734.

[113] Alessandro Abate and Sadegh Esmaeil Zadeh Soudjani. "Quantitative Approximation of the Probability Distribution of a Markov Process by Formal Abstractions". In: *Logical Methods in Computer Science* 11 (2015).

[114] Sean Summers and John Lygeros. "Verification of Discrete Time Stochastic Hybrid Systems: A Stochastic Reach-avoid Decision Problem". In: *Automatica* 46.12 (2010), pp. 1951–1961.

[115] Ernst Moritz Hahn et al. "Exploiting Robust Optimization for Interval Probabilistic Bisimulation". In: *International Conference on Quantitative Evaluation of Systems*. Springer. 2016, pp. 55–71.

[116] Morteza Lahijanian, Sean B Andersson, and Calin Belta. "Formal Verification and Synthesis for Discrete-time Stochastic Systems". In: *IEEE Transactions on Automatic Control* 60.8 (2015), pp. 2031–2045.

[117] Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate. "Aggregation and Control of Populations of Thermostatically Controlled Loads by Formal Abstractions". In: *IEEE Transactions on Control Systems Technology* 23.3 (2015), pp. 975–990.

[118] Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate. "Adaptive and Sequential Gridding Procedures for the Abstraction and Verification of Stochastic Processes". In: *SIAM Journal on Applied Dynamical Systems* 12.2 (2013), pp. 921–956.

[119] IRENA. *Renewable Power Generation Cost*. `https://www.irena.org/publications/2020/Jun/Renewable-Power-Costs-in-2019`. 2020.

[120] European Commission. *Commission Regulation (EU) 2016/631 of 14th April 2016*. `https://op.europa.eu/en/home`. 2016.

[121] *Wind Generation: Ireland*. `http://smartgriddashboard.eirgrid.com/#roi/wind`. Accessed on 10/12/2018.

[122] *Renewable Summary Report*. `http://www.eirgridgroup.com/how-the-grid-works/renewables/`. Accessed on 10/06/2020.

[123] Hassan Bevrani, Arindam Ghosh, and Gerard Ledwich. "Renewable Energy Sources and Frequency Regulation: Survey and New Perspectives". In: *IET Renewable Power Generation* 4.5 (2010), pp. 438–457.

[124] Qi Luo, Kartik B Ariyur, and Anoop K Mathur. "Control-oriented Concentrated Solar Power Plant Model". In: *IEEE Transactions on Control Systems Technology* 24.2 (2015), pp. 623–635.

[125] Yu Su et al. "Adaptive PV Frequency Control Strategy Based on Real-time Inertia Estimation". In: *IEEE Transactions on Smart Grid* (2020).

[126] Jianmin Chen, Fuwen Yang, and Qing-Long Han. "Model-Free Predictive H_∞ Control for Grid-Connected Solar Power Generation Systems". In: *IEEE Transactions on Control Systems Technology* 22.5 (2014), pp. 2039–2047.

[127] Zhongwen Li et al. "Adaptive Power Point Tracking Control of PV System for Primary Frequency Regulation of AC Microgrid with High PV Integration". In: *IEEE Transactions on Power Systems* (2021).

[128] Thanikanti Sudhakar Babu, Dalia Yousri, and Karthik Balasubramanian. "Photovoltaic Array Reconfiguration System for Maximizing the Harvested Power using Population-based Algorithms". In: *IEEE Access* 8 (2020), pp. 109608–109624.

[129] Manoj Datta and Tomonobu Senjyu. "Fuzzy Control of Distributed PV Inverters/Energy Storage Systems/Electric Vehicles for Frequency Regulation in a Large Power System". In: *IEEE Transactions on Smart Grid* 4.1 (2013), pp. 479–488.

[130] Haytham A Mostafa, Ramadan El-Shatshat, and Magdy MA Salama. "Multi-objective Optimization for the Operation of an Electric Distribution System with a Large Number of Single Phase Solar Generators". In: *IEEE Transactions on Smart grid* 4.2 (2013), pp. 1038–1047.

