Increasing the Reliability of Power and Communication Networks via Robust Optimization

Zur Erlangung des akademischen Grades Doktor-Ingenieur (Dr.-Ing.)
Genehmigte Dissertation von Allan Santos aus Santos, Brasilien
Tag der Einreichung: 21. Juni 2023, Tag der Prüfung: 11. Oktober 2023

1. Gutachten: Prof. Dr. rer. nat. Florian Steinke
2. Gutachten: Prof. Dr.-Ing. Amr Rizk, (habil.)
Darmstadt, Technische Universität Darmstadt

TECHNISCHE UNIVERSITÄT DARMSTADT

Electrical Engineering and Information Technology Department

Energy Information Networks & Systems

Increasing the Reliability of Power and Communication Networks via Robust Optimization

Accepted doctoral thesis by Allan Santos

Date of submission: 21. Juni 2023
Date of thesis defense: 11. Oktober 2023

Darmstadt, Technische Universität Darmstadt

Erklärungen laut Promotionsordnung

§ 8 Abs. 1 lit. c PromO

Ich versichere hiermit, dass die elektronische Version meiner Dissertation mit der schriftlichen Version übereinstimmt.

§ 8 Abs. 1 lit. d PromO

Ich versichere hiermit, dass zu einem vorherigen Zeitpunkt noch keine Promotion versucht wurde. In diesem Fall sind nähere Angaben über Zeitpunkt, Hochschule, Dissertationsthema und Ergebnis dieses Versuchs mitzuteilen.

§ 9 Abs. 1 PromO

Ich versichere hiermit, dass die vorliegende Dissertation selbstständig und nur unter Verwendung der angegebenen Quellen verfasst wurde.

§ 9 Abs. 2 PromO

Die Arbeit hat bisher noch nicht zu Prüfungszwecken gedient.

Darmstadt, 21. Juni 2023

A. Santos

Zusammenfassung

Unsicherheit spielt bei der Planung und dem Betrieb komplexer, vernetzter Infrastrukturen eine immer wichtigere Rolle. Die zunehmende Einbeziehung variabler erneuerbarer Energien in Stromsysteme macht die Gewährleistung grundlegender Netzanforderungen wie Übertragungsleitungsbeschränkungen und Leistungsgleichgewicht zwischen Angebot und Nachfrage komplizierter. Ebenso variiert der Datenverkehr in Kommunikationsnetzwerken stark mit Benutzerpräferenzen und Dienstverfügbarkeit, und da Kommunikationsnetzwerke aufgrund der Zunahme netzwerkfähiger Geräte mehr Verkehr als je zuvor transportieren, ist die Bewältigung der hochgradig variablen Datenflüsse zwischen Servern und Endnutzern immer herausfordernder.

In diesem Zusammenhang schlagen wir in dieser Dissertation neue anpassungsfähige Methoden zur Optimierung von Flüssen in Energie- und Kommunikationssystemen vor, die explizit die wachsende Variabilität in diesen Systemen berücksichtigen, um einen optimalen Betrieb mit einem flexiblen Grad an Zuverlässigkeit zu gewährleisten. Die vorgeschlagenen Methoden verwenden ein robustes Optimierungsframework, um Beschränkungen, die von unsicheren Faktoren abhängen, handhabbar zu machen, indem ursprünglich stochastische Bedingungen durch deterministische Gegenstücke ersetzt werden. Der Hauptvorteil robuster Methoden besteht darin, dass sie sicherstellen, dass das System für alle Werte der unsicheren Variablen innerhalb einer gegebenen kontinuierlichen Menge möglicher Realisierungen realisierbar ist. Dies kann jedoch zu übermäßig konservativen Lösungen führen. Daher untersuchen wir auch, wie die Konservativität der vorgeschlagenen Algorithmen reduziert werden kann.

Diese Arbeit konzentriert sich auf zwei Klassen von Problemen in Energie- und Kommunikationssystemen, nämlich die adaptive Flusssteuerung und die Platzierung von flusskontrollierenden Geräten. In Stromversorgungssystemen bezieht sich die Flusssteuerung auf Maßnahmen, die Änderungen der von den Leitungen übertragenen Leistung bewirken, um einen bestimmten Zielwert zu minimieren oder zu maximieren, wobei die physikalischen Beschränkungen des Stromnetzes berücksichtigt werden. Einige Beispiele für die Steuerung der Leistungsflüsse sind die Änderung des Zustands von Schaltgeräten, die Regelung der Sollwerte von Generatoren und die Verwendung von sogenannten Flexible AC Transmission Systems (FACTS). Für die beiden letztgenannten Einflussmöglichkeiten

schlagen wir einen robusten Ansatz zur Optimierung dieser vor. Bei Kommunikationsnetzen wird die (Daten-)Flusssteuerung an jedem Router des Netzes implementiert. Diese Router definieren, anhand von Routing-Tabellen, den Pfad und die Rate, mit der die Daten weitergeleitet werden. Wir zeigen, dass es möglich ist, Regelgesetze zur Anpassung dieser Routing-Tabellen zu entwerfen, die den Datenfluss im Netz entsprechend der momentanen Rate der exogenen Eingaben in das System robust optimieren. Für beide Flussprobleme verwenden wir ein robustes Optimierungsframework, in dem affin-lineare Funktionen die Regelgesetze der Flusssteuerung parametrisieren. Die parametrisierten Regelgesetze können gemäß den Systembeschränkungen durch lineare oder quadratische Programmierung effizient berechnet werden.

Darüber hinaus betrachten wir im Bereich Platzierung zur Verbesserung der Zuverlässigkeit der Netzsysteme das Problem der FACTS-Platzierung und der Einbettung virtueller Netzwerke in ein bestehendes Kommunikationssystem. Beide Probleme werden als robuste gemischt-ganzzahlige lineare Programme (MILP) formuliert. Da das Finden beweisbar optimale Lösungen in großen Netzen jedoch eine rechnerische Herausforderung darstellt, entwickeln wir Approximationsalgorithmen, die nahezu optimale Ergebnisse liefern können und dabei um ein Vielfaches schneller zu lösen sind als das ursprüngliche MILP. In dem vorgeschlagenen robusten Framework werden die Probleme der Flusssteuerung und der Platzierung von Steuergeräten gemeinsam gelöst, um die Kopplungseffekte der beiden Optimierungsmaßnahmen zu berücksichtigen.

Wir demonstrieren die vorgeschlagene Methodik in einer Reihe von Anwendungsfällen in Energie- und Kommunikationssystemen. Unter anderen betrachten wir auch Anwendungen in intelligenten Stromnetzen (Smart Grids), bei denen Kommunikation und Stromnetze eng miteinander verknüpft sind. Die Kommunikationsinfrastruktur ermöglicht beispielsweise die Überwachung des Zustands von Stromnetzen in Echtzeit und die rechtzeitige Übermittlung von Steuersignalen an Geräte zur elektrischen Flusssteuerung. Da Smart Grids sich schneller an Änderungen der Betriebsbedingungen, aufgrund der zunehmenden Zahl erneuerbarer, variabler Energiequellen, anpassen und gleichzeitig anwendungsabhängige Zuverlässigkeitsanforderungen erfüllen müssen, tragen die in dieser Arbeit entwickelten robusten Optimierungsmethoden dazu bei, die Synergien zwischen flexiblen Energie- und Kommunikationssystemen zu nutzen sowie einen sicheren und effizienten Smart-Grid-Betrieb zu ermöglichen.

Abstract

Uncertainty plays an increasingly significant role in the planning and operation of complex networked infrastructure. The inclusion of variable renewable energy in power systems makes ensuring basic grid requirements such as transmission line constraints and the power balance between supply and demand more involved. Likewise, data traffic in communication networks varies greatly with user preferences and service availability, and with communication networks carrying more traffic than ever due to the surge in network-enabled devices, coping with the highly variable data flows between server and end-users becomes more crucial for the network's overall stability.

Within this context, we propose in this thesis new adaptable methods for optimizing flows in power and communication systems that explicitly consider the growing variability in these systems to guarantee optimal operation with a flexible degree of reliability. The proposed methods use a robust optimization framework, making constraints dependent on uncertain factors tractable by replacing originally stochastic conditions with deterministic counterparts. The primary benefit of robust methods is that they ensure the system is feasible for any values of the uncertain variables within a given continuous set of possible realizations. This, however, can lead to excessively conservative solutions. Therefore, we also investigate how to reduce the conservativeness of the proposed algorithms.

This thesis focuses on two classes of problems in power and communication systems, flow control and the placement of flow-controlling devices. In power systems, flow control refers to actions that induce changes in the power carried by transmission lines to minimize or maximize a specific objective value while considering the electrical grid's physical constraints. Some examples of power flow control actions are the change of switching equipment's state, regulation of generators' set points, and the management of the so-called Flexible AC Transmission Systems (FACTS) devices. For the last two action types, we propose a robust approach to optimize the corresponding control policies. As for communication networks, (data) flow control is implemented at each router in the network. These routers define the path and the rate data is forwarded using routing tables. We show that it is possible to robustly design policies to adapt these routing tables that optimize the data flows in the network depending on the instantaneous rate of the system's exogenous inputs. For both flow problems, we employ a robust optimization framework

where affine-linear functions parametrize the flow control policies. The parametrized policies can be efficiently computed via linear or quadratic programming, depending on the system's constraints.

Furthermore, we consider the placement problems in the form of FACTS placement and the embedding of virtual networks in an existing communication network to improve the reliability of the network systems. Both problems are formulated as robust Mixed-Integer Linear Programs (MILP). However, because finding provable optimal solutions in large networks is computationally challenging, we also develop approximate algorithms that can yield near-optimal results while being several times faster to solve than the original MILP. In the proposed robust framework, the flow control and the placement of controlling-devices problems are solved together to take into account the coupling effects of the two optimization measures.

We demonstrate the proposed methodology in a series of use cases in power and communication systems. We also consider applications in Smart Grids, where communication and electric networks are closely interlinked. E.g., communication infrastructure enables real-time monitoring of the status of power grids and sending timely control signals to devices controlling the electric flow. Due to the increasing number of renewable energy resources, Smart Grids must adapt to fast changes in operating conditions while meeting application-dependent reliability requirements. The robust optimization methods introduced in this thesis can thus use the synergy between flexible power and communication systems to provide secure and efficient Smart Grid operation.

Acknowledgements

First, I would like to thank my supervisors, Prof. Dr. Florian Steinke and Prof. Dr. Amr Rizk. Florian chose me for a student associate position in 2018. Since then, he has guided me for the whole of my master's and doctoral studies, almost six years in total. I cannot describe my gratitude for all his support. From late-night answers to my emails to (intensive) joint coding sessions. They were all extremely valuable to me. Florian is an inspiring example of professionalism I will take with me throughout my entire career. I also want to thank Amr for supporting me in this thesis and our other works. He introduced me to the communications field and was always kind and patient to answer my newbie questions.

I am also grateful for my colleagues at EINS, Julia, Edwin, Kirill, Johannes, Tim, Christopher, Martin, Mario, Jonas, Andy, Sara, Andrei, Sina, and Toby. I learned so much from all of you. Thanks for the interesting lunch conversations, exciting journal clubs, and confusing German lessons. After so much time at home due to Covid, I was really glad to see you all again in the office. I also want to mention Tina and Markus. I really appreciate everything you do for our department.

Lastly, I want to thank my family and friends. Mom and Dad, thank you for always being there for me. Eu amo vocês. Thank you, Barbara. We made it.

Contents

1 Introduction

1.1 Power and Communication Systems

Power and communication systems are ubiquitous in modern society and are an integral part of our everyday lives.

Since the late 19th century, electricity has increasingly gained higher importance for human development. Global electricity production achieved its all-time high in 2021 when it reached 28,000 TWh [95]. The growth in electricity generation was only possible due to advances in power generation technology, with a recent emphasis on large-scale renewable power generation, such as wind parks and solar farms. Following the expansion of power generation units, transmission networks have also received an increasing amount of investment to enable the delivery of the additionally generated electricity.

Similarly, the broad usage of telecommunications can also be traced back to the late 19th century. Forward to today, communication networks are ever-present in our daily lives. According to the World Bank, almost 90% of the population in Germany accessed the Internet in 2020 via laptops, smartphones, gaming consoles, etc [8]. The industry sector also profits substantially from telecommunications. Communicating-capable sensors and actuators form the backbone of highly automated control systems that guarantee the fast and reliable operation of industrial processes.

The purpose of both power and communication systems could be abstractly phrased as *the transport of a commodity from an originator to a recipient using the most efficient strategy while considering the system's physical constraints.* Power systems aim to supply consumers with electricity reliably and efficiently using an extensive infrastructure of cables, transformers, and electrical circuitry. Analogously, communication networks have the goal of transporting information as fast as possible from one point to another via different communication media and devices. Applications define further requirements for both systems, such as extra reliability, minimal throughput or delay, or maximal transportation costs.

The topological structures in power and communication systems are also similar to some degree. Both systems are made of several interconnected networks with topologically

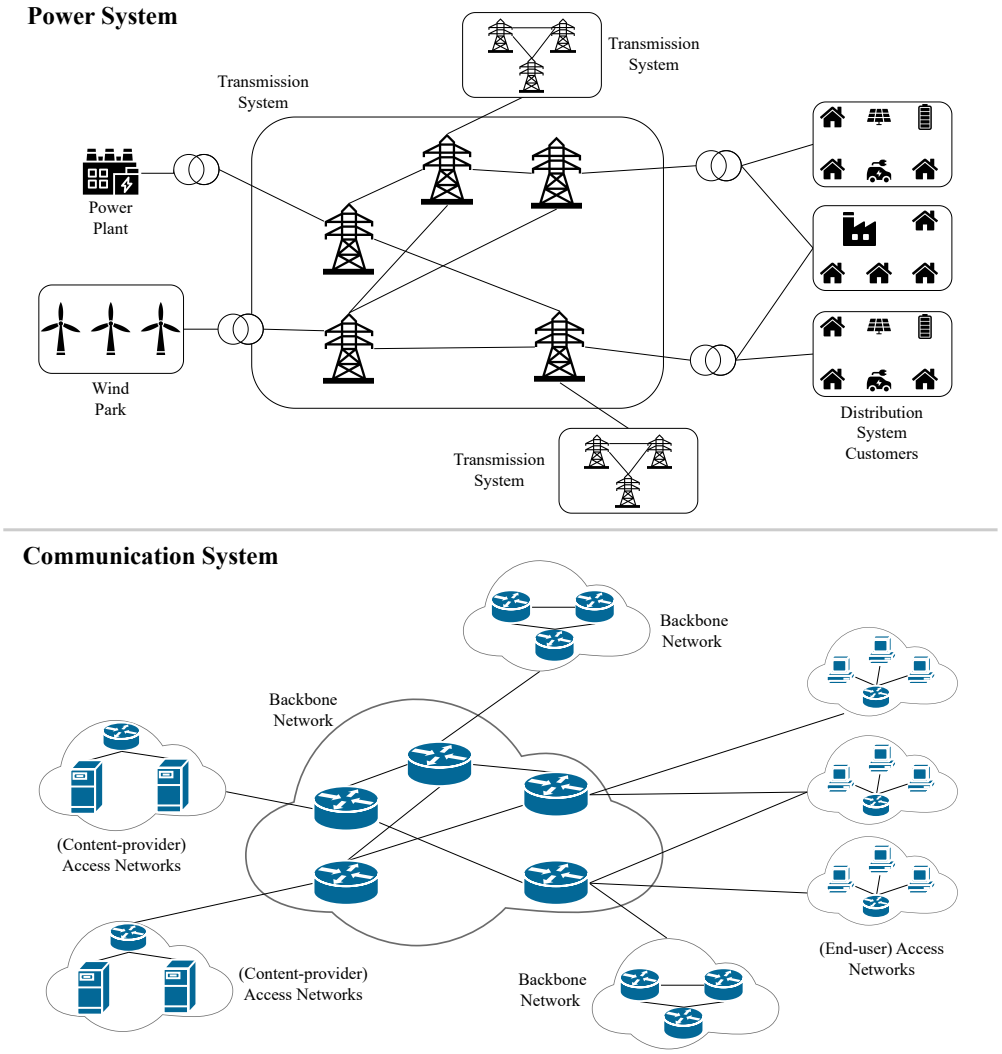

Figure 1.1: Typical architecture of power and communication systems.

similar structures. As illustrated in Fig. 1.1, both systems have in their center a meshed core network that interconnects originators and recipients. They are complemented by radial grids that connect the many end consumers to the core.

The meshed subsection of power systems is called the *transmission grid*. It transports large quantities of power, about 1.5 GW per circuit, from generators to consumers through

lines hundreds of kilometers in length at high voltage to minimize losses. Connections between two transmission grids also exist and aim to increase the reliability of the individual grids and reduce energy prices. E.g., some studies have shown that increasing the transmission capacity between neighboring transmission systems can have a significant impact in diminishing curtailment of renewable generation [92]. Power transformers are used at the intersection with the radial sub-networks, called *distribution grids*, to step down the voltage. Today with the increasing number of Decentralized Energy Resources (DER), such as small photovoltaic panels, energy can also flow from distribution grids to the transmission grid. In the case of a generation surplus in a distribution network, energy is redirected to other distribution networks where there is a residual power demand.

Likewise, meshed core communication networks, called *backbone networks*, interconnect *access networks* using high-speed optical links which transmit data at more than 100 Gbit/s. Backbone networks are linked with each other forming a communication medium, the Internet, that connects service providers to end-users all around the globe. Routers make the interface among the different sub-networks. In-hardware flow policies define the path to which the routers forward data and are also responsible for flow-control actions in the network.

However, the similarities between power and communication networks only hold for some of their characteristics. Most notably, in contrast to electricity, communication messages are non-fungible, i.e., they are not interchangeable. Thus, not only the amount of supply and demand of information must be the same, but data flows have specific sources and destinations. Another difference is that, while electricity flows continuously through transmission lines, servers and routers transmit messages in discrete packets of information. Lastly, electricity flows in a conductor from one point of higher voltage to another of lower voltage dictated by the laws of physics. On the other hand, information in communication networks has no physics-induced flow. The networking elements must actively decide on the routing of data packets.

This thesis proposes mechanisms for optimizing electricity and information flows considering the individual characteristics of power and communication systems. These mechanisms are united by a common framework whose goal is to compute adaptive optimal flow-control policies that can provide the system with different robustness levels. Additionally, this work not only considers power and communication systems individually but also investigates scenarios with close interdependence between them. Communication technology allows modern power grids to function more flexibly with the use of advanced metering and remote operation of controlling devices in the so-called Smart Grids. Example of Smart Grids functions are demand-side management [114], improved fault

detection [57], and self-healing [80]. In this context, this thesis also proposes algorithms that consider the combined flow control of power and communication systems to enable more efficient and reliable operation of Smart Grids.

1.2 The Role of Uncertainty

Uncertainty has always affected the operation of power and communication systems. However, today, with the growth in variability in energy generation and data traffic, it is more important than ever to devise mechanisms that can handle the system's uncertainty while ensuring a high degree of reliability despite the uncertain factors. Different methods have been proposed to counteract uncertainty in network systems in the last few years. But there are still underexplored topics, especially considering the novel adaptation capabilities of both power and communication systems, which have the potential to enable variable levels of robustness to the system using timely and economically advantageous responses to changes in the system's state.

The large-scale adoption of variable renewable generation at all grid levels has brought novel opportunities and challenges for power grid operators. In 2021, 17 GW of new wind and 26 GW of new solar photovoltaic (PV) capacity was added to the European energy system, further increasing the share of renewable generation in the energy mix [120, 39]. The power output of these new installations depends on the instantaneous availability of wind and solar irradiation, which can vary greatly throughout different seasons, days, and within the same day. To illustrate the variance in power output of renewable generation, we show in Fig. 1.2 the range of wind and PV generation as a fraction of the total load across a single week in a subsection of the German transmission system. The graph showcases that at times, wind and PV can provide up to 70% of the grid's energy demand, while at other times, the coverage may be as low as 10%. This variance is caused by variables that can be predicted with high accuracy and whose effect on potential power generation is well understood, such as the period of sunlight in a day, but also by factors that are harder to be taken into account in deterministic models. Wind speed, e.g., is rather difficult to predict [116]. The fact that we cannot always infer deterministic causal relationships between environmental variables and the power output of variable renewable generation leads to grid models with irreducible uncertainty. This uncertainty must be considered in the planning and operation of power systems, so that (*i*) the grid operates within its physical limitations for any realization of the uncertain power generation, and (*ii*) grid investments and operation are optimal regarding a significant performance indicator, e.g., the total expected costs or total expected CO2 emissions. Some measures that

transmission system operators have at hand to account for uncertain power production are the curtailment of renewable generation, grid expansion, and the installation and operation of flow-controlling devices. The first option is simple but leads to the waste of cheap energy and is undesired from an economic point of view. Grid expansion usually leads to long-term economic benefits, but its drawbacks are the potentially high initial investments and the resistance its implementation faces by the local population [83].

The deployment of flow-controlling devices in the transmission grid, the so-called Flexible AC Transmission Systems (FACTS), is a promising approach to cope with the growing uncertainty in power generation [78]. FACTS are static devices that can dynamically alter the power flow in the grid using fast-switching power electronic components. While FACTS have high initial investment costs, system operators can flexibly modify these devices' settings for different generation and load conditions that may arise in the future grid. This adaptability makes FACTS a cost-effective long-term solution for managing power flow, unlike grid expansions which often result in expensive and time-consuming construction projects [127].

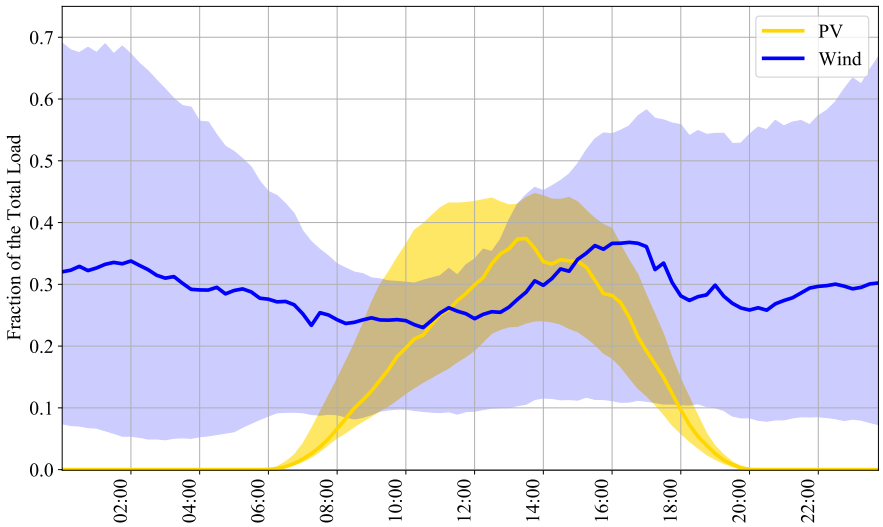

Figure 1.2: Wind and PV power generation as fraction of the total load in the 50Hz German transmission grid section from 25.08.2021-01.09.2021 (the solid lines show the average power generation, and the shaded areas represent the respective 95% percentile).[1]

[1] Source: ENTSO-E Transparency Platform (`https://transparency.entsoe.eu`)

Communication systems have also faced rapid changes in the last years with the increasing number of Internet services and connected devices. In particular, video traffic and Internet of Things (IoT) connections have led to rampant growth in Internet usage. According to Cisco's forecasts, 82% of the Internet traffic in 2021 was video content, and by 2023 almost 14.7 billion IoT devices will connect to the Internet daily [27, 26]. In this scenario, the uncertainty of data traffic grows due to technological factors, such as the possibility of connected devices changing local networks at will, and because of social aspects as the ever-changing user preferences for online services. Internet Service Providers (ISP) address the problem of maintaining adequate Quality of Service (QoE) while coping with highly variable data flows by employing fast routing hardware which can usually process data up to 100 Gbps [29] and path redundancy to increase the network capacity. Despite all the efforts from ISPs, high latency is still experienced by many users on an everyday basis, which can result from a bad experience when streaming a video to a complete failure of critical network systems, like Smart Grids.

New hardware and software solutions for data flow management in communication networks have emerged recently. In particular, the programmable data-plane technology promises to give additional flexibility to network operators by allowing routers to implement local flow-control policies. These policies are described by domain-specific programming languages, such as P4 [20], and act at line rate using different criteria to forward data packets. The programmable data-plane technology has already been proven to counteract congestions in backbone networks and reduce overhead communication with central network controllers, thus improving the overall network routing efficiency [56].

1.3 Performance and Reliability Guarantees in Network Systems

The increasing uncertainty in modern network systems poses a significant challenge to maintaining their safe operation. Thus, a critical practical question that arises is how to maintain and guarantee the reliability of such systems. Another question relevant to system operators is how different reliability levels affect the system's performance. To illustrate the complexities of these questions, we present two examples, one from power and one from communication systems. These examples demonstrate that answering these questions is not always straightforward, highlighting the critical role of flow optimization algorithms in the planning and operation of modern network systems.

To begin with, we need to define what reliability and performance mean in power and communication systems. In this work, we define the reliability of a network system as *its ability to remain within its physical limits despite the effects of uncertain factors.* System

operators often use historical service interruption data to quantify the reliability of a system. On the other hand, this thesis uses as reliability metric the range within which the uncertainty can vary while the system's constraints are still met. This approach to measuring reliability has the advantage that it does not require historical data and can, therefore, be used also in the design phase of network systems.

Performance can be generally defined as *a single or a set of metrics that indicates how good a system operates.* In power systems, performance can be measured, e.g., as the total investment and operation costs, CO2 emissions, or the amount of energy losses. In communication networks, it is usual to use metrics such as average delay, latency, cost of link utilization, or a mixture of those. In this thesis, optimality refers to the operation with highest possible performance.

Achieving a suitable level of reliability in a network system in the presence of uncertainty can be challenging, even in simple scenarios with only a single uncertain variable. To illustrate this point, we present an example of power grid with uncertain power generation. This example focuses on analyzing the steady-state behavior of the systems, i.e., it is assumed that state variables such as voltages in power systems are constant over time.

Example 1.3.1. The power grid in Fig. 1.3 features two conventional, dispatchable fossil fuel-fired power plants, one using expensive gas and the other using cheaper lignite. It also includes two wind parks, the preferred source of energy, and a single load of 4 in the per-unit system (p.u). The grid consists of three buses and three transmission lines. The gas power plant and the load are located on bus 1, the lignite power plant and a large wind park with peak production of 3 p.u. are located on bus 2, and a smaller wind park of 1 p.u is located on bus 3. The merit order algorithm is used to determine the output of the generators, with those having the lowest marginal production costs being brought online first.

Here the task is to ensure that the power flow in the grid is feasible for all possible variations in wind availability. Classical scenario-based decision rules select a limited number of "worst-case" scenarios and develop control laws for the power injections that guarantee system feasibility only for the considered cases. These methods assume that intermediate cases will cause less stress in the grid than the selected scenarios. For example, an intuitive pick for "worst-case" scenario is when the wind availability is at its maximum, meaning that the larger wind park produces 3 p.u. and the smaller one 1 p.u.. In this scenario, the power flows in lines (1,2), (1,3), and (2,3) using the DC approximation are -2, -2, and 1 p.u., respectively, thus indicating that no capacity constraints are violated in this scenario. Under the DC approximation, it is assumed that the reactive component of the impedance in any given line is significantly larger than the resistive component.

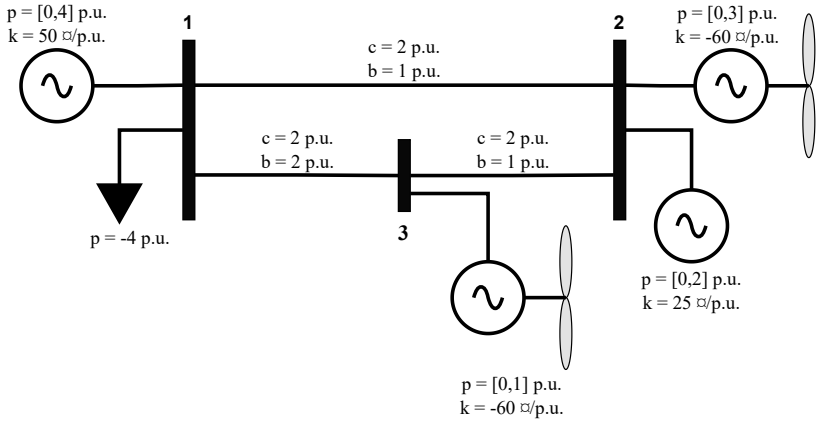

Figure 1.3: Schematic of 3-buses systems with variable wind power generation. The values for the generation upper and lower limits (p) and the line capacities (c) and susceptances (b) are shown in the image in the per-unit system (p.u.). Marginal generation costs (k) are expressed for a generic currency ¤.

As a result, the resistive portion is disregarded. Moreover, all voltage magnitudes are assumed to be at their nominal per-unit values.

The grid operator may assume that no violation of physical constraints will occur based on the previous result. However, the actual worst-case scenario with respect to the uncertainty is when wind availability is at half of its maximum value, i.e., when the larger wind park is producing 1.5 p.u. and the smaller one 0.5 p.u. In this scenario, the conventional power plant in bus 2 is also producing energy to meet the load demand as per merit order, resulting in lines (1,2) and (2,3) being overloaded by 0.2 and 0.3 p.u., respectively. Thus, if the system operator implements a control law that guarantees feasibility only when the wind availability is at its maximum, the physical constraints of the grid could be violated in at least one other scenario, which in practice could lead to contingencies in lines (1,2) and (2,3). □

As the example above illustrates, relying solely on a few extreme scenarios may not be enough to guarantee that a network system remains feasible for all possible variations of uncertain inputs. In real-world applications in power and communication systems with tens or hundreds of uncertain variables, manually selecting a limited number of scenarios becomes increasingly impractical as the number of possible combinations grows exponentially. Therefore, efficient algorithms that can guarantee the feasible operation of network systems for all possible realizations of the uncertainties are critical tools for

improving the reliability of these systems.

The next example shows how different reliability levels affect routing performance in communication systems. Because high reliability and high performance are often two contradicting goals, it is the operator's task to decide which property to prioritize according to the applications deployed onto the network.

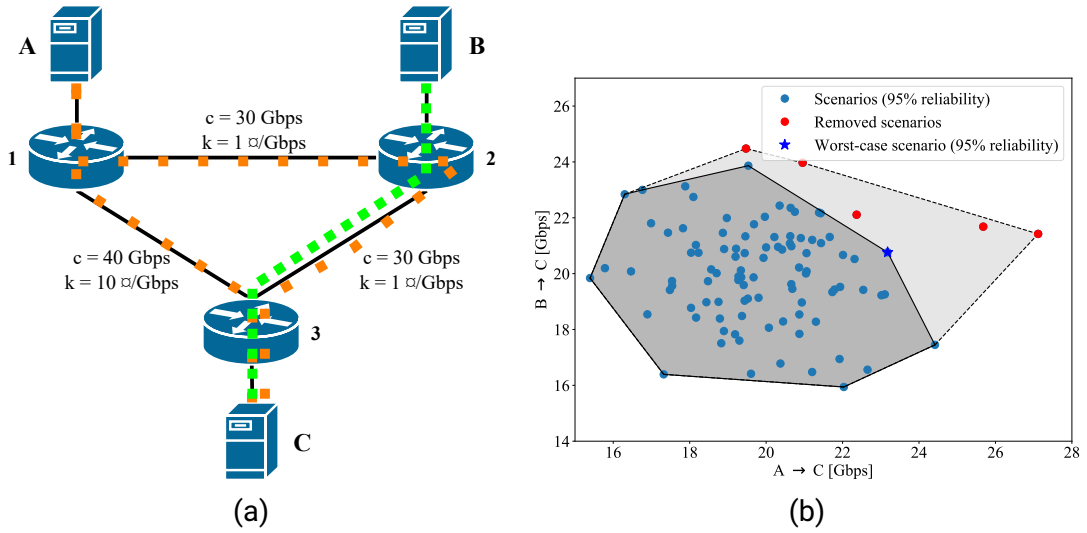

Figure 1.4: *(a)* Schematic of 3-router communication network with flows A→C (orange) and B→C (green). The capacity (c) and the cost of utilization of the links (k) in some currency ¤ are expressed in Gbps and $¤/\mathrm{Gbps}$, respectively. *(b)* Possible values for the rates of flows A→C and B→C. Points in red were removed to generate a set with 95% of the scenarios. The worst-case scenario of the new set is the scenario with the highest total flow rate. The shaded area shows the convex hull of the set with all scenarios (light gray) and with 95% of the scenarios (dark gray).

Example 1.3.2. Fig. 1.4a shows a communication network with three routers and three links forming a fully-connected topology. Two flows originate from servers A and B and have the same destination, server C. Flow A→C is transmitted via path {(1,2), (2,3)} and the expensive link (1,3). Flow B→C is transmitted only via link (2,3). The rate of the flows is assumed uncertain, and all possible realizations of the rates are shown in Fig. 1.4b.

Here the task is to minimize routing costs while ensuring that the total data flow over the network's links does not exceed their capacity. To design a routing policy that ensures

the grid is feasible for all possible flow rates originating from servers A and B, router 1 must route at least 37% of the flow A$\rightarrow$C through the expensive link (1,3). This is because, in the worst-case scenario, flows A$\rightarrow$C and B$\rightarrow$C sum to approximately 50 Gbps, while the capacity of link (2,3) is 30 Gbps.

On the other hand, if we remove 5 specific scenarios, i.e., realizations of the uncertain variables for which the routing must be feasible, it would be possible to reduce the amount of traffic through link (1,3) to 32%. In this case, the control policy would be 95% reliable, meaning that only 5% of possible values of the uncertain flows from servers A and B could cause a violation of the system's physical constraints. Unlike power systems, violations of flow constraints in communication networks do not damage the underlying physical infrastructure. Instead, they solely result in packet drops, which can be retransmitted by the message originator. Consequently, this reliability level is suitable for applications where no mission-critical information is transmitted and would significantly reduce the utilization cost compared to the routing policy with complete reliability. □

In summary, coping with uncertainty in network systems is a necessary but demanding task, as shown by the above examples. Hence, algorithms that can define the network's design and control policies according to flexible reliability and performance metrics are essential components in the toolbox of system managers and operators. With the rising amount of renewable generation and on-demand Internet services, power and communication systems face challenges that can only be tackled by automated and adaptive optimization methods.

1.4 Research Questions

In light of the above discussion on uncertainty management for the design and operation of networked power and communication systems, this dissertation delves deeper into the following key research questions (*RQ*):

(*RQ1*) How to determine adaptive flow control policies that provide a flexible trade-off between reliability and optimal operation when the system is subjected to uncertain inputs?

(*RQ2*) Where to optimally place flow-controlling devices in network systems to enable cost-effective flow control policies?

The above research questions aim to provide a deeper understanding of how to manage uncertainty in networked power and communication systems and to identify strategies to improve the flexibility, reliability, and cost-effectiveness of these systems.

Given these fundamental challenges, the thesis also considers more practical questions (*PQ*) regarding real-world implementation concerns:

(*PQ1*) How can flow control policies that are optimal and robust against input uncertainty be efficiently computed?

(*PQ2*) How to determine the optimal placement of flow-controlling devices in a reasonable time, even in large systems?

(*PQ3*) How to ensure that flow control policies are implementable in the network systems' hardware, i.e., meet their technical capabilities?

(*PQ4*) How can adaptive mechanisms be designed to ensure continuous robust operation of network systems without being overly conservative?

1.5 Contributions

Our research addresses the aforementioned research and practical questions by developing and analyzing several novel algorithms for robust optimization of flow control policies and placement of flow-controlling devices. These algorithms share two key features: (*i*) ability to ensure a flexible degree of reliability to the underlying network system and (*ii*) optimality regarding a specific metric while considering a constrained, steady state system.

We address *RQ1* in Chapters 3 and 4. In the context of power systems, we present a method to compute redispatch policies for power systems that minimize worst-case operation costs and ensure the physical constraints of the power grid for all possible values of uncertain power injections [107]. The resulting algorithm is a constrained Linear Program (LP). Chapter 4 shows the same methodology applied for communication networks with the goal of minimizing worst-case delays given uncertain data flows [105]. In this case, the algorithm for computing the optimal, robust flow control policies is a Quadratically-Constrained Quadratic Program (QCQP). In both flow control problems, affine-linear functions parametrize the policies of the system's actuators which by design respect the technical constraints of the network system's hardware devices, effectively addressing *PQ3*.

The optimal placement of flow-controlling devices problem posed in *RQ2* is investigated in Chapters 3 and 5. In Chapter 3, we present an algorithm to determine the optimal placement of FACTS in power systems [107]. The placement problem is embedded with the flow control problem and is formulated as a Mixed-Integer Linear Program (MILP).

The objective of the optimization problem is to minimize costs considering the trade-off between the investment costs due to the installation of new FACTS and the added flexibility those bring to the grid operation. In Chapter 5, we show an algorithm to compute the placement of logical elements in virtualized communication networks that leads to the lowest network utilization cost and that meets the requested reliability requirements [106]. The problem is also formulated as a MILP.

To address *PQ1* and *PQ2*, we present in Chapters 3 and 4 novel heuristics to solve MILPs and QCQPs, respectively, improving the tractability of the proposed robust optimization framework for large problems. These heuristics are orders of magnitude faster to compute compared to standard techniques, yet still able to find near-optimal solutions while ensuring the original reliability constraints [107, 105]. The heuristic for MILPs relies on solving the linear programming relaxation of the integer problem followed by a projection to the feasible set and an iterative hill-climbing search. As for QCQPs, we propose a linear relaxation of the quadratic constraints that is solved iteratively.

All of the algorithms developed in this thesis incorporate a robust framework to model the system's uncertainty. However, robust optimization approaches can be overly conservative, as noted by Ben-Tal and Nemirovski [13]. To overcome this limitation and address *PQ4*, we propose several methods to reduce the conservativeness of our algorithms.

In Chapter 3, we present a method to reduce the total amount of scenarios of the uncertain variables by exploiting the correlation between them [48, 107]. This method excludes scenarios with very low probability from the robust formulation, making the problem significantly less conservative. Chapter 4 shows an adaptive mechanism to reduce the conservativeness of routing policies in communication networks [105]. The algorithm iteratively adapts the range of the uncertain scenarios that must be considered by the robust optimization problem using real-time measurements of the uncertain data flows and computes optimal policies for each iteration. The mechanism can be implemented in routing devices found in commercial communication networks using state-of-the-art technologies such as Software-Defined Networking (SDN) and programmable data-plane.

Overall, our methods and heuristics aim to balance flexibility and conservativeness by reducing the total number of uncertain scenarios considered and using adaptive mechanisms and efficient heuristics for solving large robust optimization problems.

Several parts of this thesis have been previously published by the author and will not be specifically referenced in the rest of the dissertation. The list of these publications is:

[107] Allan Santos and Florian Steinke. "Robust placement and control of phase-shifting transformers considering redispatch measures." In: *Energies* 16.11 (2023), p. 4438.

[105] Allan Santos, Amr Rizk, and Florian Steinke. “Adaptive global coordination of local routing policies for communication networks.” In: *Computer Communications* 204 (2023), pp. 101–108.

[106] Allan Santos, Amr Rizk, and Florian Steinke. “Flexible redundancy generation for virtual network embedding with an application to smart grids.” In: *Proceedings of the Eleventh ACM International Conference on Future Energy Systems* (2020), pp. 97–105.

1.6 Related Literature

Previous research related to this thesis originates in the areas of robust optimization, flow control, and placement problems in power and communication systems. To provide context for this work, a brief overview of these fields will be given in the following. A comprehensive examination of the specific challenges and relevant literature for each algorithm presented in this thesis can be found in the corresponding chapters.

Robust Optimization. Robust optimization is a methodology for handling optimization problems with uncertainty [10]. It has been applied in several scientific fields, like control theory [130], machine learning [113], and operations research [7]. In a robust formulation, the probability distribution of the uncertain variables is not explicitly considered, which can be an advantage since identifying multi-dimensional distributions is challenging. Instead, a robust optimization program uses the set that supports these distributions, the so-called uncertainty set, to replace the probabilistic constraints with deterministic counterparts. Unlike other optimization methods for dealing with uncertainty, such as stochastic programming [100], robust optimization can enforce hard constraints in the solution space. That is, the solution to the optimization problem must satisfy the constraints for all considered realizations of the uncertain data. This feature is essential in critical applications where a single constraint violation can lead to enormous damage. Another advantage of robust optimization is that the robust counterpart of many common problems is computationally tractable and can be solved efficiently with linear programming or semidefinite programming [14, 13].

Conservativeness is usually a concern in robust optimization. Because of this, different methods to flexibly adjust the formulation to different robustness requirements have been proposed. For example, the authors of [16] propose restricting the number of uncertain coefficients that can vary simultaneously according to the desired level of robustness. [85] introduces a scenario-based approach where the portion of the probability mass enclosed by the uncertainty sets depends on adjustable robustness requirements. The works of [48,

107] also use scenario-based methods to control the uncertainty set's conservativeness but improve on [85] by considering the correlation of the uncertain parameters.

Optimal flow in power and communication systems. Network flow problems arise from several disciplines of science and engineering [4]. In this work, we focus on the minimum-cost flow problem, i.e., finding the cheapest control policy for transporting a commodity from its source to its destination. Several combinatorial algorithms to solve the minimum-cost flow problem have been developed over the years [3]. Still, linear programming has been proven to outperform them for the general case [72], and it is thus investigated in this dissertation.

In power systems, the optimal power flow is an instance of the minimum-cost flow problem, in which the system operator must define the operational set points of the grid's generators and the configuration of the flow-controllable devices to minimize costs while observing physical constraints, such as line capacities and maximum allowed voltage deviations [31]. Optimization models usually use approximations for the power flow to ensure tractability and to provide optimality guarantees. A popular approximation is the DC-power flow which assumes that the voltage magnitude on the buses is close to the grid's nominal value and that voltage angles are close to zero [99]. These conditions are observed mostly in transmission grids. Under these assumptions, the power flow equations become linear functions of the voltage angles, which can then be solved using linear programming. This approximation method can be computed significantly faster than the original AC-power flow model and can also avoid convergence issues faced by AC-power flow solvers [94, 71]. Chapter 3 covers solutions that optimize power flows in grids with uncertain power injections.

Flow optimization in communication systems, often called traffic engineering or routing optimization, is also a minimum-cost flow problem. Due to the increasing demand for Internet services, routing optimization has become necessary for avoiding bottlenecks in backbone networks [119], and it is required to prevent delays in mission-critical communications systems [19, 77, 82]. Although data flows in communication networks are discrete, continuous approximations are typically used, so methods for continuous-flow network systems can also be employed to solve the optimal routing problem [98, 2, 70]. In Chapter 4, we investigate different data flow models and algorithms to optimize routing in communication networks.

Optimal placement problems in power and communication systems. In this thesis, we cover the problems of optimal placement of FACTS devices and virtual functions in power and communication systems, respectively. FACTS are electronic devices that improve the power system's controllability [35]. Examples of FACTS are Static Var Compensators

(SVC), Thyristor-Controlled Series Capacitors (TCSC), and Thyristor-Controlled Phase-Shifting Transformers (PST). These devices, however, are expensive, so determining the optimal number and location of FACTS is necessary when planning and extending the power grid. Mixed-integer linear programming is a common approach for optimal placement of FACTS [79, 88, 111]. The main drawback of this method is that solving MILPs is, in the general case, NP-hard, which means that solving the placement problem for very large settings can be intractable. For this reason, approaches such as the Alternating Direction Method of Multipliers (ADMM), genetic algorithms, and different heuristics are used [33, 49, 108]. In Chapter 3, we formulate the FACTS placement problem as a MILP and propose a linear-relaxation heuristic to speed up computations.

In the context of network virtualization, one of the main tasks of network operators is to allocate service requests onto the underlying communication infrastructure, the so-called Virtual Network Embedding (VNE) problem. Service requests arrive in the form of Virtual Network Requests (VNR), and they specify the topology of the virtual network along with functional and non-functional requirements [45]. The optimality metric for the VNE problem varies with the use case. In this thesis, we consider the general cost minimization objective. The VNE problem can be solved to optimality via mixed-integer linear programming [21, 22, 51] or to sub-optimality using linear relaxations, heuristics, and learning algorithms [24, 129, 123]. Another line of work called Survivable VNE (SVNE) proposes increasing the resilience of virtual networks so they can still work even if some components of the underlying communication network fails [59, 110, 25]. In Chapter 5, we formulate the SVNE problem as a MILP to create routing redundancies based on the application's reliability requirements.

Based on the above literature research, the author has identified a critical research gap in the field of network system optimization. Specifically, there is a lack of methods that can simultaneously solve optimal flow and placement problems while accounting for system uncertainties in a robust fashion and respecting/leveraging the capabilities of novel flow-controlling hardware, such as FACTS and routers with programmable data-plane.

One significant reason why methods with these characteristics are highly desirable is the significant influence that flow-controlling devices' placement can have on the optimal flow in the system, as highlighted in Chapter 3. Additionally, it is crucial to consider system uncertainties, as discussed in Section 1.2, in the solution of flow and placement problems and to provide reliability guarantees despite the highly variable inputs in modern power and communication systems. Moreover, to ensure real-world applicability, it is important for optimization methods to be able to incorporate hard constraints regarding the network devices. For example, the maximum data rate of network routers or the ramp up/down constraints of power generators.

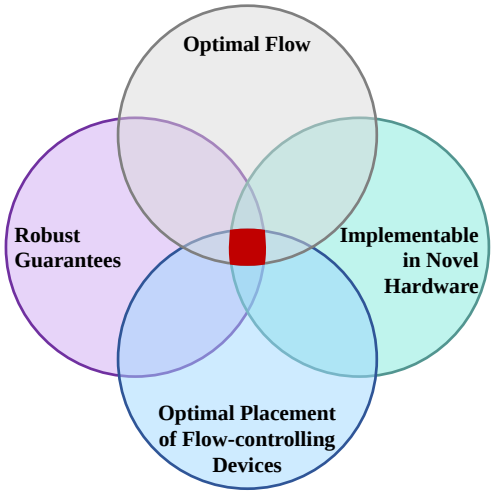

Figure 1.5: This thesis aims to address the research gap highlighted (in red): optimization methods to solve optimal flow and placement of flow-controlling devices jointly with robust guarantees on the feasibility of the network system and considering hardware capabilities of the network's flow-controlling devices.

To address this research gap, this thesis proposes an optimization framework that efficiently solves flow/placement problems in network systems to optimality while ensuring robustness against uncertainties and complying with the system's physical constraints Fig. 1.5 illustrates the different components of the research gap and where the proposed work will focus.

1.7 Thesis Outline

This thesis is structured into six chapters, as illustrated in Fig. 1.6.

Chapter 2 provides an overview of the robust optimization framework used throughout the thesis and derives closed-form, deterministic formulations for specific optimization problem classes. Additionally, it revisits the reliability versus optimality trade-off question and presents additional methodologies to address this issue.

Then the following three chapters are concerned with three novel applications of robust optimization to adaptive network systems:

- Chapter 3 addresses the placement of phase-shifting transformers and optimal flow control for the redispatch of power system generators. The problem is formulated

as a single-stage robust MILP, and redispatch policies are considered as cheaper alternatives to maintain the power system stable. A new greedy algorithm based on the robust MILP is proposed, which is orders of magnitude faster to compute compared to standard mixed-integer linear programming solvers while yielding near-optimal solutions. The algorithm is applied to a reference test grid for evaluation.

- Chapter 4 presents a novel robust QCQP formulation for the optimal routing of data packets in backbone networks. The outcome of the optimization is a set of affine-linear policies that are deployable in modern routing circuitry technology. The policies are updated periodically to avoid over-conservative routing strategies. An iterative heuristic is also devised to relax the quadratic constraints of the problem so that linear programming can be used to solve each iteration efficiently.

- Chapter 5 addresses the SVNE placement problem in communication networks and proposes a method for automated, adaptive redundancy generation. The minimum degree of redundancy is determined by the reliability requested by the VNR and underlying communication technology. The problem is formulated as a MILP. The algorithm is demonstrated in the context of voltage violation detection in Smart Grids, and it is shown to generate significant financial savings compared to fixed redundancy.

Finally, the thesis concludes with Chapter 6, which highlights the main contributions of the dissertation and suggests future research directions in the field of robust optimization in power and communication systems.

1.8 Mathematical Notation

Scalars are usually represented by lower- or uppercase, non-bold variables (example: x), while lowercase, bold variables ($\mathbf{x}$) denote vectors and uppercase, bold variables ($\mathbf{X}$) represent matrices in a Euclidean space. Sets are represented by uppercase, calligraphic variables ($\mathcal{X}$). The expression $\mathbb{P}_{\mathbf{x}\sim P}\{\mathbf{x} \in \mathcal{X}\}$ denotes the probability of the random variable $\mathbf{x} \sim P$ having a realization within the set $\mathcal{X}$. The symbol $\mathbb{R}^n$ represents the real coordinate space of dimension n, and $\mathbb{Z}^n$ represents the n-dimensional set of integers. The subscript $+$ added to either symbol represents the positive real or integer n-dimensional spaces, respectively.