[131] Mohamed E Elkhatib, Wei Du, and Robert H Lasseter. "Evaluation of Inverter-based Grid Frequency Support using Frequency-watt and Grid-forming PV Inverters". In: *2018 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE. 2018, pp. 1–5.

[132] Peter Lilienthal. *HOMER® micropower optimization model*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2005.

[133] PSS/E: Power Systems Simulation and Analysis. *version 35.2*. Munich, Germany: Siemens AG, 2020.

[134] MATLAB. *version 9.5.0 (R2018b)*. Natick, Massachusetts: The MathWorks Inc., 2018.

[135] Hamid Shaker, Hamidreza Zareipour, and David Wood. "Estimating Power Generation of Invisible Solar Sites using Publicly Available Data". In: *IEEE Transactions on Smart Grid* 7.5 (2016), pp. 2456–2465.

[136] Hamid Shaker, Hamidreza Zareipour, and David Wood. "A Data-driven Approach for Estimating the Power Generation of Invisible Solar Sites". In: *IEEE Transactions on Smart Grid* 7.5 (2015), pp. 2466–2476.

[137] Joe Brown, Alessandro Abate, and Alex Rogers. "Disaggregation of Household Solar Energy Generation using Censored Smart Meter Data". In: *Energy and Buildings* 231 (2021), p. 110617.

[138] David Angeli and Panagiotis-Aristidis Kountouriotis. "A Stochastic Approach to 'Dynamic-Demand' Refrigerator Control". In: *IEEE Transactions on Control Systems Technology* 20.3 (May 2012), pp. 581–592.

[139] Maryam Kamgarpour et al. "Modeling Options for Demand Side Participation of Thermostatically Controlled Loads". In: *2013 IREP Symposium Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid*. IEEE. 2013, pp. 1–15.

[140] Sadegh Esmaeil Zadeh Soudjani et al. "Formal Synthesis and Validation of Inhomogeneous Thermostatically Controlled Loads". In: *International Conference on Quantitative Evaluation of Systems*. 2014, pp. 57–73.

[141] Arnd Hartmanns and Holger Hermanns. "Modelling and Decentralised Runtime Control of Self-stabilising Power Micro Grids". In: *International Symposium On Leveraging Applications of Formal Methods, Verification and Validation*. Springer. 2012, pp. 420–439.

[142] Arnd Hartmanns, Holger Hermanns, and Pascal Berrang. "A Comparative Analysis of Decentralized Power Grid Stabilization Strategies". In: *Proceedings of the Winter Simulation Conference*. Winter Simulation Conference. 2012, p. 158.

[143] Holger Hermanns and Arnd Hartmanns. "An Internet Inspired Approach to Power Grid Stability". In: *it-Information Technology* 55.2 (2013), pp. 45–51.

[144] Soumya Kundu and Nikolai Sinitsyn. "Safe Protocol for Controlling Power Consumption by a Heterogeneous Population of Loads". In: *2012 American Control Conference (ACC)*. IEEE. 2012, pp. 2947–2952.

[145] Nikolai A Sinitsyn, Soumya Kundu, and Scott Backhaus. "Safe Protocols for Generating Power Pulses with Heterogeneous Populations of Thermostatically Controlled Loads". In: *Energy Conversion and Management* 67 (2013), pp. 297–308.

[146] Stephan Koch, Johanna L Mathieu, and Duncan S Callaway. "Modeling and Control of Aggregated Heterogeneous Thermostatically Controlled Loads for Ancillary Services". In: *Proc. PSCC*. Citeseer. 2011, pp. 1–7.

[147] Johanna L Mathieu, Stephan Koch, and Duncan S Callaway. "State Estimation and Control of Electric Loads to Manage Real-time Energy Imbalance". In: *IEEE Transactions on power systems* 28.1 (2012), pp. 430–440.

[148] Johanna L Mathieu and Duncan S Callaway. "State Estimation and Control of Heterogeneous Thermostatically Controlled Loads for Load Following". In: *2012 45th Hawaii International Conference on System Sciences*. IEEE. 2012, pp. 2002–2011.