Subscripts to non-bold variables (x_i) denote the element in the i-th dimension of vector $\mathbf{x}$. The index can also be the element of a non-ordered, general index set ($\mathcal{I}$). In matrices, double subscripts ($\mathbf{X}_{ij}$) denote the matrix element in row i and column j. Moreover, the

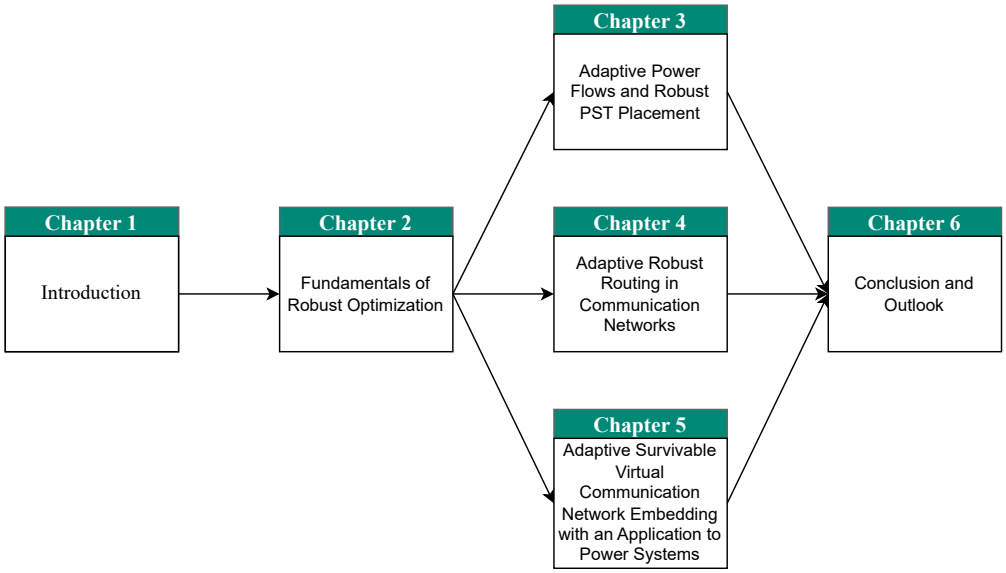

Figure 1.6: Thesis' structure.

matrix $\mathbf{X}_{(i:j)}$, $j > i$ is the submatrix made of the rows from i to j, and $\mathbf{X}_{i:}$ is the column vector of all the elements of the i-th row of $\mathbf{X}$.

The symbol $\circ$ represents the Hadamard product of two vectors or matrices, i.e., the element-wise multiplication. A diagonal matrix constructed from a vector x is denoted by $\text{diag}(\mathbf{x})$. The comparison operators $\geq$, $>$, $\leq$, and $<$ are defined element-wise for comparisons between vectors or matrices. The cardinality of a set $\mathcal{X}$ is given by $|\mathcal{X}|$, while the Cartesian product of two sets is represented by the symbol $\times$. The p-norm of $\mathbf{x} \in \mathbb{R}^n$ is denoted as $||\mathbf{x}||_p$, where $||\mathbf{x}||_p = (\sum_{i=1}^{n} |\mathbf{x}|^p)^{\frac{1}{p}}$, for $1 \leq p < \infty$.

The symbol $\mathbb{I}$ indicates an identity matrix of appropriate dimensions, and $\mathbf{1}$ is a vector of ones also with appropriate dimensions. The indicator function is represented by $\mathbb{1}_{\mathcal{X}}(x)$, where $\mathbb{1}_{\mathcal{X}}(x) = 1$ if $x \in \mathcal{X}$ and 0 otherwise.

2 Fundamentals of Robust Optimization

This chapter introduces the robust optimization framework used in this thesis to solve optimal flow and placement problems in power and communication systems. In Section 2.1, we first provide an overview of robust optimization and its relation to probabilistic methods for modeling systems under uncertainty. We then present the basic robust optimization framework utilized in this thesis for various applications, which will be explored in detail in the following chapters. In Section 2.2, we present the concept of an uncertainty set and discuss how the choice of the uncertainty set influences the solution of the robust optimization problem. Finally, in Section 2.3, we demonstrate how to reformulate standard classes of robust problems into computationally tractable forms. This is necessary because a robust optimization problem in its original form may have infinitely many constraints when the uncertainty set is continuous.

2.1 Overview of Optimization under Uncertainty

Optimization problems rely on the accurate description of the underlying system to produce meaningful solutions for decision-making. However, it is often the case that some of the inputs or parameters of the system are either unknown or intrinsically random. In science and engineering, imprecise measurements, variability in product manufacturing, and unpredictability of natural phenomena are examples of factors that deterministic models cannot capture in their entirety. This randomness can affect the model solution's optimality and the model's feasibility, as some constraints can be violated depending on the realization of the uncertain variables. Therefore, for a long time, researchers have been developing methodologies to specify model uncertainty and methods to solve these problems efficiently.

Stochastic optimization is a well-known approach to dealing with uncertainty in optimization problems [18, 100]. In this framework, the uncertainty is assumed to follow a certain probability distribution P, which can be known a priori or estimated via observations. One way of modeling a stochastic problem is via the so-called chance constraints. A classical stochastic problem is the uncertain Linear Program (LP), in which its decision

variable $\mathbf{x} \in \mathbb{R}^N$ must be minimized given random parameters $\mathbf{A} \in \mathbb{R}^{M \times N}$, $\mathbf{b} \in \mathbb{R}^N$ and $\mathbf{c} \in \mathbb{R}^N$:

$$\begin{aligned} \min_{\eta, \mathbf{x}} \ & \eta \\ \text{s.t.} \ & \mathbb{P}_{(\mathbf{A},\mathbf{b},\mathbf{c}) \sim P} \left\{ \mathbf{c}^T \mathbf{x} \leq \eta, \ \mathbf{A}\mathbf{x} \leq \mathbf{b} \right\} \geq 1 - \epsilon. \end{aligned} \tag{2.1}$$

Here, η is an auxiliary variable, ϵ specifies the maximum probability the problem's constraints can be violated, and $\mathbb{P}_{(\mathbf{A},\mathbf{b},\mathbf{c}) \sim P} \left\{ \mathbf{c}^T \mathbf{x} \leq \eta, \ \mathbf{A}\mathbf{x} \leq \mathbf{b} \right\}$ is the probability of $\mathbf{c}^T \mathbf{x} \leq \eta$ and $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ when $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ follow a distribution P.

Chance constraints, however, pose several practical difficulties when modeling and solving an optimization problem. In particular, specifying the probability distribution P is usually challenging. Unless P is already known a priori, determining it empirically requires an increasing number of samples the more dimensions the uncertain factors have. This fact makes stochastic programming impractical in many situations where not much historical data is available and/or obtaining new samples is costly. Another problem that arises is the computational tractability of such constraints, as they usually form a non-convex feasible set when no simplifications are applied [13].

In robust optimization, instead of specifying the uncertain parameters as random variables that follow a certain probability density function, we ensure that a set encloses all possible values that the uncertain variables can assume. The robust formulation of an uncertain LP is

$$\begin{aligned} \min_{\eta, \mathbf{x}} \ & \eta \\ \text{s.t.} \ & \mathbf{c}^T \mathbf{x} \leq \eta, \forall \mathbf{c} \in \mathcal{U}_c \\ & \mathbf{A}\mathbf{x} \leq \mathbf{b}, \ \forall \mathbf{A} \in \mathcal{U}_A, \ \forall \mathbf{b} \in \mathcal{U}_b, \end{aligned} \tag{2.2}$$

where $\mathcal{U}_A \subseteq \mathbb{R}^{M \times N}$, $\mathcal{U}_b \subseteq \mathbb{R}^N$ and $\mathcal{U}_c \subseteq \mathbb{R}^N$ are the *uncertainty sets* of parameters $\mathbf{A}$, $\mathbf{b}$, and $\mathbf{c}$, respectively.

If the uncertainty sets $\mathcal{U}_A$, $\mathcal{U}_b$, and $\mathcal{U}_c$ consist of a finite (and small) number of elements, solving problem (2.2) is straightforward: one can simply add a constraint for each combination of values in the uncertainty sets. However, when the uncertainty sets are continuous, the number of constraints becomes infinite, rendering the problem impractical to solve directly. Fortunately, the seminal works of El Ghaoui and Ben-Tal [36, 11] show that the robust counterpart of important generic convex problems under uncertainty are exactly tractable problems, enabling the application of efficient computational methods to solve them. Most notably, a robust LP with polytopic uncertainty sets is equivalent to a single-stage LP model, which can be solved in polynomial time.

Another advantage of robust optimization is that it is possible to specify an uncertainty set with very little information about the system. For example, one only needs to have plausible upper and lower boundaries of an uncertain variable to define an interval-based uncertainty set. On the other hand, stochastic optimization with chance constraints typically needs large amounts of data to estimate the probability distribution of the uncertain variables.

One could argue that the robust constraints (2.2) are more conservative than the chance constraints (2.1), and, thus, the robust counterpart of an uncertain optimization problem yields poorer performance compared to the stochastic formulation. However, different methods to design and adapt uncertainty sets have been developed over the years to adjust the reliability requirements of the underlying system flexibly. In Section 2.2, we present some approaches to reduce the conservatism of specific uncertainty sets.

Advancements in computing technology have expanded the applicability of robust optimization to large-scale problems in different fields, such as finance [41], logistics [126], product design [91], project planning [76], and system operation [64]. In this thesis, we focus on applying robust optimization in the planning and operation of network systems, which is an area that has received much attention in the last few years. Bertsimas et al. propose an adaptive two-level robust optimization framework to optimize the unit commitment and the dispatch levels of the generators of a power system [15]. The primary source of uncertainty in this problem is the variable renewable energy generation. The authors of [60, 61] use robust optimization to plan the extension of transmission networks also because of the variability of renewable energy resources. Furthermore, robust optimization has been employed in planning communication networks [9] and in controlling thereof [124], in which case the source of uncertainty is the rate of the exogenous data flows.

2.2 Methods for Constructing Uncertainty Sets

An important task when specifying a robust optimization problem is to define an appropriate uncertainty set for the application at hand. This set should enclose all possible values of the uncertain variables for which the model must be feasible and, at the same time, not be too large so the problem's solution is not too conservative. Usually, the larger the uncertainty set is, the poorer the model's performance, as more constraints are added to the problem. This behavior was demonstrated in Example 1.2.2. The trade-off between robustness and performance depends on how tolerable a constraint violation is in that particular instance. For example, losing data packets in a streaming service would cause

pixelated video frames and user dissatisfaction. In communication networks of critical infrastructures, such as power or transportation control systems, the loss of even a small number of data packets could trigger disruptive contingencies.

In the following, we give examples of methods to construct uncertainty sets based on observations of the uncertain variables. We focus on convex sets because they lead to robust problems that are easier to solve and whose optimality is proven.

Convex polytopic uncertainty sets are often used to describe the variance of uncertain variables. The simplest of those is the hyperrectangle, whose vertices are the combination of the maximum and minimum values of the uncertain variables. These extreme points are determined by the physical constraints of the system or are obtained from a set of observations. Let $\mathcal{U}^P$ be a convex polytopic uncertainty set that encloses the realizations of N uncertain variables $\mathbf{u} \in \mathbb{R}^N$. These variables are element-wise lower bounded by $\underline{\mathbf{u}} \in \mathbb{R}^N$ and upper bounded by $\overline{\mathbf{u}} \in \mathbb{R}^N$. If $\mathcal{U}^P$ is a hyperrectangle whose axes are the extreme values of $\mathbf{u}$, we can write it as the intersection of $2N$ half-spaces,

$$\mathcal{U}^P = \left\{ \mathbf{u} \in \mathbb{R}^N \mid \mathbf{u} \geq \underline{\mathbf{u}},\ \mathbf{u} \leq \overline{\mathbf{u}} \right\}. \tag{2.3}$$

Specifying the convex polytopic uncertainty set in (2.3) only requires the extreme realizations of the uncertain variables, which is useful when the probability distribution of the uncertainty is not known or when limited information of it is available. However, $\mathcal{U}^P$ can be excessively large because not necessarily all combinations of the extreme values are likely to occur. For example, the power output of a wind power plant will very unlikely be at its maximum while the output of another nearby wind power plant is at its minimum. Because of this, several methods to reduce the conservatism of polytopic uncertainty sets have been proposed, and we will show some of them next.

Bertsimas and Sim [16] proposed flexibly adjusting the conservatism level by restricting how much the uncertain variables can vary. The resulting uncertainty set is parametrized by a constant Γ that controls its size as

$$\mathcal{U}^D(\Gamma) = \left\{ \mathbf{u} \in \mathbb{R}^N \mid \exists \mathbf{z} \in \mathbb{R}^N_+ :\ \mathbf{u} = \overline{\mathbf{u}} + (\underline{\mathbf{u}} - \overline{\mathbf{u}}) \circ \mathbf{z},\ \mathbf{z} \leq \mathbf{1},\ \mathbf{1}^T\mathbf{z} \leq \Gamma \right\}. \tag{2.4}$$

In this method, when $\Gamma = 0$, the uncertainty set collapses to a single element $\{\overline{\mathbf{u}}\}$, i.e., the uncertainty does not vary. When $\Gamma = N$, the uncertainty set becomes the interval $[\underline{\mathbf{u}}, \overline{\mathbf{u}}]$, which is the same as the set $\mathcal{U}^P$. The authors prove that the probability of a constraint violation is bounded by $\exp\left(-\Gamma^2/2N\right)$ in the case the uncertain variables are independent and uniformly distributed in the interval $[-1, 1]$.

Margellos et al. [85] proposed a scenario-based approach to obtain an uncertainty set that encloses a flexible amount of the uncertainty's probability mass. The set is defined as

a hyperrectangle $\mathcal{U}^S(\epsilon)$, where the parameter $\epsilon \in [0,1]$ defines the probability a constraint can be violated. The vertices of $\mathcal{U}^S(\epsilon)$, $\{\underline{u}_i, \overline{u}_i\}$, $i = 1, \ldots, N$ are computed by solving N chance-constrained problems of the form:

$$\begin{aligned} &\min_{\overline{u}_i, \underline{u}_i} \ \overline{u}_i - \underline{u}_i \\ &\text{s.t.} \ \ \mathbb{P}\left(\mathbf{u} \in \mathbb{R}^N \mid \ u_i \geq \underline{u}_i, \ u_i \leq \overline{u}_i\right) \geq 1 - \epsilon_i, \end{aligned} \tag{2.5}$$

where $\sum_{i=1}^{N} \epsilon_i = \epsilon$.

The chance-constrained problems (2.5) are solved using a scenario-based approach in a stochastic-programming fashion. The chance constraints are replaced by a series of deterministic constraints, which ensure that the worst-case scenario of an observation set is contained in $\mathcal{U}^S(\epsilon)$. The minimum amount of elements S of this observation set to ensure a violation level ϵ with confidence $\beta \in [0,1]$ is

$$S = \left\lceil \frac{e}{\epsilon\,(e-1)} \left(N - 1 + \log\left(1/\beta\right)\right) \right\rceil, \tag{2.6}$$

where $\lceil x \rceil = \min\{m \in \mathbb{Z} \mid m \geq x\}$ is the ceiling function. For details on the assumptions needed for (2.6) and its guarantees, see [85, 5].

The equivalent scenario-based problem of (2.5) is

$$\begin{aligned} &\min_{\overline{u}_i, \underline{u}_i} \ \overline{u}_i - \underline{u}_i \\ &\text{s.t.} \ \ u_i^{\{k\}} \geq \underline{u}_i, \ k = 1, \ldots, S, \\ &\qquad\ \ u_i^{\{k\}} \leq \overline{u}_i, \ k = 1, \ldots, S, \end{aligned} \tag{2.7}$$

where $u_i^{\{k\}}$ is a scenario of the i-th component of the uncertain vector $\mathbf{u}$.

The hyperrectungular uncertainty set $\mathcal{U}^S(\epsilon)$, however, can be large and lead to overly conservative solutions in case the uncertainty's components are highly correlated. We exemplify this graphically in Fig. 2.1. To reduce conservatism, the authors of [48] and [107] use principal component analysis (PCA) [121] to specify a tighter hyperrectangle that is aligned with the principal components of the matrix of observations. The intersection between $\mathcal{U}^S(\epsilon)$ and the rotated hyperrectangle is then used as the uncertainty set. In both works, the high correlation of the uncertain power generation of renewable energy resources makes it possible to enclose all the realizations of the uncertainty by a much smaller polytope.

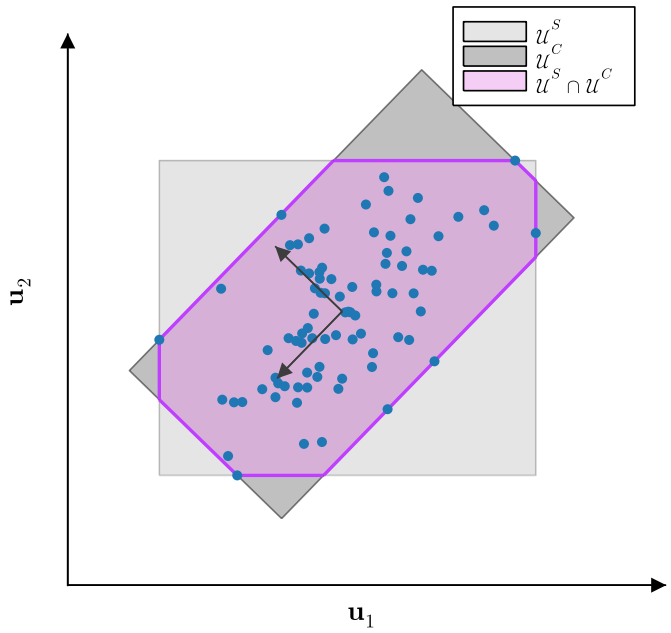

Figure 2.1: Uncertainty set defined as the intersection of the smallest non-rotated hyperrectangle $\mathcal{U}^S$ that encloses all scenarios, and a rotated hyperrectangle $\mathcal{U}^C$ in the direction of the data's principal components that also encloses all scenarios. Plot inspired by [48].

To obtain the principal components of the observation data, we first define the observation matrix $\mathbf{U} = \begin{bmatrix}\mathbf{u}^{\{1\}} & \dots & \mathbf{u}^{\{S\}}\end{bmatrix}^T \in \mathbb{R}^{S\times N}$, which contains S observations of the uncertain vector $\mathbf{u}$. With $\tilde{\mathbf{U}} \in \mathbb{R}^{S\times N}$ being the observation matrix $\mathbf{U}$ subtracted column-wise by the empirical mean of $\mathbf{u}$, we define the empirical covariance matrix as $\mathbf{\Sigma} = \frac{1}{S-1}\tilde{\mathbf{U}}^T\tilde{\mathbf{U}}$. The eigenvectors of $\mathbf{\Sigma}$ define the principal components of the observation matrix. These can be obtained via eigendecomposition, i.e., $\mathbf{\Sigma} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$, where $\mathbf{Q} \in \mathbb{R}^{N\times N}$ has the eigenvectors of $\mathbf{\Sigma}$ in its columns, and $\mathbf{\Lambda} \in \mathbb{R}^{N\times N}$ is a diagonal matrix with the corresponding eigenvalues on its diagonal.

The hyperrectangular uncertainty set aligned with the principal components of the uncertainty's observations is then specified as

$$\mathcal{U}^C = \left\{\mathbf{u} \in \mathbb{R}^N \mid \mathbf{D}\mathbf{u} \leq \mathbf{d}\right\}, \tag{2.8}$$

where $\mathbf{D} = \begin{bmatrix}\mathbf{Q} & -\mathbf{Q}\end{bmatrix}^T$ and $\mathbf{d} = \begin{bmatrix}\overline{\mathbf{d}} & -\underline{\mathbf{d}}\end{bmatrix}^T$. Vectors $\overline{\mathbf{d}} \in \mathbb{R}^N$ and $\underline{\mathbf{d}} \in \mathbb{R}^N$ are the extreme points of the observation set in the direction of the $\mathbf{Q}$ principal components and are defined element-wise as $\overline{d}_j = \max_i (\mathbf{UQ})_{ij}$ and $\underline{d}_j = \min_i (\mathbf{UQ})_{ij}$.

In Fig. 2.1, we exemplify the method used in [48, 107] for an artificially generated 2D data set. There we can clearly see that the rectangle oriented in the direction of the data set's principal components covers a much smaller area than the rectangle oriented in the canonical basis. This is because both components of the uncertain vector are highly correlated, and thus, most of the data variance is distributed along a single line. In the extreme case where both components are perfectly correlated, $\mathcal{U}^C$ would collapse into a line. In contrast, when the components have zero correlation, $\mathcal{U}^C$ would be equivalent to the non-rotated hyperrectangle $\mathcal{U}^S(\epsilon)$.

Apart from polytopic uncertainty sets, some works use ellipsoidal sets to enclose the model's uncertainty. See, e.g., [36, 12]. An ellipsoidal uncertainty set has the following general representation:

$$\mathcal{U}^Q(\rho) = \left\{\mathbf{u} \in \mathbb{R}^N \mid \hat{\mathbf{u}} + \mathbf{Vu},\ ||\mathbf{u}||_2 \leq \rho\right\}, \tag{2.9}$$

where $\hat{\mathbf{u}} \in \mathbb{R}^N$ is the center of the ellipsoid, and $\mathbf{V} \in \mathbb{R}^{N \times N}$ is a symmetric positive definite matrix. The parameter $\rho \in \mathbb{R}$ defines the radius of the ellipsoid and can be used to adjust the conservatism of the uncertainty set.

As in the polytopic case, data-based approaches can be used to determine $\mathcal{U}^Q(\rho)$. For example, one could specify $\mathbf{V}$ as the empirical covariance matrix of the observation set. Another possibility is to solve the minimum-volume covering ellipsoid problem, which is a log-determinant maximization problem subject to second-order cone constraints. It has complexity $O(N^{3.5} \log (N/\epsilon))$ [68] and can be solved via Interior-Point methods.

The maximum probability of a single constraint being violated with $\mathcal{U}^Q(\rho)$ was proven in [12] to be $\exp\left(-\rho^2/2\right)$ in the case the uncertain variables are independent, have zero mean and have support on the interval $[-1, 1]$.

All uncertainty sets discussed so far were continuous sets. However, the uncertainty of a model can also be discrete, which is the case in [106]. In this work, the failure modes of a communication network are uncertain, and because only total failures are considered, the uncertainty set is discrete. A straightforward option to deal with discrete uncertainty is to specify a continuous uncertainty set as in the continuous uncertainty case. This uncertainty set, however, can potentially induce too conservative solutions, as it adds constraints for values that are not in the support of the uncertain variables.

Another approach to applying the robust optimization framework in problems with discrete uncertain variables is to enforce the problem's constraints for each realization of the uncertainty, which is theoretically achievable as long as the number of realizations is

finite. However, the computational tractability of this method can be an issue if the number of realizations is too large because each realization results in an additional constraint. To tackle this problem, one could, e.g., eliminate realizations that are very unlikely to occur, as done in [48]. Scenario reduction techniques are common in power and communication systems. E.g., in problems that include failure modes, it is often assumed that no more than one component will fail at the same time [42, 25].

2.3 The Tractable Robust Counterpart of Common LP Models under Uncertainty

As discussed in the previous section, the uncertainty set of a robust optimization problem encloses all possible values of the uncertain variables for which the problem must be feasible. Considering a continuous uncertainty set, this means that the optimization problem has to satisfy an infinite amount of constraints, one for each value in the set, which is obviously computationally infeasible. Therefore, several methods to reformulate robust optimization problems to make them tractable have been developed.

In this section, we show the computationally tractable formulation of different classes of robust LP problems. We demonstrate that robust LP models with polytopic and ellipsoidal uncertainty sets are convex problems, namely linear problems and Second-Order Cone Programs (SOCP), respectively.

Consider the generic robust optimization problem with uncertain vectors $\mathbf{u}_j \in \mathbb{R}^Z,\ j = 1, \ldots, J$

$$\begin{aligned} &\min_{\mathbf{x}} \ f(\mathbf{x}) \\ &\text{s.t.} \quad g_j(\mathbf{x}, \mathbf{u}_j) \leq b_j,\ \forall \mathbf{u}_j \in \mathcal{U}_j,\ \forall j = 1, \ldots, J, \end{aligned} \tag{2.10}$$

where $\mathbf{x} \in \mathbb{R}^N$ is the vector of decision variables, $f : \mathbb{R}^N \to \mathbb{R}$ the objective function, and $g_j : \mathbb{R}^{NZ} \to \mathbb{R}$ the j-th constraint of optimization problem (2.10) with $b_j \in \mathbb{R}$ as its right-hand side.

It is important to notice that assuming the objective function in (2.10) to be independent of the uncertain vectors $\mathbf{u}_j$ has no loss of generality, as we can reformulate (2.10) to minimize an additional variable $\alpha \in \mathbb{R}$ that is constrained to be greater or equal than the maximum of the objective function, i.e., $\alpha \geq f(\mathbf{x}, \mathbf{u}_j),\ \forall \mathbf{u}_j \in \mathcal{U}_j,\ \forall j = 1, \ldots, J$. Furthermore, if $b_j,\ j = 1, \ldots, J$, are uncertain variables but interval bounded, we can take the lower bound of the interval as the right-hand side of the constraints, leading to

the same optimization problem (2.10) where $b_j,\ j = 1, \ldots, J$, are constants. We also assume that $\mathcal{U} = \mathcal{U}_i \times \cdots \times \mathcal{U}_j$, which also comes without loss of generality [12].

Although (2.10) is not tractable in the general case, some common problem classes like LPs with polytopic or ellipsoidal uncertainty sets are. In the following, we present the methodology for reformulating such problems to a tractable representation. For a detailed overview of tractable robust counterparts of uncertain optimization problems, see, e.g., [14].

A generic robust LP problem can be represented as

$$\begin{aligned} \min_{\mathbf{x}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \mathbf{A}_{j:}^T\mathbf{x} \leq b_j,\ \forall \mathbf{A}_{j:} \in \mathcal{U}_j,\ \forall j = 1, \ldots, J, \end{aligned} \tag{2.11}$$

with $\mathbf{c} \in \mathbb{R}^N$ being the cost vector and $\mathbf{A}_{j:} \in \mathbb{R}^N$ the parameters of the j-th constraint. In the case where $\mathcal{U}_j,\ j = 1, \ldots, J$, are polytopes with exterior representation $\mathcal{U}_j = \left\{\mathbf{u} \in \mathbb{R}^N \mid \mathbf{D}_j\mathbf{u} \leq \mathbf{d}_j\right\}$, with $\mathbf{D}_j \in \mathbb{R}^{M_j \times N},\ \mathbf{d}_j \in \mathbb{R}^{M_j}$, (2.11) is equivalent to the min-max problem

$$\begin{aligned} \min_{\mathbf{x}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \left\{ \begin{array}{ll} \max\limits_{\mathbf{A}_{j:}} & \mathbf{A}_{j:}^T\mathbf{x} \\ \text{s.t.} & \mathbf{D}_j\mathbf{A}_{j:} \leq \mathbf{d}_j \end{array} \right\} \leq b_j,\ \forall j = 1, \ldots, J. \end{aligned} \tag{2.12}$$

Proof. The equivalence of (2.11) and (2.12) can be proved with the following. Let $\mathbf{x}^* \in \mathbb{R}^N$ be a solution of (2.12). For any $\mathbf{A}_{j:} \in \mathcal{U}_j$, we have that ${\mathbf{A}_{j:}}^T\mathbf{x}^* \leq \max_{\mathbf{A}_{j:} \in \mathcal{U}_j} \mathbf{A}_{j:}^T\mathbf{x}^* \leq b_j,\ \forall j = 1, \ldots, J$. Thus, (2.11) is feasible for $\mathbf{x} = \mathbf{x}^*$ and has the same objective value as (2.12). On the other hand, a solution $\mathbf{x}' \in \mathbb{R}^N$ of (2.11) must satisfy the problem's constraints for all values of $\mathbf{A}_{j:}$, including for $\arg\max_{\mathbf{A}_{j:} \in \mathcal{U}_j} \mathbf{A}_{j:}^T\mathbf{x}',\ \forall j = 1, \ldots, J$. Thus, (2.12) is feasible for $\mathbf{x} = \mathbf{x}'$, and both problems have the same objective value. □

Now, assuming that $\mathcal{U}_j,\ \forall j = 1, \ldots, J$ are non-empty, strong duality holds for the inner maximization problem of (2.12), as it is an LP problem. Therefore, because the primal and dual optimal objective of (2.12) are equal, we can rewrite it into a dual-stage minimization problem:

$$\begin{aligned} \min_{\mathbf{x}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \left\{ \begin{array}{ll} \min\limits_{\boldsymbol{\lambda}_j} & \boldsymbol{\lambda}_j^T\mathbf{d}_j \\ \text{s.t.} & {\mathbf{D}_j}^T\boldsymbol{\lambda}_j = \mathbf{x} \\ & \boldsymbol{\lambda}_j \geq 0 \end{array} \right\} \leq b_j,\ \forall j = 1, \ldots, J, \end{aligned} \tag{2.13}$$

where $\boldsymbol{\lambda}_j \in \mathbb{R}^{M_j}$ are the dual variables of the j-th inner problem of (2.13).

Finally, we can then show that problem (2.13) is equivalent to the single-stage minimization problem

$$\begin{aligned} \min_{\mathbf{x},\boldsymbol{\lambda}_1,\ldots,\boldsymbol{\lambda}_j} \quad & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \quad & \boldsymbol{\lambda}_j^T\mathbf{d}_j \leq b_j,\ \forall j = 1,\ldots,J, \\ & \mathbf{D}_j{}^T\boldsymbol{\lambda}_j = \mathbf{x},\ \forall j = 1,\ldots,J, \\ & \boldsymbol{\lambda}_j \geq 0,\ \forall j = 1,\ldots,J. \end{aligned} \tag{2.14}$$

Proof. Let $\mathbf{x}^* \in \mathbb{R}^N,\ \boldsymbol{\lambda}_j^* \in \mathbb{R}^{M_j},\ \forall j = 1,\ldots,J$ be a solution of (2.14). Problem (2.13) is feasible for $\mathbf{x} = \mathbf{x}^*$ because there is at least one value of $\boldsymbol{\lambda}_j$, namely $\boldsymbol{\lambda}_j^*$, which makes both inner and outer minimization problems feasible, as $\min_{\boldsymbol{\lambda}_j} \boldsymbol{\lambda}_j^T\mathbf{d}_j \leq \boldsymbol{\lambda}_j^{*T}\mathbf{d}_j \leq b_j,\ \forall j = 1,\ldots,J$. Moreover, the objective value is equal in (2.14) and (2.13). On the other hand, if $\mathbf{x}' \in \mathbb{R}^N$ is a solution of (2.13), there must be a $\boldsymbol{\lambda}_j' \in \mathbb{R}^{M_j},\ \forall j = 1,\ldots,J$ that verifies the constraints of the inner minimization problem. The single-stage problem (2.14) is feasible for $\mathbf{x} = \mathbf{x}',\ \boldsymbol{\lambda}_j = \boldsymbol{\lambda}_j',\ \forall j = 1,\ldots,J$, and the objective value is the same of (2.13). Therefore, we can show that a robust LP problem with a polytopic uncertainty set is equivalent to another LP problem. □

It is worth noting that while strong duality does not hold for other classes of problems in general, the equivalence between equations (2.11) and (2.14) remains valid for problems that involve only a linear relation between the uncertain terms and the continuous decision variables. In such cases, the inner optimization problem in equation (2.12) is still an LP problem, and therefore, strong duality holds. The authors of [107] applied this reformulation to solve a robust MILP and in [105] to solve a robust QCQP problem.

For the particular case in which the uncertainty is interval-bounded, i.e., the uncertainty set is a hyperrectangle whose axes align with the canonical basis, the robust counterpart of an LP model has an alternative formulation as shown in [90].

Consider (2.11) with $\mathcal{U}_j = \left\{\mathbf{u} \in \mathbb{R}^N \mid \mathbf{u} \geq \underline{\mathbf{u}}_j,\ \mathbf{u} \leq \overline{\mathbf{u}}_j\right\},\ j = 1,\ldots,J$. With $\mathbf{H} \in \mathbb{R}^{J\times N}$ being an auxiliary matrix, we can tightly upper bound the left side of the problem's constraints as

$$\begin{aligned} & \mathbf{H}_{ji} \geq \underline{u}_{j_i}\mathbf{x}_i,\ \forall i = 1,\ldots,N,\ j = 1,\ldots,J, \\ & \mathbf{H}_{ji} \geq \overline{u}_{j_i}\mathbf{x}_i,\ \forall i = 1,\ldots,N,\ j = 1,\ldots,J. \end{aligned} \tag{2.15}$$

The robust problem (2.11) with an interval-bounded uncertainty set can then be written as the following LP problem:

$$\begin{aligned} \min_{\mathbf{x},\mathbf{H}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \mathbf{H}_{j:}^T\mathbf{1} \leq b_j, \ \forall j = 1,\dots,J, \\ & \mathbf{H}_{ji} \geq \underline{u}_{j_i}\mathbf{x}_i, \ \forall i = 1,\dots,N, \ \forall j = 1,\dots,J, \\ & \mathbf{H}_{ji} \geq \overline{u}_{j_i}\mathbf{x}_i, \ \forall i = 1,\dots,N, \ \forall j = 1,\dots,J. \end{aligned} \tag{2.16}$$

The formulation above offers an advantage over the one presented in (2.14), as it has the potential to result in fewer constraints. This reduction in constraints can lead to faster computation of the optimal solution, which is crucial in real-time deployments, such as those discussed in [105].

Robust LP problems with an ellipsoidal uncertainty set also have a tractable formulation. Consider the ellipsoids $\mathcal{U}_j = \left\{\mathbf{u} \in \mathbb{R}^N \mid \hat{\mathbf{u}} + \mathbf{V}_j\mathbf{u}, \ ||\mathbf{u}||_2 \leq 1\right\}, \ j = 1,\dots,J$. Here, the radius of the ellipsoid is specified by the scaling of $\mathbf{V}_j \in \mathbb{R}^{N\times N}$. It can be proved that the inner maximization problem has an explicit solution [12], and thus, the optimization problem can be rewritten as the following SOCP:

$$\begin{aligned} \min_{\mathbf{x}} \ & \mathbf{c}^T\mathbf{x} \\ \text{s.t.} \ & \hat{\mathbf{u}}^T\mathbf{x} + ||{\mathbf{V}_j}^T\mathbf{x}||_2 \leq b_j, \ \forall j = 1,\dots,J. \end{aligned} \tag{2.17}$$

The problem above can be efficiently solved via Interior-Point methods.

Classes of problems other than LPs can also have tractable robust counterparts. For example, robust QCQP and SOCP problems with simple ellipsoidal uncertainty, i.e., whose uncertainty set is a single ellipsoid, are equivalent to a semidefinite optimization problem [14]. However, if polytopic uncertainty sets are considered instead, the robust counterpart is not easily solved (NP-hard) [11].

2.4 Summary

In this chapter, we provided an overview of robust optimization concepts and techniques that are crucial for the methods presented in this thesis. Robust optimization is an alternative approach to addressing model uncertainty, which, in contrast to scenario-based methods, provides a deterministic representation of the uncertain problem by considering all possible realizations of the model's uncertain variables. The uncertainty set, which is defined as the set of all possible realizations of the uncertain variables, is a key design parameter for the overall solution performance.

We presented the definition of some of the most common classes of uncertainty sets, namely polytopes and ellipsoids, which are used to represent uncertainty in the form of bounds or regions. We also discussed methods used in literature to reduce conservatism in the solution space, such as using PCA to specify a smaller uncertainty set.

We have also demonstrated that while robust LPs with polytopic and ellipsoidal uncertainty sets must satisfy infinitely many constraints, they still have tractable formulations. Specifically, robust LP models with polytopic uncertainty sets are equivalent to an LP problem, while those with ellipsoidal uncertainty sets can be formulated as a SOCP.

In the following three chapters, we present three applications of robust optimization in power and communication systems. In Chapter 3, a robust MILP with a polytopic uncertainty set is solved to determine the optimal redispatch of a power transmission grid and the location and control policies of PSTs. Because of the strong correlation of variable renewable generation of the same energy source, we use PCA to specify a much smaller uncertainty set. Chapter 4 models the routing problem in communication networks as a QCQP. A hyperrectangle encloses the uncertainty, and an analogous formulation to (2.16) is derived. Lastly, Chapter 5 uses robust optimization to determine the optimal placement of logical functions in a Smart Grid communication infrastructure with uncertain, discrete failure modes. In this case, the uncertainty set of the robust optimization problem is discrete, which means that it can be solved by adding a set of constraints for each failure scenario.

3 Adaptive Power Flows and Robust PST Placement

Chapter 3 of this thesis presents a new algorithm to determine the minimum number and location of PSTs and the optimal redispatch policy such that the power system can operate robustly for any realization of the (active) power set points from a known, continuous uncertainty set. In Section 3.1, we provide an overview of the FACTS placement problem and present in detail the advantage of our robust approach in comparison to existing work. In Section 3.2, we present the employed power system model and, in Section 3.3, we define the robust PST placement and redispatch problem. The solution algorithm is presented in Section 3.4 and is formulated as a robust MILP model using the robust reformulations of Chapter 2. In Section 3.5, we also present a hill-climbing like algorithm to enable fast computation of a sub-optimal solution of the placement problem, which is suitable for initial power grid analyses. The proposed algorithm is evaluated in Section 3.6 for a small demonstrative 3-bus example and the IEEE 39 bus test system. Table 3.1 shows the main symbols used to describe this chapter's main optimization problem presented in Section 3.4.

3.1 Overview of the FACTS Placement Problem

Variable renewable energies from wind and sun are important to reduce the global carbon footprint of energy systems. At the same time, their variability creates novel challenges for both the design and the operation of transmission grids. Especially when large renewables capacities are localized in few regions with good natural resources their power generation may lead to line overloads in the grid. Redispatch is a common measure to counteract these congestions, but it may be expensive. Grid capacity extensions are a last resort, but typically require long time to implement and are subject to societal resistances. Thus, utilizing active power flow control technologies in the grid is an attractive and often cost-effective option for grid operators. FACTS utilize power electronics to enhance the

Table 3.1: List of symbols of main optimization problem of Chapter 3.

Parameters	
$N \in \mathbb{Z}_+$	Number of buses
$L \in \mathbb{Z}_+$	Number of lines
$P \in \mathbb{Z}_+$	Number of generators/loads
$\mathbf{A} \in \mathbb{R}^{L \times N}$	Incidence matrix of power grid
$\mathbf{B} \in \mathbb{R}^{L \times L}$	Diagonal matrix with power grid's line susceptances
$\mathbf{C} \in \mathbb{R}^{N \times P}$	Matrix mapping generators/loads to buses
$\mu \in \mathbb{R}_+$	PST installation cost relative to redispatch cost
$\mathbf{c} \in \mathbb{R}_+^L$	Maximum transport capacity of the power lines
$\overline{\phi} \in \mathbb{R}_+$	Maximum phase shift of a PST
$\overline{\mathbf{H}} \in \mathbb{R}^{P \times P}$, $\overline{\mathbf{h}} \in \mathbb{R}^P$	Matrix and vector defining upper limitations of power injections
$\underline{\mathbf{H}} \in \mathbb{R}^{P \times P}$, $\underline{\mathbf{h}} \in \mathbb{R}^P$	Matrix and vector defining lower limitations of power injections
$\boldsymbol{\kappa} \in \mathbb{R}_+^P$	Unit costs of generators and loads
$M \in \mathbb{Z}_+$	Number of halfspaces whose intersection defines $\mathcal{X}$
$\mathbf{D} \in \mathbb{R}^{M \times P}$, $\mathbf{b} \in \mathbb{R}^M$	Matrix and vector of exterior representation of $\mathcal{X}$
Sets	
$\mathcal{X} \subset \mathbb{R}^P$	Uncertainty set of generators/loads' set points
Continuous Variables	
$\mathbf{p}^F \in \mathbb{R}^L$	Power line flows
$\mathbf{p}^I \in \mathbb{R}^N$	Power injections
$\boldsymbol{\theta} \in \mathbb{R}^N$	Voltage phase angles
$\boldsymbol{\phi} \in \mathbb{R}^L$	Phase shifts added by PSTs
$\mathbf{x} \in \mathcal{X}$	Power set points
$\mathbf{y} \in \mathbb{R}^P$	Redispatch power injections
$\mathbf{S} \in \mathbb{R}^{L \times P}$, $\mathbf{w} \in \mathbb{R}^L$	Matrix and vector of phase shift affine control law
$\mathbf{T} \in \mathbb{R}^{P \times P}$, $\mathbf{q} \in \mathbb{R}^P$	Matrix and vector of redispatch affine control law
$\gamma \in \mathbb{R}$	Worst-case redispatch cost
Integer Variables	
$\mathbf{z} \in \{0,1\}^L$	Indicator of PST placement

grid's controllability [35] and are an integral part of various grid development plans, e.g.,

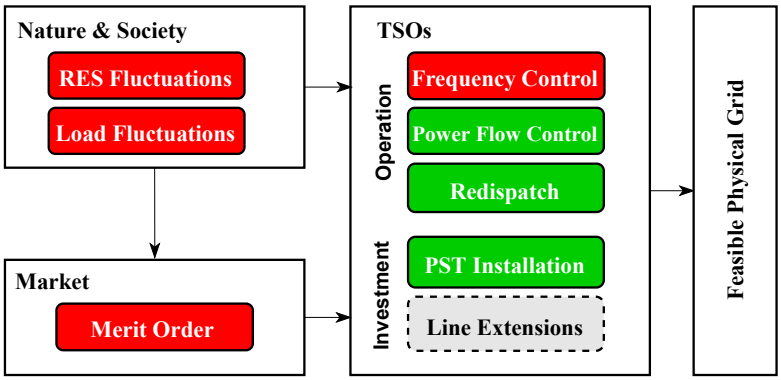

Figure 3.1: We propose an efficient planning algorithm to improve power flow controllability of transmission grids taking into account the various operational options of the TSO (green). The approach is robust against key influencing factors (red).

in Germany [1].

3.1.1 Existing Approaches

An important question that naturally arises is where to place such FACTS devices in the power grid and how to control them. Maximizing the loadability of the system for one demand scenario considering various FACTS devices can be solved by genetic algorithms [49] or by ADMM [33]. Mixed-integer linear programming can be used for PST [79] or SVC [88] placement in the same one scenario setting. Several scenarios are considered in approaches based on two-stage stochastic programming [122, 128, 46] which allows to account for the expected scenario-dependent generation costs as well. All of these approaches do not guarantee system feasibility or the (re-)dispatch costs for situations outside the (typically few) considered scenarios. In our experimental section we show that relying only on extreme generation situations may be misleading about the grid's feasibility for intermediate in-feeds.

[1] www.netzentwicklungsplan.de

3.1.2 Contributions

The main contribution of this chapter is an algorithm to find the minimal number and location of PSTs that can guarantee system feasibility and limit redispatch costs for all scenarios in a continuous uncertainty set. Considered uncertainty factors are renewable and load fluctuations as well as the corresponding dispatch decisions determined by markets or load-frequency control, see Fig. 3.1. During PST placement we take into account the grid operators' situation-dependent control options, such as specifying PST settings and redispatch orders to the attached power plants. Our algorithm thus produces as a byproduct of the PST placement task control policies to dynamically determine the PST settings and the redispatch orders. These policies are assumed as affine linear in the uncertain factors in our work. They are chosen optimally in the sense that they minimize the worst-case redispatch cost in this class of controls. Since we consider the often strongly non-linear dispatch decisions by markets and frequency control as part of the uncertainty set, the linearity assumption for small to medium-sized redispatch corrections is plausible.

In our proposed method, we model the power flow in the grid as a set of linear equations, similar to the approach used in [103]. However, we use a robust optimization framework instead of relying on chance constraints to model the uncertainty in power injections, as in [103, 17]. This framework enables us to write the continuous part of the placement and control problem as a Linear Program (LP), making it easier to solve than the second-order cone programs derived from chance-constrained formulations. The authors of [117] also use a robust optimization approach to model reserve scheduling in power systems with highly uncertain power injections due to renewable generation. In our approach, the location of the PSTs is modeled as binary variables, thus making the robust optimal placement and redispatch problem a MILP. The uncertainty set considered in the robust optimization should be large enough to cover at least a selection of plausible scenarios and all mixtures of these. At the same time, the uncertainty set should not be too large in order to avoid excessive conservatism of the solution. We use PCA [121] based on a set of given scenarios to define the uncertainty set as a rotated hyperrectangle as in [48]. While the solution might be conservative, it is guaranteed to yield feasible system states for this continuous set.

We also propose a greedy algorithm that yields near-optimal solutions while being much faster to solve than the MILP. It iteratively checks for improvement in the objective function when adding a PST to a transmission line in a hill-climbing fashion. In each iteration, the algorithm solves the linear programming relaxation of the proposed MILP and projects the fractional solution onto the feasible set of the original problem. While the maximum number of iterations grows linearly with the number of transmission lines in the grid, a parametrized stopping condition can drastically reduce the algorithm's search space. We

demonstrate that for a realistically-sized transmission grid, the greedy algorithm is more than 40 times faster to compute than the MILP and finds a solution with an objective value 17% greater. Despite a relatively big optimality gap for a planning problem, the solution of the greedy algorithm is still useful, e.g., in preliminary grid analyses.

3.2 Power Flow Model

We use the common DC approximation [74] to linearly model the power flow in a transmission grid with N buses and L lines. Power line flows are denoted as $\mathbf{p}^F \in \mathbb{R}^L$, voltage phase angles as $\boldsymbol{\theta} \in \mathbb{R}^N$, and the phase shifts potentially added by the PSTs as $\boldsymbol{\phi} \in \mathbb{R}^L$. We then have

$$\mathbf{p}^F(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbf{B}\left(\mathbf{A}\boldsymbol{\theta} + \boldsymbol{\phi}\right), \tag{3.1}$$

where $\mathbf{A} \in \mathbb{R}^{L \times N}$ is the grid's incidence matrix and $\mathbf{B} \in \mathbb{R}^{L \times L}$ is a diagonal matrix with the line susceptances.

Let $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^P$ be the uncertain vector of power set points of the P generators/loads of the grid where $\mathcal{X}$ is the corresponding uncertainty set. Moreover, let $\mathbf{y} \in \mathbb{R}^P$ denote the vector of power adjustments caused by redispatch actions. The power injections at the buses of the grid $\mathbf{p}^I(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^N$ are then given by

$$\mathbf{p}^I(\mathbf{x}, \mathbf{y}) = \mathbf{C}(\mathbf{x} + \mathbf{y}), \tag{3.2}$$

where $\mathbf{C} \in \mathbb{R}^{N \times P}$ is the matrix that maps generators/loads to the buses at which they are connected to the grid. According to Kirchhoff's first law, the nodal power injections match the sum of the power flows on the connected lines, i.e., $\mathbf{p}^I(\boldsymbol{\theta}, \boldsymbol{\phi}) = \mathbf{A}^T\mathbf{p}^F(\boldsymbol{\theta}, \boldsymbol{\phi})$. By substituting (3.1) into the previous relationship and equalizing it with (3.2), it is possible to express the voltage phase angles of the buses $\boldsymbol{\theta}$ as

$$\boldsymbol{\theta}(\mathbf{x}, \mathbf{y}, \boldsymbol{\phi}) = \left(\mathbf{A}^T\mathbf{B}\mathbf{A}\right)^+ \left(\mathbf{C}\left(\mathbf{x} + \mathbf{y}\right) - \mathbf{A}^T\mathbf{B}\boldsymbol{\phi}\right). \tag{3.3}$$

Since a constant shift of the phase angles does not change the physical situation, $\mathbf{A}^T\mathbf{B}\mathbf{A}$ is not invertible, and we use the pseudo-inverse to obtain a minimum norm solution. Substituting (3.3) into (3.1), the power line flows $\mathbf{p}^F$ can be written as a linear function of the power set points $\mathbf{x}$, the redispatch $\mathbf{y}$ and the phase shifts added by the PSTs $\boldsymbol{\phi}$ as

$$\mathbf{p}^F(\mathbf{x}, \mathbf{y}, \boldsymbol{\phi}) = \underbrace{\mathbf{E}\mathbf{C}}_{\tilde{\mathbf{C}}}(\mathbf{x} + \mathbf{y}) + \underbrace{\left(\mathbb{I} - \mathbf{E}\mathbf{A}^T\right)\mathbf{B}}_{\tilde{\mathbf{B}}}\boldsymbol{\phi}, \tag{3.4}$$

with $\mathbf{E} = \mathbf{B}\mathbf{A}\left(\mathbf{A}^T\mathbf{B}\mathbf{A}\right)^+$. This model allows us to express line capacity constraints as a linear function of the modeled uncertainties and the control decisions.

3.3 Robust Optimization of PSTs Placement and Redispatch Policies

The optimization of the number and placement of PSTs needs to consider the control actions to be performed with them as well as the effect of redispatch measures that are important additional operational measures available to the grid operators. We model the control policies for setting the phase shifts $\boldsymbol{\phi}$ and the power adjustments $\mathbf{y}$ as affine linear functions of the system state, i.e., they depend linearly on the uncertain power set points $\mathbf{x}$. The expression for the control policies are then given by

$$\boldsymbol{\phi}(\mathbf{x}) = \mathbf{S}\mathbf{x} + \mathbf{w}, \tag{3.5}$$

$$\mathbf{y}(\mathbf{x}) = \mathbf{T}\mathbf{x} + \mathbf{q}, \tag{3.6}$$

where $\boldsymbol{\phi}(\mathbf{x})$ is parametrized by $\mathbf{S} \in \mathbb{R}^{L\times P}$ and $\mathbf{w} \in \mathbb{R}^{L}$, and $\mathbf{y}(\mathbf{x})$ by $\mathbf{T} \in \mathbb{R}^{P\times P}$ and $\mathbf{q} \in \mathbb{R}^{P}$.