[149] Sergio Grammatico et al. "A Mean Field Control Approach for Demand Side Management of Large Populations of Thermostatically Controlled Loads". In: *2015 European Control Conference (ECC)*. IEEE. 2015, pp. 3548–3553.

[150] Luminita Cristiana Totu, Rafael Wisniewski, and John Leth. "Demand Response of a TCL Population using Switching-rate Actuation". In: *IEEE Transactions on Control Systems Technology* 25.5 (2016), pp. 1537–1551.

[151] Luminita C Totu, John Leth, and Rafael Wisniewski. "Control for Large Scale Demand Response of Thermostatic Loads". In: *2013 American Control Conference*. IEEE. 2013, pp. 5023–5028.

[152] Yanbo Che et al. "Demand Response From the Control of Aggregated Inverter Air Conditioners". In: *IEEE Access* 7 (2019), pp. 88163–88173.

[153] Sean P Meyn et al. "Ancillary Service to the Grid using Intelligent Deferrable Loads". In: *IEEE Transactions on Automatic Control* 60.11 (2015), pp. 2847–2862.

[154] Axel Legay, Benoit Delahaye, and Saddek Bensalem. "Statistical Model Checking: An Overview". In: *International conference on runtime verification*. Springer. 2010, pp. 122–135.

[155] Koushik Sen, Mahesh Viswanathan, and Gul Agha. "On Statistical Model Checking of Stochastic Systems". In: *International Conference on Computer Aided Verification*. Springer. 2005, pp. 266–280.

[156] Edmund M Clarke and Paolo Zuliani. "Statistical Model Checking for Cyber-physical Systems". In: *International Symposium on Automated Technology for Verification and Analysis*. Springer. 2011, pp. 1–12.

[157] Gul Agha and Karl Palmskog. "A Survey of Statistical Model Checking". In: *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 28.1 (2018), pp. 1–39.

[158] Alexandre David et al. "An Evaluation Framework for Energy Aware Buildings using Statistical Model Checking". In: *Science China information sciences* 55.12 (2012), pp. 2694–2707.

[159] Gerd Behrmann et al. "Uppaal 4.0". In: (2006).

[160] Peter Bulychev et al. "UPPAAL-SMC: Statistical Model Checking for Priced Timed Automata". In: *arXiv preprint arXiv:1207.1272* (2012).

[161] Toni Mancini et al. "Demand-aware Price Policy Synthesis and Verification Services for Smart Grids". In: *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE. 2014, pp. 794–799.

[162] Ahsan Shahid. "A Cyber-physical Approach for Stochastic Hybrid Control and Safety Verification of Smart Grids". In: *2014 IEEE Innovative Smart Grid Technologies-Asia (ISGT ASIA)*. IEEE. 2014, pp. 721–725.

[163] Souymodip Chakraborty et al. "Modelling and Statistical Model Checking of a Microgrid". In: *International Journal on Software Tools for Technology Transfer* 17.4 (2015), pp. 537–554.

[164] Toni Mancini et al. "Parallel Statistical Model Checking for Safety Verification in Smart Grids". In: *2018 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE. 2018, pp. 1–6.

[165] Gilles Notton, V Lazarov, and L Stoyanov. "Optimal Sizing of a Grid-connected PV System for Various PV Module Technologies and Inclinations, Inverter Efficiency Characteristics and Locations". In: *Renewable Energy* 35.2 (2010), pp. 541–554.

[166] B. Lazpita M. Jung O. Wiss. *Analyses et conclusions – Tests en sous-frequence.* Tech. rep. DTS/RT/2O16/072. Cea Tech, 2016.

[167] B. Lazpita M. Jung O. Wiss. *Analyses et conclusions – Tests en sur-frequence.* Tech. rep. DTS/RT/2016/008. Cea Tech, 2016.

[168] John G. Kassakian et al. "The Future of the Electric Grid". In: *Massachusetts Institute of Technology, Tech. Rep* (2011), pp. 197–234.