The optimization task that we aim to solve is to determine the minimal set of PSTs, along with their location and control policies (3.5) and (3.6) such that the total worst-case redispatch costs and the PST installations costs are minimized while ensuring that the grid operates within its feasible region for any realization of the uncertain power set points $\mathbf{x}$. With $\mathbf{z} \in \{0,1\}^{L}$ indicating the placement of a PST at a transmission line and $\gamma \in \mathbb{R}$ being the worst-case redispatch cost, the optimization problem can then be written as

$$\begin{aligned}
\min_{\mathbf{z},\gamma,\mathbf{S},\mathbf{w},\mathbf{T},\mathbf{q}} \quad & \mu\mathbf{1}^T\mathbf{z} + \gamma \\
\text{s.t.} \quad & |\mathbf{p}^F(\mathbf{x},\mathbf{y},\boldsymbol{\phi})| \leq \mathbf{c},\ \forall \mathbf{x} \in \mathcal{X}, \\
& |\boldsymbol{\phi}(\mathbf{x})| \leq \overline{\phi}\mathbf{z},\ \forall \mathbf{x} \in \mathcal{X}, \\
& \mathbf{1}^T\mathbf{y}(\mathbf{x}) = 0,\ \forall \mathbf{x} \in \mathcal{X}, \\
& \underline{\mathbf{h}}(\mathbf{x}) \leq \mathbf{y}(\mathbf{x}) \leq \overline{\mathbf{h}}(\mathbf{x}),\ \forall \mathbf{x} \in \mathcal{X}, \\
& \boldsymbol{\kappa}^T\mathbf{y}(\mathbf{x}) \leq \gamma,\ \forall \mathbf{x} \in \mathcal{X}, \\
& (3.4),\ (3.5),\ (3.6),\ \forall \mathbf{x} \in \mathcal{X}.
\end{aligned} \tag{3.7}$$

Here, $\mu \in \mathbb{R}_+$ is a weighting factor between the cost for placing the PSTs in the grid and the worst-case redispatch cost. It should be chosen depending on the assumed probability of the worst-case redispatch situation. Vector $\mathbf{c} \in \mathbb{R}_+^L$ denotes the maximum transport capacity of the lines, and $\overline{\phi} \in \mathbb{R}_+$ is the maximum phase shift of a PST. Redispatch power adjustments $\mathbf{y}$ sum to zero and should not violate the situation-dependent physical limitations $\overline{\mathbf{h}}(\mathbf{x})$, $\underline{\mathbf{h}}(\mathbf{x})$ of generation or consumption for each element of $\mathbf{x}$. An example

of such bounds is $\underline{h}(\mathbf{x}) = -\mathbf{x}$ and $\overline{h}(\mathbf{x}) = 0$ for wind and solar power plants. We generally assume a linear relation, i.e., $\overline{\mathbf{h}}(\mathbf{x}) = \overline{\mathbf{h}} + \overline{\mathbf{H}}\mathbf{x}$ and $\underline{\mathbf{h}}(\mathbf{x}) = \underline{\mathbf{h}} + \underline{\mathbf{H}}\mathbf{x}$, where $\overline{\mathbf{H}}$ and $\underline{\mathbf{H}}$ are diagonal matrices. Moreover, $\boldsymbol{\kappa} \in \mathbb{R}^P_+$ represents the cost of increasing the power output of a generator/load by one unit.

3.4 PST Placement and Redispatch Optimization as a Robust MILP

In this section, we formulate a MILP model to solve the optimization task defined in Section 3.3. By assuming a polytopic uncertainty set $\mathcal{X}$, we are able to express the infinite number of robustness constraints as a finite set of linear constraints using the theory presented in Chapter 2.

To this end, we first use expression (3.4) and the parametrizations (3.5) and (3.6) to rewrite the constraints of (3.7) in compact matricial form, i.e.,

$$\tilde{\mathbf{A}}(\mathbf{S}, \mathbf{T})\mathbf{x} \leq \tilde{\mathbf{c}}(\mathbf{z}, \mathbf{w}, \mathbf{q}, \gamma), \ \forall \mathbf{x} \in \mathcal{X}, \tag{3.8}$$

with $\tilde{\mathbf{A}}(\mathbf{S}, \mathbf{T})$ and $\tilde{\mathbf{c}}(\mathbf{z}, \mathbf{w}, \mathbf{q}, \gamma)$ being defined as

$$\tilde{\mathbf{A}}(\mathbf{S}, \mathbf{T}) = \begin{bmatrix} \begin{bmatrix} \tilde{\mathbf{C}} \\ -\tilde{\mathbf{C}} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ -\overline{\mathbf{H}} \\ \underline{\mathbf{H}} \\ \cdot \end{bmatrix} + \begin{bmatrix} \tilde{\mathbf{B}} \\ -\tilde{\mathbf{B}} \\ \mathbb{I} \\ -\mathbb{I} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{bmatrix} \mathbf{S} + \begin{bmatrix} \tilde{\mathbf{C}} \\ -\tilde{\mathbf{C}} \\ \cdot \\ \cdot \\ \mathbf{1}^T \\ -\mathbf{1}^T \\ \mathbb{I} \\ -\mathbb{I} \\ \boldsymbol{\kappa}^T \end{bmatrix} \mathbf{T} \end{bmatrix}, \ \tilde{\mathbf{c}}(\mathbf{z}, \mathbf{w}, \mathbf{q}, \gamma) = \begin{bmatrix} \mathbf{c} - \tilde{\mathbf{B}}\mathbf{w} - \tilde{\mathbf{C}}\mathbf{q} \\ \mathbf{c} + \tilde{\mathbf{B}}\mathbf{w} + \tilde{\mathbf{C}}\mathbf{q} \\ \overline{\phi}\mathbf{z} - \mathbf{w} \\ \overline{\phi}\mathbf{z} + \mathbf{w} \\ -\mathbf{1}^T\mathbf{q} \\ \mathbf{1}^T\mathbf{q} \\ \overline{\mathbf{h}} - \mathbf{q} \\ -\underline{\mathbf{h}} + \mathbf{q} \\ \gamma - \boldsymbol{\kappa}^T\mathbf{q} \end{bmatrix}, \tag{3.10}$$

(3.9)

where $\cdot$ represents zero entries.

Next, we assume that the uncertainty set $\mathcal{X}$ of the power set points $\mathbf{x}$ is a polytope, which can be defined as the intersection of M halfspaces. That is, given the exterior representation of $\mathcal{X}$

$$\mathcal{X} = \left\{ \mathbf{x} \in \mathbb{R}^P | \ \mathbf{D}\mathbf{x} \leq \mathbf{b} \right\}, \tag{3.11}$$

with $\mathbf{D} \in \mathbb{R}^{M \times P}$ and $\mathbf{b} \in \mathbb{R}^M$, we assume that $\exists \upsilon \in \mathbb{R}_+$ such that $||\mathbf{x}||_2 \leq \upsilon$, $\forall \mathbf{x} \in \mathcal{X}$.

In our experiments, we create the uncertainty set $\mathcal{X}$ using the scenario-based approach used in [48, 107]. This method leverages PCA to reduce the conservativeness of the uncertainty set. Details on this approach can be found in Section 2.2.

Using (3.8) and (3.11), our optimization task is equivalent to the following min-max optimization problem:

$$\begin{aligned}\eta = \min_{\mathbf{z},\gamma,\mathbf{S},\mathbf{w},\mathbf{T},\mathbf{q}} \quad & \mu \mathbf{1}^T \mathbf{z} + \gamma \\ \text{s.t.} \quad & \left\{ \begin{array}{cc} \max_{\mathbf{x}} & \tilde{\mathbf{A}}_{j:}^T(\mathbf{S},\mathbf{T})\mathbf{x} \\ \text{s.t.} & \mathbf{D}\mathbf{x} \leq \mathbf{b} \end{array} \right\} \leq \tilde{c}_j(\mathbf{z},\mathbf{w},\mathbf{q},\gamma), \\ & \forall j \in 1,\ldots,K, \end{aligned} \tag{3.12}$$

where K is the number of rows of $\tilde{\mathbf{A}}(\mathbf{S},\mathbf{T})$.

As shown in Section 2.3, it is possible to transform a min-max problem into a single-stage program using duality theory if the inner maximization problem is an LP and if it is always feasible. As $\mathcal{X}$ is nonempty and bounded, the inner optimization problem is always feasible, and hence its primal and dual problems have the same objective value by strong duality. Thus, (3.12) is equivalent to

$$\begin{aligned}\eta = \min_{\mathbf{z},\gamma,\mathbf{S},\mathbf{w},\mathbf{T},\mathbf{q}} \quad & \mu \mathbf{1}^T \mathbf{z} + \gamma \\ \text{s.t.} \quad & \left\{ \begin{array}{cc} \min_{\boldsymbol{\lambda}_j} & \mathbf{b}^T \boldsymbol{\lambda}_j \\ \text{s.t.} & \mathbf{D}^T \boldsymbol{\lambda}_j = \tilde{\mathbf{A}}_{j:}(\mathbf{S},\mathbf{T}) \\ & \boldsymbol{\lambda}_j \geq 0 \end{array} \right\} \leq \tilde{c}_j(\mathbf{z},\mathbf{w},\mathbf{q},\gamma), \\ & \forall j \in 1,\ldots,K \end{aligned} \tag{3.13}$$

where $\boldsymbol{\lambda}_j \in \mathbb{R}^M$ are the dual variables of the j-th inner problem.

The optimization problem (3.13) can then be formulated as the following single-level minimization problem

$$\begin{aligned}\eta = \min_{\mathbf{z},\gamma,\mathbf{S},\mathbf{w},\mathbf{T},\mathbf{q},\boldsymbol{\lambda}} \quad & \mu \mathbf{1}^T \mathbf{z} + \gamma \\ \text{s.t.} \quad & \mathbf{b}^T \boldsymbol{\lambda}_j \leq \tilde{c}_j(\mathbf{z},\mathbf{w},\mathbf{q},\gamma),\ \forall j \in 1,\ldots,K, \\ & \mathbf{D}^T \boldsymbol{\lambda}_j = \tilde{\mathbf{A}}_{j:}(\mathbf{S},\mathbf{T}),\ \forall j \in 1,\ldots,K, \\ & \boldsymbol{\lambda}_j \geq 0,\ \forall j \in 1,\ldots,K. \end{aligned} \tag{3.14}$$

Despite MILP models being NP-complete problems, modern mathematical solvers like Gurobi [53] can efficiently handle integer constraints by using various branch-and-bound

and cutting-plane methods. However, solving times can grow significantly as the size of the problem increases, which is the case in our problem with a large number of transmission lines and buses. Therefore, in the next section, we propose a heuristic approach that can provide solutions much faster using LP approximations.

3.5 Greedy Algorithm for Solving Robust MILP

As the size of the grid increases, the runtime for solving (3.14) to optimality becomes too long, as shown in our experiments (see Section 3.6). We, thus, propose a hill-climbing-like heuristic that iteratively checks whether adding a PST to a line decreases the value of the objective function. Each iteration relies on solving the linear programming relaxation of (3.14) and projecting the fractional solution onto the feasibility set of the original MILP model. The algorithm performs at most L iterations, thus yielding polynomial time complexity. However, the solution of the proposed method is only guaranteed to be a local minimum on the solution space of the original mixed-integer problem by nature of the hill-climbing procedure.

The algorithm first determines a lower bound for the objective function of (3.14) by solving its linear programming relaxation, as any solution of the (mixed-)integer program is a feasible solution of its linear relaxation. If the optimal solution of the relaxed problem has all values of $\mathbf{z}$ as 0 or 1, it will also be the optimal solution to the MILP model. To obtain an upper bound, we can solve (3.14) with $\mathbf{z}$ fixed to the value of the solution to the relaxed problem rounded to the closest integer.

The algorithm proceeds in a hill-climbing fashion to improve the integer solution. With $\mathbf{z}^*$ being the value of $\mathbf{z}$ from the solution to the relaxed problem, the algorithm places a PST into the line corresponding to the largest element of $\mathbf{z}^*$, i.e., it fixes element $j = \arg\max_i \mathbf{z}_i^*$ of $\mathbf{z}$ to 1. If the line already has a PST, the algorithm adds a PST to the line corresponding to the next largest element of $\mathbf{z}^*$. With one of the elements of $\mathbf{z}$ fixed, the algorithm solves the relaxed problem and uses the rounded solution to update the upper bound of the objective function. The algorithm continues adding PSTs to the grid until the value of the objective function stops decreasing between iterations or if the largest element of $\mathbf{z}^*$ is smaller than some parameter $\epsilon \in [0, 1]$. The pseudocode for the proposed greedy algorithm is presented in Algorithm 1, where $\mathcal{L}$ is the set of lines with PST and $\lfloor \cdot \rceil$ is the round operator.

Algorithm 1 Greedy Algorithm for Solving Robust MILP

procedure MAIN(ϵ)
  $\eta^*, \mathbf{z}^* \leftarrow$ SOLVE($\cdot, \varnothing$)
  $\overline{\eta}^*, \cdot \leftarrow$ SOLVE($\lfloor \mathbf{z}^* \rceil, \{1, \ldots, L\}$)
  $\mathcal{L} \leftarrow \varnothing$
  while $\max_{i \notin \mathcal{L}} \mathbf{z}_i^* > \epsilon$ **and** $\mathbf{z}^* \notin \{0,1\}^L$ **do**
    $\mathcal{L} \leftarrow \mathcal{L} \cup \{\arg\max_{i \notin \mathcal{L}} \mathbf{z}_i^*\}$
    $\cdot, \hat{\mathbf{z}}^* \leftarrow$ SOLVE($\{1\}^{|\mathcal{L}|}, \mathcal{L}$)
    $\hat{\overline{\eta}}^*, \cdot \leftarrow$ SOLVE($\lfloor \hat{\mathbf{z}}^* \rceil, \{1, \ldots, L\}$)
    if $\hat{\overline{\eta}}^* \geq \overline{\eta}^*$ **then**
      return $\lfloor \mathbf{z}^* \rceil$
    end if
    $\overline{\eta}^* \leftarrow \hat{\overline{\eta}}^*, \mathbf{z}^* \leftarrow \hat{\mathbf{z}}^*$
  end while
  return $\lfloor \mathbf{z}^* \rceil$
end procedure

procedure SOLVE($\mathbf{z}^*, \mathcal{Q}$)
  solve relaxation of MILP (3.14) with $\mathbf{z}_{i \in \mathcal{Q}}$ fixed to $\mathbf{z}^*$
  return value of η, value of $\mathbf{z}$
end procedure

3.6 Numerical Experiments

We demonstrate our proposed approach for two examples, a small 3-bus test grid and the IEEE 39 bus system. Our implementation uses the modeling language JuMP [34] and Gurobi [53] to solve the MILP model (3.14).

3.6.1 3-Bus Test Grid

We first apply the proposed algorithm to the grid in Fig. 3.2a. The figure also contains the values of the loads, generator limits and unit costs, as well as line capacities and susceptances. The unit cost of wind production has a negative value, mirroring precedence for renewable power generation over conventional one, as well as standard subsidy policies. Power is represented in the per unit system (p.u.) for some base power value and costs in a generic currency ¤. The example is a stylized sketch of the situation in Germany, with much wind and lignite generation in the East and the load center and many gas power plants in the Southwest [101]. Bus 3 would then represent neighboring countries.

In Fig. 3.2, we also examine the operation patterns of the test grid when varying the

available wind power from zero to its maximum capacity. Both wind sites are scaled proportionally. For each wind in-feed we determine the cost optimal dispatch using the merit order algorithm. Demands that cannot be covered from wind alone are supplied by the lignite power plant first and then by the gas power plant. For each scenario we then obtain a cost-optimal redispatch by solving (3.7).

As can be seen in Fig. 3.2b and Fig. 3.2c, no redispatch is needed when the wind is minimal or maximal: the merit-order dispatch is directly feasible. With no wind, the generation pattern is 2 p.u. at bus 1 and 2 p.u. at bus 2. With maximum wind, the generation is 3 p.u. at bus 2 and 1 p.u. at bus 3. A grid analysis based only on these two scenarios could conclude that the grid can be feasibly operated for all possible wind outputs without additional measures such as PST placement or redispatch. However, this would be a wrong conclusion since for 2 p.u. of wind, the generation pattern is 3.5 p.u. at bus 2 and 0.5 p.u. at bus 3, which leads to a capacity violation on line (2,3). This example thus proves that relying on few scenarios, as it is often done for two-stage stochastic programming approaches due to computational limitations, is not enough to ensure universal grid feasibility. In contrast, our robust approach considers a continuous uncertainty with all (convex) scenario mixtures included. While for this low-dimensional example one can obviously include average scenarios into a scenario-based analysis, this would be much more difficult in more complex, high-dimensional setups where the number of required scenarios could be huge. Note that our approach scales only in the dimension of the uncertainties but not in the number of scenarios, which can grow exponentially with the number of uncertainty dimensions if an even cover of the scenario space is desired.

In Fig. 3.2b and Fig. 3.2c, we also show the redispatch policies computed by solving (3.14). For Fig. 3.2b we choose μ such that no PST is installed, but higher redispatch costs are preferable, whereas in Fig. 3.2c the parameter μ is chosen small enough such that one PST on line (1,3) is optimal. Note that the derived values of the computed redispatch policy are not linear in the wind power production. However, they are affine linear in the dispatch including the non-linear merit order results. For both cases, with and without PST, the implied worst-case cost by the computed affine linear policy matches the worst-case cost of the scenario-wise optimized redispatch schedule, as shown in Fig. 3.2d.

This example also shows that the redispatch cost can be reduced by increasing the controllability of the power flow in the grid with the addition of a PST. Fig. 3.2e shows the angle shift added by the PST following the affine control policy as well as the amount of power flow redirected from branch (2,3) to (2,1). These settings ensure a feasible grid power flow in all situations. It is worth noting that in the “no wind” situation, the PST also redirects energy from line (2,3) to line (1,2), which helps increase the distance to the feasibility border, since without PST interaction the line (2,3) would be fully loaded.

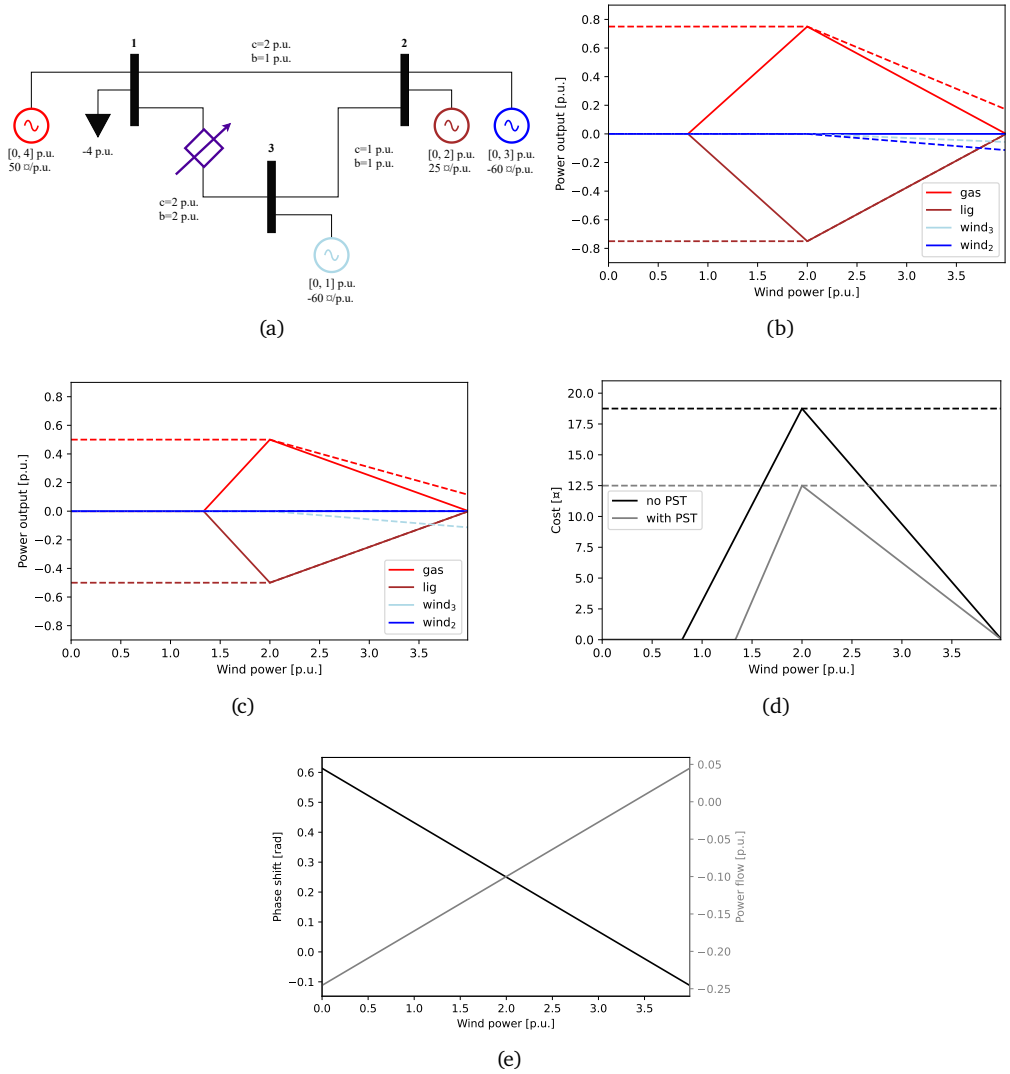

Figure 3.2: (a) 3-bus test grid. A gas power plant (red) and a load are connected to bus 1, while a lignite power plant (brown) and a wind park (blue) are connected to bus 2. A second, smaller wind park (light blue) is connected to bus 3. A PST is added to line (1,3) for a small enough installation cost μ. For varying the total wind power in-feed for the 3-bus test grid in Fig. 3.2a. (b) Scenario-wise optimized redispatch (solid) and affine-linear redispatch policy (dashed) without a PST and (c) with PST. (d) The corresponding redispatch costs. (e) Optimal affine PST policy (black) in case a PST is installed and redirected power flow (gray).

3.6.2 IEEE 39 Bus Test Case

We now examine a larger example, namely the IEEE 39 test case, which contains 39 buses, 46 transmission lines, 10 generators, and 21 loads, as shown in Fig. 3.3a. The properties of the transmission lines, generators, and loads are based on the MATPOWER test cases [131]. To generate diverse usage scenarios we proceed as follows. We start with an hourly load profile curve aggregated for a single country. Here we use the Danish load from Sept. 28$^{\text{th}}$ of 2020 and Sept. 27$^{\text{th}}$ of 2021 to generate 8735 scenarios [37]. To obtain nodal load time series that are both realistically correlated but are also partly independent of each other, we normalize the given load curve by its maximum value to yield $\ell(t)$, $t = \{1, \ldots, 8735\}$, and add randomness as follows:

$$d_i(t) = d_i^N \left(\ell(t) + u_i(t)\right), \tag{3.15}$$

where d_i^N is the nominal value given in the IEEE test case for load i, and $u_i(t)$ is a sample of a uniform distribution with probability density between -0.2 and 0.2. Merit order is used for the dispatch of the generators where the dispatch cost of a generator is taken proportional to the inverse of its maximum active power output.

In Fig. 3.3b, we show the dispatch of the generators for the first 200 consecutive time steps. The system without redispatch or PSTs is infeasible for the generated scenarios in the sense that the power flow over some transmission lines is larger than their capacity. We show in Fig. 3.3a the transmission lines with the largest violations in red.

In Fig. 3.3c, we show the power output of the generators when following the computed redispatch policy without any PST. This is achieved for very high PST installation costs in the MILP (3.14) problem. The optimal redispatch policy makes the power flow feasible for all scenarios, and the worst-case redispatch cost is 3.30 ¤. The cheap generator G1 produces less power than is economically optimal without considering grid constraints.

In Fig. 3.3d, we solve (3.14), now lowering the PST installation costs significantly. With these settings, the algorithm proposes to install one PST between buses 1 and 39 as shown in Fig. 3.3a. Given this decision, it is possible to reduce the worst-case redispatch cost to 1.93 ¤ by using more power from the cheaper generator G1.

Optimality versus computation time is examined in Table 3.2, comparing the MILP and the greedy algorithm proposed in Section 3.5 for the same setting. The greedy algorithm is run with two different values of ϵ. With ϵ equal to 0.2, the algorithm does not propose the installation of any PST. If the value of ϵ is decreased to 0.06, the algorithm finds a better solution that yields 2.25 ¤ for the worst-case redispatch cost, an increase of 17% compared to the optimal solution of the MILP model. In this solution, a PST is installed between buses 2 and 3.

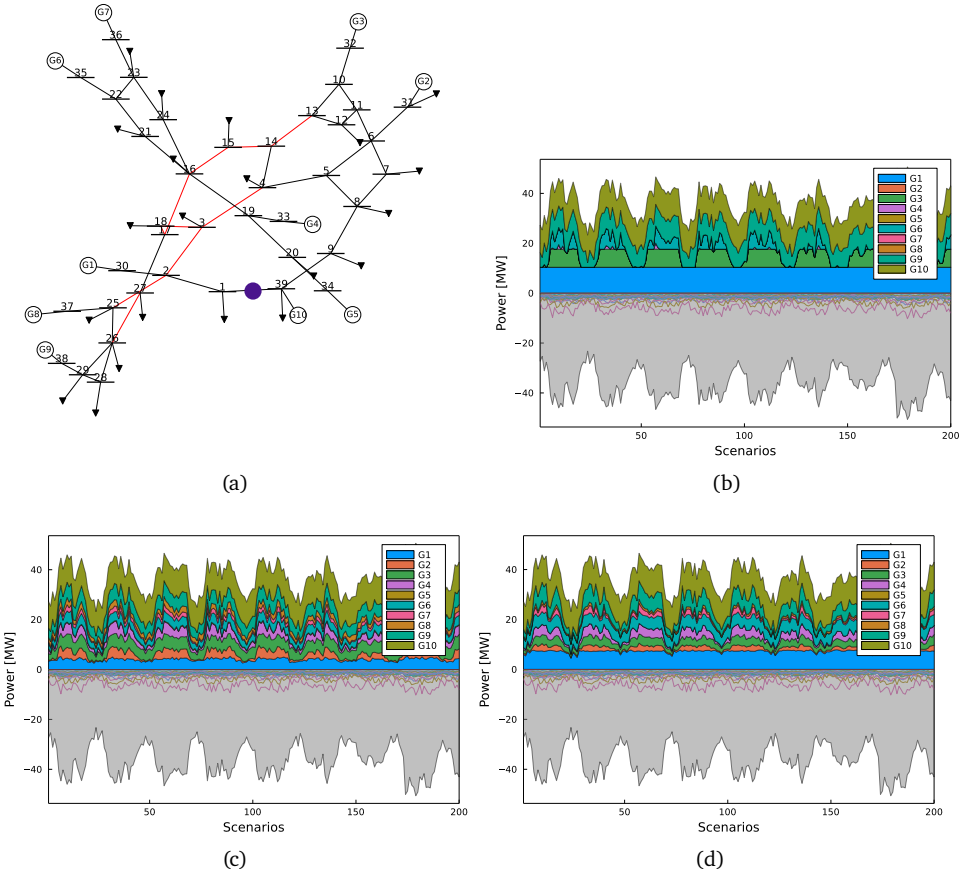

Figure 3.3: (a) Topology of IEEE 39 bus test case. The lines marked in red have the largest capacity violations when solving the power flow without redispatch or PSTs. The shaded circle denotes the location proposed by our algorithm for installing a PST. Power output of the ten generators of the IEEE 39 system for 200 consecutive time steps. The gray area represents the total power demand. (b) The power output is defined by merit order where the production cost is inversely proportional to the generator maximum active power output. (c) The power output is defined by the dispatch plus the redispatch generated by the optimal redispatch policy found using the proposed algorithm. (d) Same as (c) but with a PST installed between buses 1 and 39.

Table 3.2: Performance evaluation of MILP and greedy algorithm $\mathcal{G}(\epsilon)$.

	cost [¤]	**time** [min]	**# PSTs**
MILP	1.93	348	1
$\mathcal{G}(0.2)$	3.30	3	0
$\mathcal{G}(0.06)$	2.25	8	1

Using a local compute node with 22 cores, MILP optimality with a 10% gap was proven after 5 hours and 48 min. While acceptable for grid planning purposes, the computation time is relatively high. In the configuration with ϵ equal to 0.2, the computation time of the greedy algorithm is 116 times faster than solving the MILP model, and it is 43.5 times faster for ϵ equal to 0.06. One could further improve the solving time of both the mixed-integer problem and the proposed algorithm by, e.g., restricting the affine linear maps to depend on principal components of the in-feeds only, thereby reducing the dimensions of $\mathbf{S}$ and $\mathbf{T}$.

The results show that the hill-climbing approach could be a valuable tool in preliminary analyses of PST placement where sub-optimal solutions are acceptable. Then, once the refined model is found, one can solve the MILP model to find the optimal placement.

3.7 Summary

This chapter presented a novel algorithm to find the minimal set of PSTs and the optimal phase shift and redispatch policies that enable the feasible operation of the transmission system given uncertain in-feed scenarios. The algorithm leverages the robust optimization framework presented in Chapter 2 to guarantee feasibility for all combinations of scenarios within a continuous uncertainty set.

Using duality theory, we formulated the robust problem as a MILP model with a finite number of constraints, which can be solved by commercial solvers. Additionally, we developed a greedy algorithm based on the hill-climbing optimization approach to speed-up computations, especially when solving the placement problem for large power systems. While the proposed heuristic provides sub-optimal solutions, one could use this heuristic for preliminary analyses.

We demonstrated our algorithms with two examples proving that only considering a few extreme scenarios may not be sufficient to guarantee feasible grid states under all conditions and that the PST placement can significantly reduce the worst-case redispatch cost.

4 Adaptive Robust Routing in Communication Networks

This chapter presents a novel robust formulation for the optimal routing of data packets in backbone communication networks. After an introduction to the problem setting (Section 4.1), the chapter presents in Section 4.2 a continuous-time model of a communication network, which is based on the fluid flow model [109]. The optimal routing problem is then formulated as a robust optimization problem that must be solved in real-time whenever the uncertain data flows are close to the boundary of the uncertainty set in Section 4.3. We utilize the techniques presented in Chapter 2 to derive a tractable representation of the robust optimization problem, which takes the form of a QCQP, in Section 4.4. An iterative heuristic to solve the robust QCQP is presented in Section 4.5. This algorithm makes the approach practical and applicable to real-world deployment, where the routing policies must be computed within just a few minutes. As shown in Section 4.6, this heuristic enables the efficient solution of the robust optimization problem, even for large numbers of data flows. To evaluate the proposed approach, we compare its performance against standard benchmarks and assess its scalability using an example based on an actual data center backbone network. The results demonstrate the effectiveness of the novel robust QCQP formulation for the optimal routing of data packets in backbone networks and its potential for practical implementation in real-world communication systems. Table 4.1 shows the main symbols used to describe this chapter's main optimization problem presented in Section 4.4.

4.1 Problem Setting

In communication networks, data is transmitted in the form of packets. These packets not only contain the actual information to be conveyed but also other crucial data such as the originator and recipient addresses. Large amounts of information are divided into multiple packets, creating data flows. These data flows must traverse various links and networks to

Table 4.1: List of symbols of main optimization problem of Chapter 4.

Parameters	
$N \in \mathbb{Z}_+$	Number of routers
$L \in \mathbb{Z}_+$	Number of communication links
$\mathbf{A} \in \mathbb{R}^{N^2 \times NL}$	Connectivity matrix indicating outgoing flows
$\mathbf{B} \in \mathbb{R}^{N^2 \times NL}$	Connectivity matrix indicating incoming flows
$\mathbf{C} \in \mathbb{R}^{N^2 \times N^2}$	Connectivity matrix indicating routers with external flows
$\mathbf{S} \in \mathbb{R}^{NL \times N^2}$	Selector matrix embedding network's locality constraints
$\mathbf{D} \in \mathbb{R}^{L(N+1) \times NL}$	Connectivity matrix mapping flows to physical links
$\mathbf{c} \in \mathbb{R}^{L(N+1)}$	Maximum transport capacity of links
$\boldsymbol{\delta} \in \mathbb{R}^{L}$	Average link delays
$\overline{\mathbf{r}} \in \mathbb{R}^{N^2}$	Upper limit of incoming data rates
$\underline{\mathbf{r}} \in \mathbb{R}^{N^2}$	Lower limit of incoming data rates
Sets	
$\mathcal{U} \subset \mathbb{R}^{N^2}$	Uncertainty set of incoming data rates
Continuous Variables	
$\mathbf{v} \in \mathbb{R}^{NL}$	Data rates of flows within the network
$\mathbf{r} \in \mathbb{R}^{N^2}$	Incoming data rates
$\mathbf{G} \in \mathbb{R}^{NL \times N^2}$	Matrix defining local decision rules
$\eta \in \mathbb{R}$	Total network delay

reach their final destination. The route taken by these data flows is determined by routers, which are placed at the intersections of different links. The routing policies stored in the router hardware aim to deliver packets to their recipients while maximizing a utility function, such as the total network data throughput.

In the past, routing primarily operated in a fully-distributed manner, and policies remained unchanged for extended periods of time. This was because every adaptation necessitated routers to engage in numerous information exchange rounds to achieve convergence. However, with the widespread use of Internet-connected mobile devices and on-demand Internet services, static policies are no longer sufficient. Data flows are subject to fluctuations based on end-user preferences and the growing number of mobile devices further increases the uncertainty in data traffic. Thus, characterizing Internet traffic has been a crucial step for enabling more efficient, reliable and faster communication

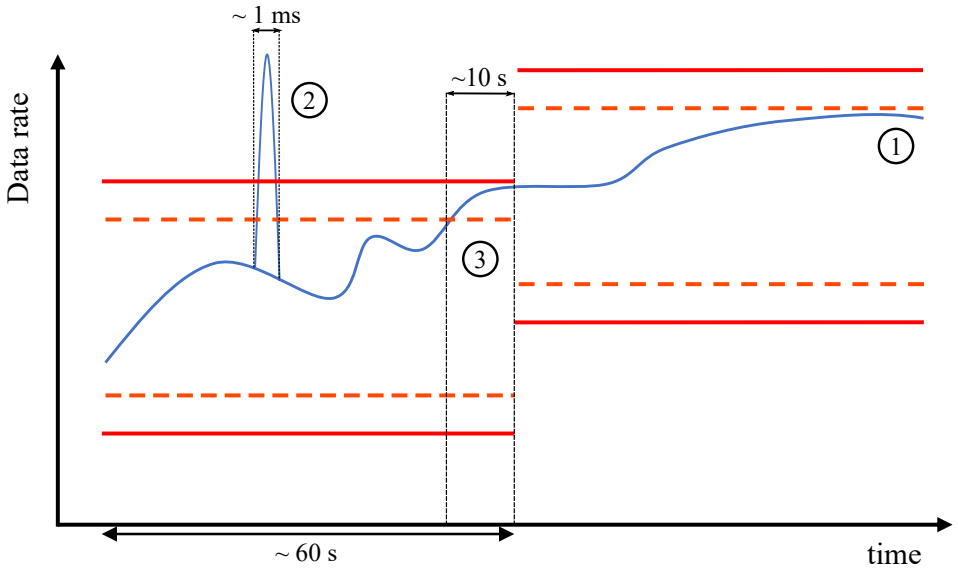

Figure 4.1: Data flow rates (blue) are assumed to contain slowly varying trends (1: variation in tens of s) as well as fast bursts (2: variation in few ms). Bursts are handled by router buffers while slower variations are handled by local control policies that guarantee global flow feasibility within bandwidth limits (red, solid). If the warning level (orange, dashed) is exceeded, a global re-coordination with an adapted uncertainty set is triggered, and the new controller gains are applied soon afterward (3).

networks.

It is well-known that the aggregated traffic at backbone networks has long-term correlations [104], while some works such as [66] show that on very fine time scales a convergence to Poisson traffic can be observed. Short-term traffic rate variations (usually in the order of ms's) can be attributed to the stochastic burstiness of TCP connections [112]. Fig. 4.1 sketches such a traffic profile in a backbone network.

To address the growing uncertainty in data traffic in communication networks, we proposed a new adaptive algorithm for data routing [105]. The mechanism consists of local routing policies that route data packets at line rate and a global coordination step that optimizes the local policies based on the long-term characterization of the data flows. With this two-level scheme, the algorithm ensures optimal routing while being able to adapt quickly to changes in data flows.

4.1.1 Existing Approaches

A part of the existing adaptive routing algorithms relies on a global view of the network, using different optimization techniques to optimize local routing policies. For example, [115, 44, 118] formulated the optimal routing problem as a numerical network utility optimization problem. On the other hand, authors in [97, 43, 63] made use of heuristics and demonstrated the effectiveness of their methods experimentally. However, one of the main drawbacks of global solutions is their inability to adapt to rapidly changing data traffic. Additionally, the overhead of monitoring the entire network and aggregating information at a single control center can be significant [55].

In contrast, other methods utilize a distributed optimization strategy, where each router can adapt based on local information, leading to quicker adaptation. These approaches make decisions based on small units of locally available data, such as single packets [30, 50] or packet bursts known as flowlets [6, 67]. Yet, these local policies use no or only limited information about the global network state and, hence, may lead to sub-optimal behavior.

Few works have explored a hybrid approach to optimize data flow routing, combining global optimization with adaptive local policies to achieve the global optimum with small adaptation latency. The authors of [1] use a robust approach based on the $\mathcal{H}_\infty$-control framework to devise parametrized local policies that are robust against varying communication delays and which ensures the network's physical constraints, such as link bandwidth. However, this approach requires the linear bandwidth constraints to be replaced by linear matrix inequality conditions that are sufficient but not necessary for many realizations of the data flows, which may lead to overly conservative solutions. Additionally, conservatism is further increased by not considering adaptive (global) optimization, which mandates using large uncertainty sets.

Reduced conservatism in routing can be accomplished through the use of different control strategies, such as the lookahead control scheme [75]. This approach involves solving a global optimization problem at each time step, considering not only the current data flow rates but also their future changes over an observation window. The authors of [58] use this method to determine the optimal placement of SDN controllers in a backbone network to minimize operation costs and the drawbacks of recurrent reconfigurations. However, this method assumes precise traffic prediction over all time steps, which is challenging in practice.

4.1.2 Contributions

The main contribution of the proposed method in [105] is a flow-control scheme specifically designed to meet the demands and constraints of backbone communication networks. The scheme consists of two parts: a global optimization step and local flow-controlling policies. The global optimization step is a robust optimization problem whose goal is to minimize a certain metric given the physical constraints of the system, such as the bandwidth of the communication links. The solution to the problem is robust against the long-term variations of the data flows, which are assumed to be contained by a hyperrectangular uncertainty set. Short-term fluctuations in the range of ms's are assumed to be absorbed by the buffers of the routers' interfaces and, thus, are not considered in the controller design. Indeed, for a typical $100\ \mathrm{Gbps}$ aggregate traffic rate, a rate increase of 10% amounts to $10\ \mathrm{Mb}$ of additional data in $1\ \mathrm{ms}$. Therefore, bursts of few ms's can be absorbed with plausible buffer sizes in the range of a few Mb's.

The robust counterpart of the global robust optimization problem is a (non-convex) QCQP, which can be solved by standard solvers and yields a set of local routing policies implemented. These policies are proportional controllers dependent solely on the current rate of incoming data flows. This type of controller complies with the capabilities of modern programmable routers and can, thus, be deployed in real-world networks.

In addition, the chapter presents an adaptation mechanism to trigger the re-optimization of the local routing policies based on the proximity of the instantaneous data rates to the boundary of the uncertainty set from the global robust optimization problem. This mechanism balances the need for reliable operation with the reduction of conservatism, as local policies must ensure feasible operation only for variations in the data rates within a small observation window, typically a few minutes, rather than all possible rate values.

As the number of flows and network size increases, solving the global robust optimization problem becomes increasingly challenging due to the absence of a polynomial-time algorithm to solve non-convex QCQPs. To overcome this, the authors have developed a heuristic approach that solves the global optimization problem iteratively, relaxing the quadratic constraints to linear constraints until convergence is achieved. This approach provides close-to-optimal results with minimal bandwidth violations in our experiments.

4.2 Data Flow Model

The network model used throughout this chapter is based on the so-called fluid-flow model [109]. In this model, the flows of data packets are represented as continuous flows, characterized by their source, sink, and information transfer rate measured in

bits per second. Given a communication network with N routers and L links, let $v_{ij}^k(t)$ be the instantaneous input rate at time t of the aggregated data flow arriving at the output interface of link $(i,j) \in \mathcal{L}$ whose destination is router $k \in \mathcal{N}$, and $y_{ij}^k(t)$ the output rate packets are transmitted over link (i,j). Assuming output buffering, i.e., packets are stored in buffers at the output interface before transmission, the amount of data $q_{ij}^k(t)$ in interface buffer (i,j) headed for destination node k is governed by the following differential equation

$$\dot{q}_{ij}^k(t) = v_{ij}^k(t) - y_{ij}^k(t) - d_{ij}^k(t), \tag{4.1}$$

where $d_{ij}^k(t)$ is the packet drop rate at the interface.

The actual transmission rate $y_{ij}^k(t)$ is determined by the buffer occupancy and the maximum rate at which data can be processed, represented by σ_{ij}, using the following equation:

$$y_{ij}^k(t) = \sigma_{ij} \frac{q_{ij}^k(t)}{\sum_{l \in \mathcal{N}} q_{ij}^l(t) + \epsilon}, \tag{4.2}$$

where ϵ is a small positive real constant added to prevent division by zero in the event that all buffers are empty. This equation assumes equal priority for all traffic over the link (i,j).

If the interface buffer becomes full, packets will be dropped. The rate at which packets are dropped, given the buffer size λ_{ij}, is given by

$$d_{ij}^k(t) = \max(v_{ij}^k(t) - y_{ji}^k(t), 0)\, \mathbb{1}_{\sum_{k \in \mathcal{N}} q_{ij}^k(t) \geq \lambda_{ij}}. \tag{4.3}$$

The routing controller determines the forwarding rates $v_{ij}^k(t)$ based on measurements of the transmitted rates $y_{ij}^k(t)$, the buffer states $q_{ij}^k(t)$, and the rates $r_i^k(t)$ of data entering the network at node i for destination k. To ensure consistency, the following condition must hold for all nodes $i \in \mathcal{N}$ except the sink node k:

$$\sum_{j \in \mathcal{A}_i} v_{ij}^k(t) = \sum_{j \in \mathcal{B}_i} y_{ji}^k(t) + r_i^k(t), \tag{4.4}$$

with $\mathcal{A}_i,\ \mathcal{B}_i \subset \mathcal{N}$ being the set of downstream and upstream neighbors of i, respectively.

4.3 Two-level Controller for Optimal Routing

Adapting all routing decisions continuously based on the instantaneous values of all possible measurements in the network would require continued fast communication and

computation within the network's control plane, which is an inefficient and practically challenging task. Therefore, we propose a two-level hierarchical flow-controlling scheme that uses typical data traffic characteristics in backbone networks to optimize the local routing decision rules efficiently.

Local routing policies are implemented at each router as proportional controllers using locally available measurements to split the incoming data flows among the routers' neighbors. Acting in a slower time scale, a *robust global coordination* algorithm determines the gain parameters of all local controllers such that the routing leads to a globally feasible and optimal behavior. That is, the routing should guarantee compliance with the bandwidth limits of the network and a low value of metrics such as the network average delay.

A robust framework is employed for the coordination step to account for potential changes in incoming data rates while the gain parameters are active. The formulation uses a hyperrectangular uncertainty for the flow rates. Re-coordination is triggered whenever the incoming flow rates surpass a warning level, i.e., when they are too close to the boundaries of the uncertainty set. Fig. 4.1 exemplifies the process of updating the local policies.

4.3.1 Local Routing Policies

We propose a proportional controller for the actuated variables $v_{ij}^k(t)$ using the rates of incoming data flows $r_i^k(t)$ for all destinations $k \in \mathcal{N}$ as feedback signals. To reduce the effect of short-term traffic variations, we use a moving average filter $\mathcal{F}$ with a window in the range of 100 ms on the incoming rates. The filtered feedback signal $\tilde{r}_i^k(t)$ is defined as

$$\tilde{r}_i^k(t) = \mathcal{F}\left\{\sum_{j\in\mathcal{B}_i} y_{ji}^k(t) + r_i^k(t)\right\}. \tag{4.5}$$

Choosing the size of the sliding window of the filter $\mathcal{F}$ requires consideration. A smaller window provides faster adaptability, but also increases the risk of frequent triggering of the global coordination algorithm due to fast variances in the filtered signal. On the other hand, a larger window results in fewer re-coordinations, but the guarantee that buffers can absorb typical packet bursts would become invalid, potentially causing packet drops.

We assume a general proportional control scheme in which the flow rate $v_{ij}^k(t)$ is a linear combination of all incoming rates $\tilde{r}_i^{\tilde{k}}(t)$, $\forall \tilde{k} \in \mathcal{N}$. Note that the computation of the local policies is restricted to the local rate measurements in router i. This way, real-time exchange of data rate information between routers is not required, thereby reducing communication overhead. With $g_{ij}^{k\tilde{k}}(t)$ being the multiplicative factor that determines

the flow rate $v_{ij}^{k}(t)$ as a function of the incoming rate $\tilde{r}_{i}^{\tilde{k}}(t)$, the proposed proportional controller scheme can be stated as

$$v_{ij}^{k}(t) = \sum_{\tilde{k} \in \mathcal{N}} g_{ij}^{k\tilde{k}}(t)\, \tilde{r}_{i}^{\tilde{k}}(t). \tag{4.6}$$

Although having a restricted operation set, modern routers with programmable data planes can execute both the filtering operation (4.5) and the linear algebra in (4.6) at line rate [47]. Hence, we ensure with this formulation that the proposed routing scheme is implementable in a real-world setting.

4.3.2 Robust Global Coordination

The target of the global coordination algorithm is to determine suitable values for the local controller gains $g_{ij}^{k\tilde{k}}(t)$ such that the bandwidth limits in the network are not violated and that a given performance metric, e.g., the average packet delay, is minimized. We compute the controller gains centrally using the filtered value of the instantaneous traffic rates at the moment the coordination was triggered to build the robust optimization's uncertainty set. In this framework, network's capacity constraints are enforced for any realization of the uncertain incoming flow rates by construction.

For the controller design, the routed data rates $v_{ij}^{k}(t)$ and the transmission rates $y_{ij}^{k}(t)$ are considered to be identical, replacing dependency (4.2). This is because we target feasible flow policies without line overloads. Buffers will then be always empty, and no packet drop will occur since packets can be processed and retransmitted with a rate larger or equal to the incoming data rate. Note that control policies with limited, but small bandwidth violations can be designed by adding a dummy node connected to all other nodes in the network via links with unlimited capacity but with very high cost of utilization.

Let $\mathbf{v} \in \mathbb{R}^{NL}$ be the vector of time-averaged data rates of all flows within the network and $\mathbf{r} \in \mathbb{R}^{N^2}$ that of the incoming data rates, respectively. During the execution of a global coordination, we consider the flows as static given the current optimization time t and, thus, we drop time indices in the following. Condition (4.4) can then be rewritten in matrix form as

$$\mathbf{A}\mathbf{v} = \mathbf{B}\mathbf{v} + \mathbf{C}\mathbf{r}, \tag{4.7}$$

where $\mathbf{A} \in \mathbb{R}^{N^2 \times NL}$, $\mathbf{B} \in \mathbb{R}^{N^2 \times NL}$ and $\mathbf{C} \in \mathbb{R}^{N^2 \times N^2}$ are appropriate matrices with zeros and ones that express the network's connectivity.

Moreover, the local decision rules of (4.6) can be jointly represented via a gain matrix $\mathbf{G} \in \mathbb{R}^{NL \times N^2}$ that embeds all multiplicative factors $g_{ij}^{k\tilde{k}}(t)$ as

$$\mathbf{v} = \mathbf{G}\left(\mathbf{B}\mathbf{v} + \mathbf{C}\mathbf{r}\right). \tag{4.8}$$

Matrix $\mathbf{G}$ contains many zeros to express locality constraints so that traffic can only be routed to adjacent links and local decisions depend only on locally available measurements. Using a suitable selector matrix $\mathbf{S} \in \mathbb{R}^{NL \times N^2}$ with only zero and one entries, these constraints can be expressed as

$$\mathbf{S} \circ \mathbf{G} = \mathbf{0}. \tag{4.9}$$

The bandwidth limits of the network links and the non-negativity of the data flows are enforced by linear inequality constraints as

$$\mathbf{D}\mathbf{v} \leq \mathbf{c}, \tag{4.10}$$

with $\mathbf{D} \in \mathbb{R}^{L(N+1) \times NL}$ and $\mathbf{c} \in \mathbb{R}^{L(N+1)}$. The capacity constraints are represented in the first L rows of $\mathbf{D}$, which summarize all flows transmitted via a certain link. The remaining LN rows represent the non-negativity constraints.

The objective of the global coordination algorithm is to optimize a given metric such as the average packet delay, which can be written as the following linear algebra:

$$\eta = \boldsymbol{\delta}^T \mathbf{D}_{(1:L)} \mathbf{v}, \tag{4.11}$$

where $\boldsymbol{\delta} \in \mathbb{R}^L$ is the vector of link delays and $\eta \in \mathbb{R}$ is the indicator of the total network delay.

It's noteworthy to mention that the proposed method can also be applied to any network metric that can be expressed as a convex quadratic function of the actuated variable $\mathbf{v}$. This is because, as demonstrated in the following, the robust global coordination problem will be transformed into a QCQP. For example, minimizing the squared relative link utilization would favor an even distribution of the flows over the available network links.

During the time interval that a given gain matrix $\mathbf{G}$ is actively applied, the incoming data rates may fluctuate within a certain range. We assume that this range of variation can be estimated from past observations and we aim for a gain matrix $\mathbf{G}$ that keeps the network flows within the bandwidth limits for all possible values of the data rates within this range. More specifically, we define the uncertainty set $\mathcal{U}$ for the rates $\mathbf{r}$ as

$$\mathcal{U} = \left\{ \mathbf{u} \in \mathbb{R}^{N^2} \middle| \, \underline{\mathbf{r}} \leq \mathbf{u} \leq \overline{\mathbf{r}} \right\},$$

where $\underline{\mathbf{r}}$ and $\overline{\mathbf{r}}$ are the lower and upper boundaries of $\mathbf{r}$, respectively. In our implementation of the global coordination algorithm, we define the uncertainty set's boundaries

symmetrically with the traffic rate when the coordination algorithm is triggered as its center point.

Using the previous conditions, we can state the global coordination algorithm as a robust optimization problem whose solution is the optimal gain matrix $\mathbf{G}^* \in \mathbb{R}^{NL \times N^2}$ that yields the minimum total average delay $\eta^* \in \mathbb{R}$:

$$\begin{aligned} \mathbf{G}^*, \eta^* \in \underset{\mathbf{G},\eta}{\arg\min} \quad & \eta \\ \text{s.t. } & (4.9) \ \wedge \ \forall \mathbf{r} \in \mathcal{U} : \exists \mathbf{v} : \\ & (4.7),\ (4.8),\ (4.10),\ (4.11). \end{aligned} \tag{4.12}$$

In Chapter 2, it was explained that optimization problems with the form of (4.12) are not tractable as for every value of $\mathbf{r} \in \mathcal{U}$, a set of constraints must be fulfilled, which would result in an infinite number of constraints. Fortunately, the robust counterpart of (4.12) is a finite non-convex QCQP, which can be solved by existing optimization tools. In the following, we demonstrate how to reformulate (4.12) into a tractable optimization problem.