[169] M. Pathare et al. "Designing and Implementation of Maximum Power Point Tracking (MPPT) Solar Charge Controller". In: *2017 International Conference on Nascent Technologies in Engineering (ICNTE).* Jan. 2017, pp. 1–5.

[170] Moacyr Aureliano Gomes De Brito et al. "Evaluation of the Main MPPT Techniques for Photovoltaic Applications". In: *IEEE transactions on industrial electronics* 60.3 (2013), pp. 1156–1167.

[171] Boualem Bendib, Hocine Belmili, and Fateh Krim. "A survey of the Most Used MPPT Methods: Conventional and Advanced Algorithms Applied for Photovoltaic Systems". In: *Renewable and Sustainable Energy Reviews* 45 (2015), pp. 637–648.

[172] Boualem Bendib, Hocine Belmili, and Fateh Krim. "A survey of the most used MPPT methods: Conventional and Advanced Algorithms Applied for Photovoltaic Systems". In: *Renewable and Sustainable Energy Reviews* 45 (May 2015), pp. 637–648.

[173] Eftichios Koutroulis, Kostas Kalaitzakis, and Nicholas C Voulgaris. "Development of a Microcontroller-based, Photovoltaic Maximum Power Point Tracking Control System". In: *IEEE Transactions on power electronics* 16.1 (2001), pp. 46–54.

[174] Nicola Femia et al. "Optimization of Perturb and Observe Maximum Power Point Tracking Method". In: *IEEE transactions on power electronics* 20.4 (2005), pp. 963–973.

[175] Kenji Kobayashi, Ichiro Takano, and Yoshio Sawada. "A Study of a Two Stage Maximum Power Point Tracking Control of a Photovoltaic System Under Partially Shaded Insolation Conditions". In: *Solar energy materials and solar cells* 90.18-19 (2006), pp. 2975–2988.

[176] Hiren Patel and Vivek Agarwal. "Maximum Power Point Tracking Scheme for PV Systems Operating Under Partially Shaded Conditions". In: *IEEE transactions on industrial electronics* 55.4 (2008), pp. 1689–1698.

[177] Chian-Song Chiu. "TS Fuzzy Maximum Power Point Tracking Control of Solar Power Generation Systems". In: *IEEE Transactions on Energy Conversion* 25.4 (2010), pp. 1123–1132.

[178] E Karatepe, T Hiyama, et al. "Artificial Neural Network-polar Coordinated Fuzzy Controller based Maximum Power Point Tracking Control under Partially Shaded Conditions". In: *IET Renewable Power Generation* 3.2 (2009), pp. 239–253.

[179] Mohammed A Elgendy, Bashar Zahawi, and David J Atkinson. "Assessment of Perturb and Observe MPPT Algorithm Implementation Techniques for PV Pumping Applications". In: *IEEE transactions on sustainable energy* 3.1 (2011), pp. 21–33.

[180] Trishan Esram and Patrick L Chapman. "Comparison of Photovoltaic Array Maximum Power Point Tracking Techniques". In: *IEEE Transactions on energy conversion* 22.2 (2007), pp. 439–449.

[181] Dean Banerjee. *PLL Performance, Simulation and Design*. Dog Ear Publishing, 2006.

[182] Yang Han et al. "Comparative Performance Evaluation of Orthogonal-signal-generators-based Single-phase PLL Algorithms - A Survey". In: *IEEE Transactions on Power Electronics* 31.5 (2015), pp. 3932–3944.

[183] Arnold O Allen. *Probability, Statistics, and Queueing Theory*. Academic press, 2014.

[184] Edwin T Jaynes. *Probability Theory: The Logic of Science*. Cambridge university press, 2003.

[185] European Commission. *Commission Regulation (EU) 2017/1485 of 2 August 2017*. `https://op.europa.eu/en/home`. 2017.

[186] Operations Directorate of Energy Networks Association. "Recommendations for the Connection of Type Tested Small-scale Embedded Generators (Up to 16A per Phase) in Parallel with Low-Voltage Distribution Systems". In: (2012).