4.4 Global Coordination as a Robust QCQP

The global coordination problem (4.12) is a robust optimization problem with an included linear response policy. We reformulate the problem constraints such that they do not depend explicitly on the realizations of $\mathbf{r}$ anymore so we can use the technique shown in Section 2.3 to upper bound the problem's constraints by the extreme values of the set $\mathcal{U}$. In this formulation, the number of constraints grows linearly with the dimension of $\mathbf{r}$.

To begin, we rearrange the local decision rules (4.8) by introducing an auxiliary variable $\tilde{\mathbf{G}} \in \mathbb{R}^{NL \times N^2}$. This variable enables to express the rates of internal flows $\mathbf{v}$ as a linear combination of the uncertain incoming flow rates $\mathbf{r}$:

$$\mathbf{v} = \tilde{\mathbf{G}}\mathbf{r}. \tag{4.13}$$

Assuming $(\mathbb{I} - \mathbf{GB})$ is invertible, it follows from (4.8) and (4.13) that $\tilde{\mathbf{G}}$ can be determined from $\mathbf{G}$ via the bilinear equation

$$(\mathbb{I} - \mathbf{GB})\tilde{\mathbf{G}} = \mathbf{GC}. \tag{4.14}$$

Next we show that traffic conservation can be expressed as a linear constraint on $\tilde{\mathbf{G}}$ that is independent of $\mathbf{v}$ and $\mathbf{r}$ as

$$(\mathbf{A} - \mathbf{B})\,\tilde{\mathbf{G}} - \mathbf{C} = \mathbf{0}. \tag{4.15}$$

Without loss of generality, let $\mathbf{r} = \underline{\mathbf{r}} + \Delta\mathbf{r}$, where $\Delta\mathbf{r} \in \tilde{\mathcal{U}}$, $\tilde{\mathcal{U}} = \{\tilde{\mathbf{u}} \in \mathbb{R}^{N^2} | \ \mathbf{0} \leq \tilde{\mathbf{u}} \leq \overline{\mathbf{r}} - \underline{\mathbf{r}}\}$. By combining (4.7) and (4.13), we have that

$$(\mathbf{A}\tilde{\mathbf{G}} - \mathbf{B}\tilde{\mathbf{G}} - \mathbf{C})(\underline{\mathbf{r}} + \Delta\mathbf{r}) = \mathbf{0}.$$

Since this equation must hold for all $\Delta\mathbf{r} \in \tilde{\mathcal{U}}$, it must also hold for $\Delta\mathbf{r} = \mathbf{0}$, and, thus,

$$(\mathbf{A}\tilde{\mathbf{G}} - \mathbf{B}\tilde{\mathbf{G}} - \mathbf{C})\underline{\mathbf{r}} = \mathbf{0}.$$

Assuming that $\tilde{\mathcal{R}}$ is not degenerate, i.e., $\overline{\mathbf{r}} > \underline{\mathbf{r}}$, we can conclude that, if $(\mathbf{A}\tilde{\mathbf{G}} - \mathbf{B}\tilde{\mathbf{G}} - \mathbf{C})\Delta\mathbf{r} = \mathbf{0}, \ \forall\Delta\mathbf{r} \in \tilde{\mathcal{U}}$, then (4.15) must hold.

The capacity constraint (4.10) can be tightly upper bounded using an auxiliary matrix $\mathbf{H} \in \mathbb{R}^{NL \times N^2}$ whose entries fulfill

$$\begin{aligned} \mathbf{H} &\geq \mathbf{D}\tilde{\mathbf{G}}\text{diag}\left(\underline{\mathbf{r}}\right), \\ \mathbf{H} &\geq \mathbf{D}\tilde{\mathbf{G}}\text{diag}\left(\overline{\mathbf{r}}\right). \end{aligned} \tag{4.16}$$

Then, the capacity constraints can be rewritten without depending on the single realizations of $\mathbf{r}$ as

$$\mathbf{H}\mathbf{1} \leq \mathbf{c}. \tag{4.17}$$

The total average delay (4.11) can be similarly upper bounded with variable $\mathbf{H}$:

$$\eta = \boldsymbol{\delta}^T \mathbf{H}_{(1:L)}\mathbf{1}. \tag{4.18}$$

The above steps allow us to write the global coordination problem (4.12) as a finite-dimensional non-convex QCQP:

$$\begin{aligned} \mathbf{G}^*, \tilde{\mathbf{G}}^*, \eta^*, \mathbf{H}^* \in \underset{\mathbf{G}, \tilde{\mathbf{G}}, \eta, \mathbf{H}}{\arg\min} \quad & \eta \\ \text{s.t.} \quad & (4.9), \ (4.14) - (4.18). \end{aligned} \tag{4.19}$$

QCQPs like (4.19) can be solved by commercial solvers, e.g., Gurobi [53]. Despite the problem being non-convex, Gurobi is able to solve it to global optimality by converting the non-convex quadratic constraints into bilinear ones and applying spatial branching (branching on continuous variables) and adaptive McCormick relaxations.

However, because (4.19) is non-convex, solving it for large networks with tens of routers and hundreds of flows within acceptable time is challenging even for state-of-the-art solvers. Therefore, in the following, we present an iterative method for solving (4.19) that improves its tractability for large networks. In each iteration, the algorithm solves a convex relaxation of (4.19), which can be solved in polynomial time [73].

4.5 Iterative Method for solving QCQP

As the network size increases the solution time of the proposed QCQP (4.19) becomes impractically high. We thus propose a heuristic, iterative solution algorithm that experimentally yields results close to the global optimum while guaranteeing the locality constraints of (4.9). The algorithm consists of repeatedly solving a (linearly constrained) convex quadratic program obtained by linearizing and relaxing the bilinear term (4.14). The goal is to minimize both the objective function and the violation of (4.14) in an iterative way until a certain feasibility tolerance is achieved.

The key step of this approach is to linearize (4.14) around the current values of the optimized gain matrices $\mathbf{G}$ and $\tilde{\mathbf{G}}$, here denoted as $\mathbf{G}^*, \tilde{\mathbf{G}}^* \in \mathbb{R}^{NL \times N^2}$, respectively. Assuming small deviations and neglecting bilinear terms, (4.14) can be approximated by

$$(\mathbb{I} - \mathbf{G}^*\mathbf{B})\tilde{\mathbf{G}} + \mathbf{G}^*\mathbf{B}\tilde{\mathbf{G}}^* = \mathbf{G}(\mathbf{C} + \mathbf{B}\tilde{\mathbf{G}}^*). \tag{4.20}$$

Small deviations can be enforced during optimization with the addition of the linear constraints

$$\begin{aligned} |\mathbf{G} - \mathbf{G}^*| &\leq \upsilon, \\ |\tilde{\mathbf{G}} - \tilde{\mathbf{G}}^*| &\leq \upsilon, \end{aligned} \tag{4.21}$$

where υ is a small positive constant.

Since (4.20) tightly couples $\mathbf{G}$ and $\tilde{\mathbf{G}}$, we relax constraint (4.20) and solve the problem iteratively until a certain feasibility tolerance is reached. To this end, we introduce the auxiliary matrix $\aleph \in \mathbb{R}^{NL \times N^2}$ to represent the feasibility violations of (4.20), i.e.,

$$\aleph = (\mathbb{I} - \mathbf{G}^*\mathbf{B})\tilde{\mathbf{G}} + \mathbf{G}^*\mathbf{B}\tilde{\mathbf{G}}^* - \mathbf{G}(\mathbf{C} + \mathbf{B}\tilde{\mathbf{G}}^*). \tag{4.22}$$

We also extend the objective function by a measure of the infeasibility of (4.20), here expressed by the Eucledian norm of $\aleph$ as

$$\tilde{\eta} = \boldsymbol{\delta}^T \mathbf{H}_{(1:L)} \mathbf{1} + \lambda ||\aleph||_2^2, \tag{4.23}$$

where $\lambda \in \mathbb{R}_+$ is a linear scalarization term of the multi-objective function (4.23). The larger the value of λ, the more focus is given on finding a feasible solution. It was observed experimentally that fast convergence and low objective value can be obtained by doubling the value of λ after each iteration of the algorithm.

Replacing (4.14) by the described linearization procedure, we obtain the following

sub-routine:

$$\begin{aligned} \mathbf{G}^*, \tilde{\mathbf{G}}^*, \tilde{\eta}^*, \mathbf{H}^*, \boldsymbol{\aleph}^* \in \underset{\mathbf{G}, \tilde{\mathbf{G}}, \tilde{\eta}, \mathbf{H}, \boldsymbol{\aleph}}{\arg\min} \quad & \tilde{\eta} \\ \text{s.t. } & (4.9),\ (4.15) - (4.17), \\ & (4.21) - (4.23). \end{aligned} \tag{4.24}$$

We iterate until $\max(\boldsymbol{\aleph}^*) \leq \tau$ for a given stopping criterion $\tau \in \mathbb{R}_+$. Initial values for $\tilde{\mathbf{G}}^*$ can be obtained, e.g., by solving (4.19) without the bilinear constraint (4.14), and for $\mathbf{G}^*$ by minimizing $||(\mathbb{I} - \mathbf{GB})\tilde{\mathbf{G}}^* - \mathbf{GC}||_2$ over $\mathbf{G}$ such that the locality constraints (4.9) are still fulfilled.

Although the optimality of the solution of the proposed iterative algorithm is not guaranteed, we show in Section 4.6 that its solution has an objective value less than one percent higher than the optimal solution of (4.19). Moreover, almost no line capacity constraint is violated in any realization of the uncertain flow rates.

4.6 Experiments

We validate our approach with a network simulation following the fluid flow model of Section 4.2 using differentiable approximations of equations (4.1)-(4.6). Two study cases are considered. First, the adaptive global coordination is demonstrated for a test network with three routers. Second, scalability is verified for a real medium-sized data center backbone network.

The network simulator was developed using the DifferentialEquations.jl package [102]. The QCQP and the iterative algorithms were implemented using JuMP.jl [34] as modeling language and Gurobi [53] as optimization solver. For running the experiment over the first topology we employed an Intel Core i7 laptop with 8 GB of RAM. The experiment over the second topology was run with the same hardware when optimizing for four flows, while a compute server with similar single-core performance but increased RAM memory was used for the real backbone network example.

We compare our approach against Random Packet Spraying (RPS) [30] and the Shortest-Path (SP) solution. RPS is a simple local routing policy that splits the flows equally among the available paths. This algorithm can be seen as an instance of the proposed approach, where the split ratios at all routers are fixed to 0.5 regardless of measurements the incoming rate measurements. SP is the solution with the lowest average delay for low traffic routes, but it can lead to excessive packet drops in case bandwidth constraints are violated. The proposed framework can also replicate the SP solution by fixing the split ratios of the links on the shortest path to 1 and setting the ratios on all other links to 0.

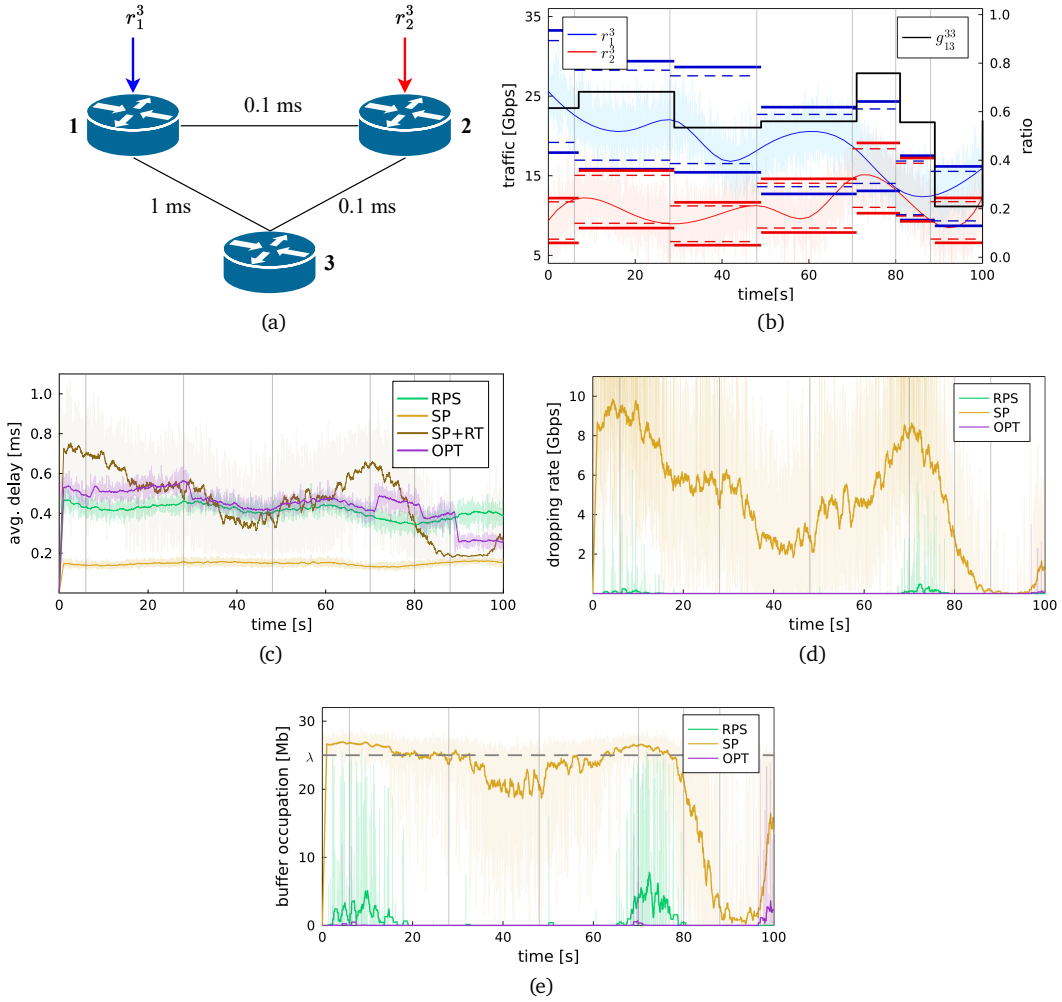

Figure 4.2: (a) 3-router test network with different link delays. (b) Simulated time profiles of the flows $r_1^3(t)$ and $r_2^3(t)$ in the network from (a) (blue/red) as well as the used uncertainty intervals (bold, solid) and warning levels (bold, dashed). The black line marks the fraction of the data flow $r_1^3(t)$ that is routed through the slow link (1,3). (c) Average delay, (d) dropping rate, and (e) buffer occupation for the flows from Fig. 4.2b applied to the network in Fig. 4.2a. Solid, bold lines: moving average of $1\ \mathrm{s}$ for clearer visualization.

4.6.1 3-Router Test Network

Consider the communication network shown in Fig. 4.2a. The bandwidth of all links is assumed to be $25\ \mathrm{Gbps}$, and the buffer size of all interfaces $25\ \mathrm{Mb}$. We assume that the average packet delay in the network is constant, and its value for each link is depicted in the diagram. Fig. 4.2a also shows two incoming flows, namely $r_1^3(t)$ and $r_2^3(t)$, which enter the network via routers 1 and 2, respectively, and have router 3 as the destination node.

The simulated traffic profile of both incoming flows is shown in Fig. 4.2b. It was generated by sampling two superposed Gaussian processes representing both a slow and a fast varying component according to the discussion in Section 4.1. Both processes have the Radial Basis Function (RBF) kernel as the covariance function but with different parameter values. The average rate of $r_1^3(t)$ and $r_2^3(t)$ range between 17 and $26\ \mathrm{Gbps}$ and 10 and $15\ \mathrm{Gbps}$, respectively, which means that routing without packet drop over the fastest path {(1,2), (2,3)} is not possible at all times.

We define the time-variable bounds for the robust set and the warning level as $\pm 30\%$ and $\pm 25\%$ of the filtered incoming flow rates, respectively. Whenever the warning level is violated, the reconfiguration of the local policies is triggered, and the bounds are updated based on current measurements. A fixed delay of 1 s between the optimization trigger and the update of the local policies is added to emulate the time required for reconfiguring the data plane in an actual deployment.

Fig. 4.2b also shows the fraction of flow $r_1^3(t)$ that is routed through the slower path (1,3) instead of the faster path {(1,2), (2,3)}, which is equivalent to the gain g_{13}^{33} in (4.6). Usage of the slow link (1,3) is minimized by the proposed algorithm, and it occurs only when the capacity of the fast link (2,3) is not sufficient to cover both incoming flows.

We compare our approach with RPS, i.e., half of the flow $r_1^3(t)$ is routed through (1,2) and the other half through (1,3), and with the SP solution, i.e., only the path {(1,2), (2,3)} is utilized independently if the total rate is larger than the capacity of the link (2,3). Fig. 4.2c shows that our approach has a variable average delay, which depends on both the rate values and the local controllers. When the total traffic rate is low, e.g., in the last $10\ \mathrm{s}$ of the simulation window, more data packets can be forwarded via the faster link (1,2). The average delay in the network decreases to $0.25\ \mathrm{ms}$, which is around 60% smaller than routing using RPS but 65% greater than using SP. When the total rate is high, e.g., in the first $25\ \mathrm{s}$ of the simulation or between seconds 70 and 80, the robust policy shows a higher average delay than RPS and SP due to the global coordination algorithm enforcing bandwidth limits, which induces a more conservative policy.

The advantage of our algorithm can be seen in Fig. 4.2d. While the SP policy causes 13% of the flows to be dropped because of the limited bandwidth of link (2,3), the proposed

robust policy leads to drops of only 0.001% of the total flow, which is equivalent to $20\ \mathrm{Mb}$ of data. Note that drops occur for our policy only when the assumption of limited burstiness of data flows is slightly violated. RPS routing shows a much lower dropping rate than SP, but with $600\ \mathrm{Mb}$ of dropped data over the simulation horizon, it still drops 30 times more data than our proposed algorithm.

It is worth noting that in real-world deployments, even low dropping rates may have a significant impact on the observed performance, e.g., on latency, due to packet retransmissions that are triggered in the most prevalent transport protocols. Therefore, in practice, RPS and SP routing policies would lead to significantly higher average delays than the values resulting from the flow simulation. To demonstrate the impact of packet retransmissions, we added to the average delay comparison the SP solution with the retransmission of dropped packets. We lower-bounded the average delay of the retransmitted packets as two times the delay of the slower link. We assumed that retransmissions did not result in additional packet drops. As shown in Fig. 4.2c, during periods of high packet drops, SP with retransmission has a higher average delay compared to the proposed method, demonstrating a clear advantage of our approach.

Fig. 4.2e also shows that our approach uses the buffers much less, with an average occupancy of 0.31% compared to 2.63% for RPS and 83.08% for the SP policies. These values are also a result of the hard constraint on the bandwidth limits imposed by the global coordination algorithm.

In this exemplary setup, we demonstrated the capability of the proposed routing scheme to handle a high amount of data flow in a capacity-limited network. We also showed that the algorithm proposes routing with an average delay comparable to routing with RPS but with much less packet drop. In the following, we evaluate the algorithm's scalability when employed in a network with tens of nodes and hundreds of possible source/destination pairs, which is often the case in real-world backbone communication networks. We assess the heuristic presented in Section 4.5 in terms of its solution's optimality gap and constraint violations.

4.6.2 Real Data Center Backbone Network

We evaluate the proposed routing framework on a more extensive, real data center backbone network shown in Fig. 4.3a that extends over North America, Europe, and East Asia [63]. The network consists of 12 nodes and 19 bi-directional links, with four of them connecting nodes placed on different continents. Link delays for intercontinental connections are assumed to be $1\ \mathrm{ms}$, while intracontinental ones are $0.1\ \mathrm{ms}$.

For this network size, the proposed QCQP problem (4.12) has too many variables and constraints to be directly solvable within the coordination time window. For this example,

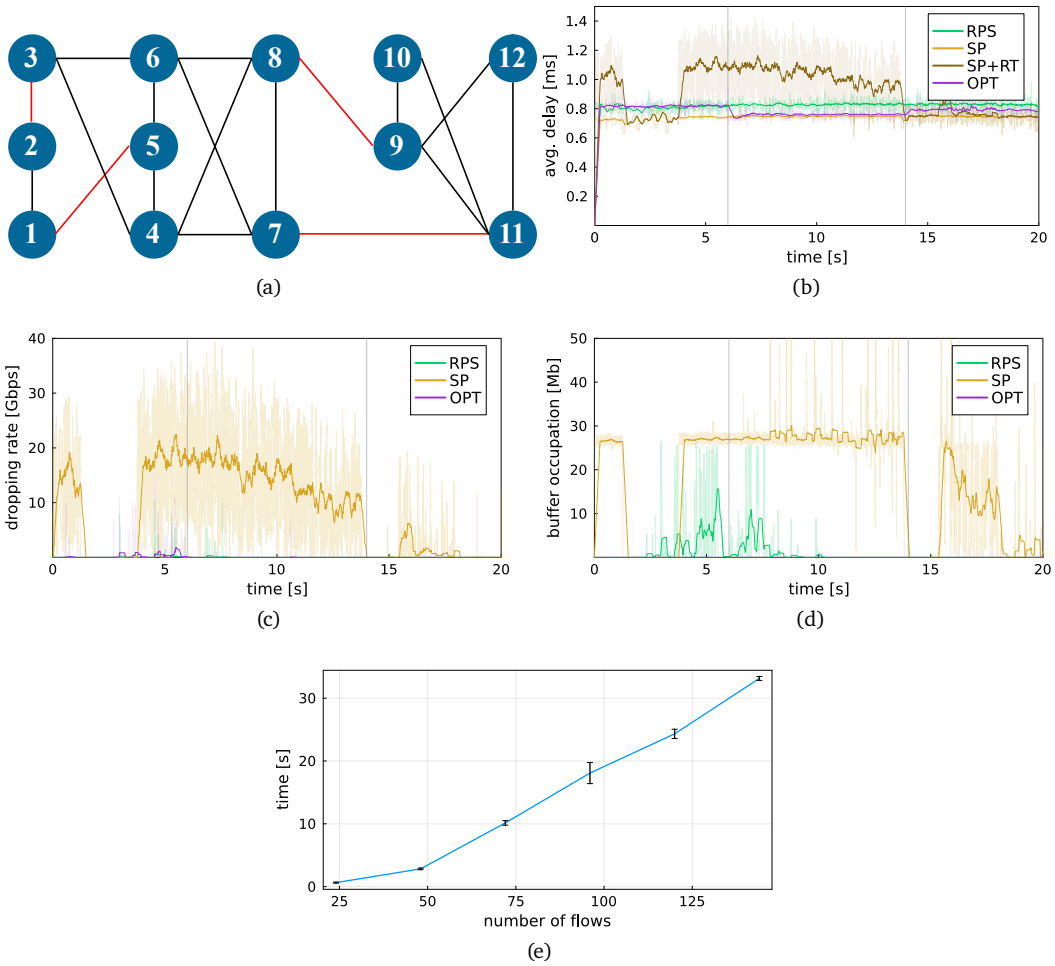

Figure 4.3: (a) Real data center backbone network taken from [63]. Intercontinental links (red) are assumed to have ten times higher delay compared to intracontinental ones (black). (b) Average delay, (c) dropping rate, and (d) buffer occupation for the flows $r_2^6(t)$, $r_3^6(t)$, $r_{11}^8(t)$ and $r_{12}^8(t)$ in the real data center backbone network of (a). Solid, bold lines: moving average of $250\ \mathrm{ms}$ for clearer visualization. (e) Time to solve the global coordination problem using the proposed heuristic for a certain number of flows. Results averaged over five runs and network topology from (a).

we define the coordination time ceiling as 5 min, which according to measurements shown in [86], would allow for variations of less than $\pm 20\%$ in the incoming rates for a typical backbone network. When considering only four incoming flows, namely $r_2^6(t)$, $r_3^6(t)$, $r_{11}^8(t)$ and $r_{12}^8(t)$, we can guess and validate the optimal solution analytically, and we use the solution to evaluate the solution quality of the iterative method to solve the global coordination problem proposed in Section 4.5.

The iterative algorithm converges in 2.7 s. From the computed locally-implementable gain matrix $\mathbf{G}^*$, we derive $\tilde{\mathbf{G}}^*$ via (4.19) and use this to determine the maximum violations of the bandwidth constraints (4.10) via (4.16)-(4.17). Even for the worst-case input rates, rate violations are below 0.003% of line capacity for all links. Moreover, the solution of the iterative algorithm has only 0.62% higher objective value than the analytical solution.

After verifying that the optimality gap of our proposed iterative algorithm is small, we compare its performance against RPS and SP. The traffic profile of the flows $r_2^6(t)$, $r_3^6(t)$, $r_{11}^8(t)$ and $r_{12}^8(t)$ were generated using the same method as described in Subsection 4.6.1. The average rate of the flows was set at 20 Gbps.