[187] Comitato Elettrotecnico Italiano. "Regola Tecnica di Riferimento per la Connessione di Utenti Attivi e Passivi alle Reti BT delle Imprese Distributrici di Energia Elettrica". In: (2020).

[188] ENTSO-E. *Assessment of the System Security with Respect to Disconnection Rules of Photovoltaic Panels*. `https://www.entsoe.eu/data/data-portal/`. 2012.

[189] *National Grid Status*. `https://www.gridwatch.templar.co.uk/`. Accessed: 10/12/2018.

[190] ENTSO-E. *Network Code on Operational Security*. `https://www.entsoe.eu/data/data-portal/`. 2013.

[191] E. B. Iversen, J. M. Morales, and J. K. Moller. "Probabilistic Forecasts of Solar Irradiance by Stochastic Differential Equations". In: *Envirometrics* 25 (Apr. 2014).

[192] ENTSO-E. *Load-frequency Control and Performance, Appendix 1*. `https://eepublicdownloads.azureedge.net/clean-documents/pre2015/publications/entsoe/Operation_Handbook/Policy_1_Appendix_final.pdf`. 2018.

[193] ENTSO-E. *Continental Europe Operation Handbook – Glossary*. `https://www.entsoe.eu/data/data-portal/`. 2004.

[194] David Jones. "Dynamic System Parameters for the National Grid". In: *IEEE Proceedings - Generation, Transmission and Distribution* 152.1 (Jan. 2005), pp. 53–60.

[195] Prabha Kundur, Neal J Balu, and Mark G Lauby. *Power System Stability and Control*. Vol. 7. McGraw-hill New York, 1994.

[196] Johan Morren et al. "Wind Turbines Emulating Inertia and Supporting Primary Frequency Control". In: *IEEE Transactions on Power Systems* 21.1 (Feb. 2006), pp. 433–434.

[197] Tractebel Engineering S.A. *Eurostag*. `http://www.eurostag.be`. 1997–2020.

[198] Richard C Dorf and Robert H Bishop. "Modern Control Systems". In: (2011).

[199] Katsuhiko Ogata et al. *Discrete-time Control Systems*. Vol. 2. Prentice Hall Englewood Cliffs, NJ, 1995.

[200] GR Aghajani, HA Shayanfar, and H Shayeghi. "Demand Side Management in a Smart Micro-grid in the Presence of Renewable generation and Demand Response". In: *Energy* 126 (2017), pp. 622–637.

[201] T. K. Chau et al. "Demand-Side Regulation Provision From Industrial Loads Integrated With Solar PV Panels and Energy Storage System for Ancillary Services". In: *IEEE Transactions on Industrial Informatics* 14.11 (2018), pp. 5038–5049.

[202] M Zinaddinov and S Mil'shtein. "Solar Tracking with Anti-Tracking Support for Ancillary Service". In: *2019 IEEE 46th Photovoltaic Specialists Conference (PVSC)*. IEEE. 2019, pp. 2091–2096.

[203] ENTSO-E. *Common Grid Model Exchange Standard*. `https://www.entsoe.eu/digital/cim/cim-for-grid-models-exchange/`. 2017.

[204] Norman L Johnson, Samuel Kotz, and Narayanaswamy Balakrishnan. *Univariate continuous distributions*. 1994.

[205] Erwin H Bareiss. "Computational Solutions of Matrix Problems over an Integral Domain". In: *IMA Journal of Applied Mathematics* 10.1 (1972), pp. 68–104.

[206] Olaf Stursberg and Bruce H Krogh. "Efficient Representation and Computation of Reachable Sets for Hybrid Systems". In: *International Workshop on Hybrid Systems: Computation and Control*. Springer. 2003, pp. 482–497.

[207] Winston A Richards et al. "An Efficient Polynomial Approximation to the Normal Distribution Function and its Inverse Function". In: *Journal of Mathematics Research* 2.4 (2010), p. 47.

[208] Milton Abramowitz and Irene A Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Vol. 55. US Government printing office, 1948.

[209] Inc. Wolfram Research. *Mathematica, Version 12.0*. Champaign, IL, 2019.