The results of the simulation are displayed in Fig. 4.3. During the 20 s simulation window, the global coordination algorithm was triggered twice, 6 s and 14 s after the start of the simulation. Fig. 4.3b shows that our proposed iterative algorithm consistently outperforms RPS in terms of average delay throughout the simulation period, and at times even reaches the same level as SP. When considering packet retransmissions, our approach has much lower average delay than SP during most part of the simulation window. Additionally, Fig. 4.3c and Fig. 4.3d demonstrate that our algorithm effectively reduces packet drops and buffer occupancy, which helps to avoid packet retransmissions.

We also evaluate the scalability of the proposed iterative optimization algorithm in the presence of an increasing number of incoming flows. The results depicted in Fig. 4.3e indicate that, despite the fact that the number of constraints grows quartically with the number of nodes if flows from every node to every other node are allowed, resulting in 144 flows here, the optimization method still converges within 33 s. This is significantly below the assumed 5 min limit for reconfiguring local policies.

The experimentally observed good scalability property of the iterative algorithm is attributed to the ability of modern LP solvers, such as Gurobi, to very efficiently eliminate redundant constraints. In the case with 144 flows, the presolve eliminates more than 82% of the constraints in the problem. This makes our algorithm a suitable solution for networks with up to tens of nodes and hundreds of aggregated flows. This is sufficient for many real backbone or data center networks.

4.7 Summary

This chapter presented a new two-level coordination scheme for routing in backbone communication networks using the robust optimization framework proposed in this thesis. Using techniques presented in Chapter 2, we designed an adaptive flow-controlling mechanism whose design is compatible with modern SDN technologies and which ensures low buffer occupancy and packet dropping.

The global coordination algorithm used to optimize the local decision rules was formulated as a robust QCQP, which could be solved by commercial solvers. To improve the method's tractability when applied to networks with tens of nodes and hundreds of aggregated flows, we devised an iterative heuristic that is significantly faster to solve than the QCQP and yields near-optimal results.

We evaluated the proposed mechanism in a small, exemplary communication network and a realistic data center topology. The results showed that our robust approach outperforms standard routing protocols regarding total dropping rate and buffer occupancy. Moreover, due to its adaptation mechanism, the algorithm had at times of low traffic rate average delay as low as the SP solution. Thus, we demonstrated the usefulness of our robust optimization framework in flow-controlling in backbone communication networks.

5 Adaptive Survivable Virtual Communication Network Embedding with an Application to Power Systems

This chapter discusses how the proposed robust framework can be applied to Smart Grids for the minimization of communication costs while considering the power grid constraints. More specifically, we address the problem of optimal placement of Local Control Units (LCU) coupled with On-Load Tap Changer (OLTC) transformers to reduce the likelihood of voltage violations in the presence of contingencies of communication links. We formulate this problem as an SVNE task with a flexible number of redundancies using MILP. Section 5.1 introduces the SVNE problem and describes our proposed approach. In Section 5.2, we present the VNE model, and in Section 5.3, we define the SVNE problem with flexible redundancy. We formulate the SVNE problem as a robust MILP model in Section 5.4, where the uncertainty set comprises failure modes of the underlying network. To demonstrate the effectiveness of our approach, we apply the algorithm to a toy example in Section 5.5 and a Smart Grid application. In the last experiment, we show that the placement computed by our algorithm is more cost-efficient than fixed redundancy, as it can flexibly adapt to different power generation and consumption patterns. Table 5.1 shows the main symbols used to describe this chapter's main optimization problem presented in Section 5.4.

5.1 Overview of the SVNE problem

Network Virtualization (NV) and the closely related field of SDN [32, 52, 93] are well-established in the field of communication systems. They are key technologies for the Internet [96] and cloud infrastructures [62] that enable increased reliability of the provided services, efficient usage of the communication hardware, and the ability to perform quick adaptations in communication networks [54].

Table 5.1: List of symbols of main optimization problem of Chapter 5.

Parameters	
$V \in \mathbb{Z}_+$	Number of substrate nodes
$L \in \mathbb{Z}_+$	Number of substrate links
$\tilde{V} \in \mathbb{Z}_+$	Number of virtual nodes
$\tilde{L} \in \mathbb{Z}_+$	Number of virtual links
$K \in \mathbb{Z}_+$	Degree of redundancy
$c_i \in \mathbb{R}^V$	Cost for substrate node allocation
$p_i \in \mathbb{R}^V$	Maximum substrate node processing capacity
$c_{ijn} \in \mathbb{R}^L$	Cost for substrate link allocation
$b_{ijn} \in \mathbb{R}^L$	Maximum substrate link bandwidth capacity
$r_i \in \mathbb{R}^V,\ r_{ijn} \in \mathbb{R}^L$	Substrate node and link reliabilities
$p^u \in \mathbb{R}^{\tilde{V}}$	Required virtual node processing capacity
$b^{uv} \in \mathbb{R}^{\tilde{L}}$	Required virtual link bandwidth capacity
$r^u \in \mathbb{R}^{\tilde{V}},\ r^{uv} \in \mathbb{R}^{\tilde{L}}$	Required virtual node and link reliabilities
$\pi_\omega \in [0,1]$	Probability of failure mode
$\lambda^V,\ \lambda^L \in \mathbb{R}$	Adaptation cost for node and link mappings
Sets	
$\mathcal{V} \subset \mathbb{Z}_+$	Set of substrate network nodes (index i)
$\mathcal{L} \subset \mathbb{Z}_+^3$	Set of substrate network links (indices i, j, n)
$\tilde{\mathcal{V}} \subset \mathbb{Z}_+$	Set of virtual network request nodes (index u)
$\tilde{\mathcal{L}} \subset \mathbb{Z}_+^2$	Set of virtual network request links (indices u, v)
$\mathcal{K} \subset \mathbb{Z}_+$	Set of embedding indices (index k)
$\Omega \subset \{0,1\}^{VL}$	Uncertainty set of failure modes (index ω)
Integer Variables	
$\mathtt{x}^u_{i,k} \in \{0,1\}^{\tilde{V}VK}$	Indicator of virtual node placement
$\mathtt{y}^{uv}_{ijn,k} \in \{0,1\}^{\tilde{L}LK}$	Indicator of virtual link placement
$\xi_{i,\omega} \in \{0,1\}^{V^2L}$	Indicator of failure of substrate node in the scenario
$\xi_{ijn,\omega} \in \{0,1\}^{VL^2}$	Indicator of failure of substrate link in the scenario
$\eta^u_{k,\omega} \in \{0,1\}^{\tilde{V}KVL}$	Indicator of failure of virtual node within an embedding in the scenario
$\eta^{uv}_{k,\omega} \in \{0,1\}^{\tilde{L}KVL}$	Indicator of failure of virtual link within an embedding in the scenario

These properties make NV and SDN also well suited for Smart Grids, which integrate power and communication networks [125]. For instance, NV enabled by SDN technology can provide scalable and efficient solutions for virtual power plants [125]. SDN can also support IEC 61850 implementations [89] and is an important tool for increasing the resilience [65] and cyber security [23] of Smart Grids.

At the core of many NV tasks lies the VNE problem: Networks of inter-connected logical functional blocks, so-called Virtual Network Requests (VNR), are to be mapped for execution onto physical networks of compute resources and communication links, the so-called Substrate Networks (SN). VNE is a widely studied problem [45], whose one of its subproblems, SVNE [59], focuses on increasing the reliability of the embedded logic in case of SN failures.

5.1.1 Existing Approaches

Various mechanisms have been proposed to achieve survivability in VNE. Reactive approaches involve redirecting communication traffic to an alternative link or reinstantiating a compute job on another node after a network failure has occurred [110]. These approaches require sufficient backup capacities for all potential failures before the contingency to be effective [69]. The backup capacities, however, could be used for other purposes until the failure occurs. Consequently, reactive approaches can lead to service interruptions during the transition process, resulting in reduced quality of service that can be unacceptable in critical infrastructure such as Smart Grids.

Another class of SVNE methods called proactive mechanisms keeps redundant backup resources online at all times to ensure immediate service continuity in the event of a failure [25]. However, these approaches have the drawback that the backup capacities must be unused at all times to ensure seamless adaptation, thus leading to inefficient and costly embedding.

5.1.2 Contributions

In this chapter, we formulate and solve the SVNE problem proactively using a robust MILP formulation. Unlike previous approaches such as [25], we do not provide a fixed degree of redundancy, e.g., two-fold instantiation for all VNR elements. Instead, we flexibly embed each VNR element one or more times onto the SN, satisfying the individual reliability requirements of the virtual elements at the minimal embedding cost. This reduces SN resource usage for uncritical elements while at the same time allowing critical elements to be secured with high degrees of redundancy. The automated redundancy generation process is well-suited for adapting the embedding and its redundancy degree when the

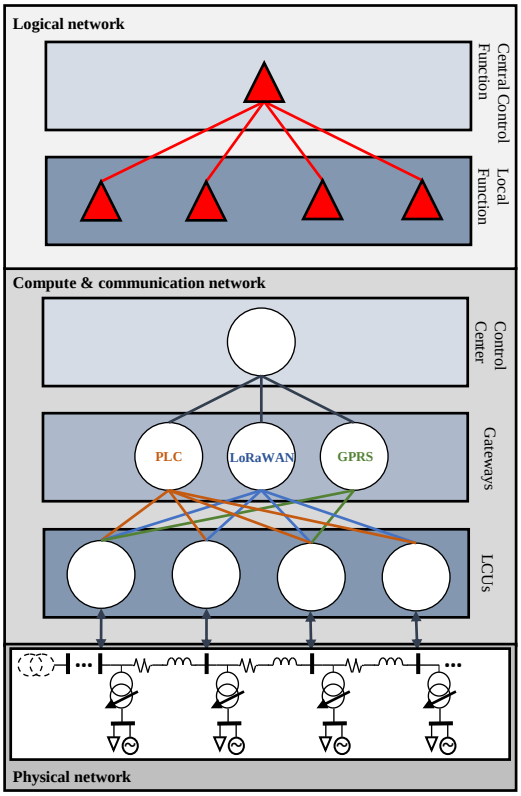

Figure 5.1: Overview of Smart Grid setup with different communication technologies. The logical network represents a voltage control application.

reliability requirements of the VNR or the reliability guarantees of the SN change over time.

We show how the proposed approach can suitably be employed to reduce communication costs in smart distribution grids. Renewable energy integration leads to potential voltage violations, especially in Medium Voltage (MV) feeders [38], and various monitoring and control schemes have been proposed for this purpose, e.g., [81]. We construct a VNR for a voltage monitoring scheme that connects OLTC power transformers with the control center. We embed the VNR into an SN consisting of different communication technologies, such as Power Line Communication (PLC), Long Range Wide Area Network (LoRaWAN), and General Packet Radio Service (GPRS) networks, see Fig. 5.1. Unlike the dedicated communication networks of transmission grid operators, the considered links are highly

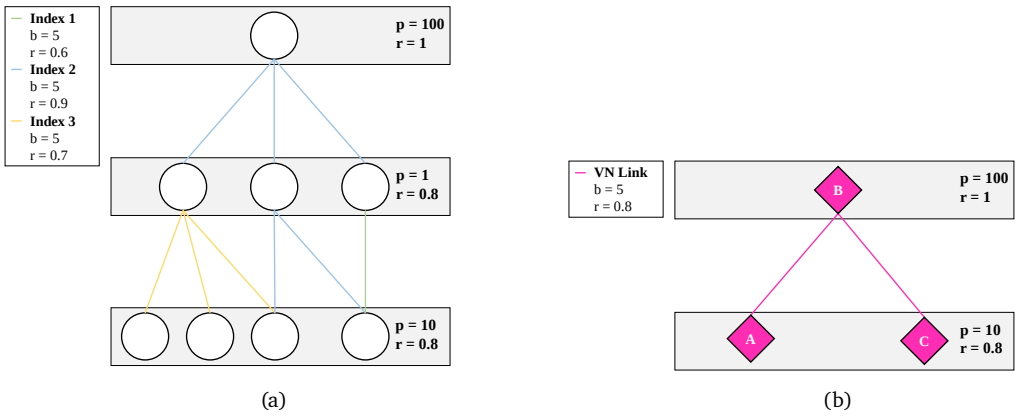

Figure 5.2: (a) SN where each color represents a different technology n of the substrate links (i, j, n), (b) VNR with specific reliability requirements.

prone to failures and outages individually, requiring a certain extent of redundancy to guarantee reliable service execution. We compute the required reliability for each link given the current grid state and use our proposed approach to generate minimal redundancy for each generation/load pattern. Compared to a fixed dual redundancy, as is, e.g., done in [84] using PLC and wireless mesh networks, we can reduce average operating costs by more than half in our simulated experiments.

5.2 SN and VNR Models

Let an SN with V nodes and L links be defined as an undirected multigraph $G = (\mathcal{V}, \mathcal{L})$ with substrate nodes $\mathcal{V}$ and substrate links $\mathcal{L}$. Each substrate node $i \in \mathcal{V}$ has a cost of allocation c_i and processing capacity p_i. Substrate links are identified by the triple $(i, j, n) \in \mathcal{L}$, with i and j being the end nodes of the link and n being the link index between both nodes. We use the multigraph notation to represent heterogeneous networks, where different communication technologies connect the same nodes. Each link has a cost of allocation c_{ijn} and available bandwidth b_{ijn}.

In a non-perfect SN, one or more nodes and links can fail anytime. Any such failure mode of the SN can be encoded by a vector of failure indicators, one for each network element. We do not consider correlations between failures in time and assume the failure modes to be independently and identically distributed according to a known probability distribution over this finite set of vectors. To simplify our outline, we describe the failure

distributions over the SN only with their marginal failure probabilities for each element and assume independence between the failures of different elements. Specifically, we define for each substrate node a reliability r_i and for each substrate link a reliability r_{ijn} indicating that the node i and link (i, j, n) are available that fraction of time, respectively.

Analogously, a VNR with $\tilde{V}$ virtual nodes and $\tilde{L}$ virtual links is defined as an undirected graph $\tilde{G} = (\tilde{\mathcal{V}}, \tilde{\mathcal{L}})$ with virtual nodes $\tilde{\mathcal{V}}$ and virtual links $\tilde{\mathcal{L}}$. A virtual node $u \in \tilde{\mathcal{V}}$ requires a computing capacity p^u for processing data, and a virtual link $(u, v) \in \tilde{\mathcal{L}}$ requires a data bandwidth b^{uv} to transmit processed data from u to v.

When a VNR represents a critical service, as is the case for many Smart Grid applications, it is necessary to specify suitable reliability requirements. We thus define a reliability requirement r^u for each virtual node u and a reliability requirement r^{uv} for each virtual link (u, v), which determine the allowed minimum fraction of time each component of the request must be available. Fig. 5.2 shows an example of an SN and a VNR along with their capacity and reliability parameters.

5.3 SVNE Problem with Flexible Redundancy

VNE is an allocation problem that involves mapping each virtual node of a VNR onto a substrate node and each virtual link onto a path in the SN. This process must consider the VNR's connectivity requirements, as well as capacity and bandwidth constraints. SVNE builds upon VNE by adding an additional requirement: ensuring the operational reliability of the SN for each individual virtual element. To achieve this while ensuring service continuity, we allow the SVNE to embed virtual elements onto several redundant copies in the SN network. This approach helps to ensure the required reliability even if a VNR element has a greater reliability requirement than provided by any single SN element, as the probability of multiple copies failing at the same time is significantly reduced (assuming that the failures of individual SN elements are fully independent).

To account for the variable amount of redundancy required for each VNR element, we determine up to K full VNEs, where $K \geq 2$ is a user-defined parameter, which are not required to be disjoint. If a virtual node is mapped onto the same SN node in several of these VNEs, we assume that only one copy is executed in the SN environment, and that SN resources are used only once. As a result, cost minimization drives many copies to be located in the same place, implying a degree of redundancy smaller than K, while still satisfying the reliability requirement. The same holds true for the embedding of virtual links onto SN paths.

However, it's possible that the paths are only partly identical, i.e., they split or merge at

certain nodes. In the case of splitting, we assume that the SDN-based network replicates the communication messages along both splitting paths. In the case of path mergers, we assume that the data traffic from one incoming link is dropped, as long as the other one is available. These procedures enable us to operate with varying degrees of redundancy for each element of the VNR, while still ensuring reliable communication.

In this thesis, our objective is to determine K feasible VNEs that jointly fulfill the reliability requirements posed to the VNR. We also aim for minimal VNE changes in response to slight changes in the VNR or SN parameters that do not alter the topology.

5.4 SVNE Problem as a Robust MILP

In this section, we formulate the SVNE problem as a MILP model. The problem's objective is to minimize the resource utilization while ensuring the reliability requirement of each VNR element. This is achieved by constraining the number of allowed failure modes based on their respective probabilities of occurrence. Furthermore, to avoid service disruption, we also aim to minimize the number of VNR element changes in subsequent embeddings.

Given a VNR $\tilde{G} = (\tilde{\mathcal{V}}, \tilde{\mathcal{L}})$, an SN $G = (\mathcal{V}, \mathcal{L})$, and a set of embedding indices $\mathcal{K} = \{1, \ldots, K\}$, we define binary decision variables $\mathrm{x}^u_{i,k}$ for all nodes $u \in \tilde{\mathcal{V}}$, $i \in \mathcal{V}$, and $k \in \mathcal{K}$, which indicate whether a node u in the VNR is mapped to node i in the SN under the k-th embedding, i.e.,

$$\mathrm{x}^u_{i,k} = \begin{cases} 1, \text{ if } u \text{ embedded onto } i \text{ in embedding } k, \\ 0, \text{ otherwise.} \end{cases}$$

Additionally, we introduce binary decision variables $\mathrm{y}^{uv}_{ijn,k}$ for all links $(u, v) \in \tilde{\mathcal{L}}$, $(i, j, n) \in \mathcal{L}$, and $k \in \mathcal{K}$. These variables represent whether a link (u, v) in the VNR is mapped to a path (i, j, n) in the SN under the k-th embedding as

$$\mathrm{y}^{uv}_{ijn,k} = \begin{cases} 1, \text{ if } (i, j, n) \text{ is part of the } k\text{-th embedding of } (u, v), \\ 0, \text{ otherwise.} \end{cases}$$

To ensure that each virtual node is mapped onto exactly one substrate node, we enforce the following constraint:

$$\sum_{i \in \mathcal{V}} \mathrm{x}^u_{i,k} = 1, \forall u \in \tilde{\mathcal{V}}, \forall k \in \mathcal{K}. \tag{5.1}$$

We also impose via (5.2) that all virtual links are embedded onto a substrate path with minimum length of one link. Constraint (5.3), also called the multi-commodity flow

constraint [40], ensures a continuous substrate path for each virtual link embedding:

$$\sum_{(i,j,n)\in\mathcal{L}} \mathtt{y}^{uv}_{ijn,k} \geq 1, \forall (u,v) \in \tilde{\mathcal{L}}, \forall k \in \mathcal{K}, \tag{5.2}$$

$$\sum_{j\in\mathcal{N}_i} \left(\mathtt{y}^{uv}_{ijn,k} - \mathtt{y}^{uv}_{jin,k}\right) = \mathtt{x}^{u}_{i,k} - \mathtt{x}^{v}_{i,k}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall i \in \mathcal{V}, \forall k \in \mathcal{K}, \tag{5.3}$$

where, $\mathcal{N}_i$ is the set of neighbors of SN node i.

As multiple embeddings of a single VNR element onto the same SN element imply that only one instance is executed in the SN, we introduce the following auxiliary variables: $\hat{\mathtt{x}}^u_i \in \{0,1\}$, which indicates whether VNR node $u \in \tilde{\mathcal{V}}$ is mapped onto SN node $i \in \mathcal{V}$ in any of the K mappings, and $\hat{\mathtt{y}}^{uv}_{ijn} \in \{0,1\}$, which indicates whether SN link $(i,j,n) \in \mathcal{L}$ is part of any of the K SN paths to which VNR link $(u,v) \in \tilde{\mathcal{L}}$ is mapped. To enforce these definitions, we include the following constraint into the optimization problem:

$$\hat{\mathtt{x}}^u_i \geq \mathtt{x}^u_{i,k}, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \forall k \in \mathcal{K}, \tag{5.4}$$

$$\hat{\mathtt{y}}^{uv}_{ijn} \geq \mathtt{y}^{uv}_{ijn,k}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \forall k \in \mathcal{K}. \tag{5.5}$$

The capacity and bandwidth limitations are modeled using the auxiliary variables $\hat{\mathtt{x}}^u_i$ and $\hat{\mathtt{y}}^{uv}_{ijn}$ as

$$\sum_{u\in\tilde{\mathcal{V}}} p^u \hat{\mathtt{x}}^u_i \leq p_i, \forall i \in \mathcal{V}, \tag{5.6}$$

$$\sum_{(u,v)\in\tilde{\mathcal{L}}} b^{uv} \hat{\mathtt{y}}^{uv}_{ijn} \leq b_{ijn}, \forall (i,j,n) \in \mathcal{L}. \tag{5.7}$$

To account for uncertainty in the reliability of virtual elements, we need to ensure that the probability of all K embeddings of a virtual element failing simultaneously does not exceed its maximal allowed failure level, which is defined as one minus its required reliability level. To formulate these reliability constraints within our robust framework, we ensure that the optimal solution of the SVNE problem is feasible for a certain discrete, finite uncertainty set that contains scenarios of the substrate elements' failures.

We begin by introducing constraints to ensure the reliability of VNR nodes, followed by constraints for VNR links. Let Ω be the uncertainty set containing all possible single-failure modes of the substrate elements, including an additional scenario for the undisturbed

case, and let π_ω denote the probability of scenario $\omega \in \Omega$. We define parameter $\xi_{i,\omega}$ to indicate whether node $i \in \mathcal{V}$ fails under scenario ω as

$$\xi_{i,\omega} = \begin{cases} 1, \text{ if } i \text{ failed in scenario } \omega, \\ 0, \text{ otherwise.} \end{cases}$$

To indicate whether the substrate node i that virtual node $u \in \tilde{\mathcal{V}}$ is mapped onto in embedding k fails in scenario ω, we use binary variable $\eta^u_{k,\omega}$. This is expressed by the following constraint:

$$\begin{aligned} \eta^u_{k,\omega} \geq \xi_{i,\omega} \mathtt{x}^u_{i,k}, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \\ \forall k \in \mathcal{K}, \forall \omega \in \Omega. \end{aligned} \tag{5.8a}$$

Then, we summarize $\eta^u_{k,\omega}$ over all K embeddings to determine whether all K embeddings of virtual node u fail in scenario $\omega \in \Omega$, which is expressed by the constraint

$$\zeta^u_\omega \geq \sum_{k \in \mathcal{K}} \eta^u_{k,\omega} - (K-1), \forall u \in \tilde{\mathcal{V}}, \forall \omega \in \Omega. \tag{5.8b}$$

Finally, we ensure that the sum of the probabilities of all scenarios where all embeddings fail are less or equal than the maximum allowed probability of failure of that virtual node:

$$\sum_{\omega \in \Omega} \pi_\omega \zeta^u_\omega \leq 1 - r^u, \forall u \in \tilde{\mathcal{V}}. \tag{5.8c}$$

For virtual links we construct analogous constraints (5.9) using the link element failure indicators $\xi_{ijn,\omega}$ and the auxiliary embedding variables $\eta^{uv}_{k,\omega}$ and ζ^{uv}_ω to express the failure of one or all embeddings for a link in a given scenario, respectively. This results in

$$\begin{aligned} \eta^{uv}_{k,\omega} \geq \xi_{ijn,\omega} \mathtt{y}^{uv}_{ijn,k}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \\ \forall k \in \mathcal{K}, \forall \omega \in \Omega, \end{aligned} \tag{5.9a}$$

$$\zeta^{uv}_\omega \geq \sum_{k \in \mathcal{K}} \eta^{uv}_{k,\omega} - (K-1), \forall (u,v) \in \tilde{\mathcal{L}}, \forall \omega \in \Omega, \tag{5.9b}$$

$$\sum_{\omega \in \Omega} \pi_\omega \zeta^{uv}_\omega \leq 1 - r^{uv}, \forall (u,v) \in \tilde{\mathcal{L}}. \tag{5.9c}$$

The previous conditions imply that communication can still be routed through a node, even if it failed. This may sometimes be plausible, but we present the following "node-failure-implies-link-failure" conditions that can be used in addition to (5.9a) to ensure

that a node failure implies also the failure of its adjacent links, or equivalently, that a virtual link embedding fails if any of the nodes on its path fails:

$$\eta^{uv}_{k,\omega} \geq \xi_{i,\omega} \mathtt{y}^{uv}_{ijn,k}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \forall k \in \mathcal{K}, \forall \omega \in \Omega, \tag{5.10a}$$

$$\eta^{uv}_{k,\omega} \geq \xi_{j,\omega} \mathtt{y}^{uv}_{ijn,k}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \forall k \in \mathcal{K}, \forall \omega \in \Omega. \tag{5.10b}$$

We aim to minimize the embedding cost considering only the actually executed instances of the virtual elements. Thus, we can write the SVNE problem as the following MILP model:

$$\eta = \min \sum_{u \in \tilde{\mathcal{V}}} \sum_{i \in \mathcal{V}} c_i \hat{\mathtt{x}}^u_i + \sum_{(u,v) \in \tilde{\mathcal{L}}} \sum_{(i,j,n) \in \mathcal{L}} c_{ijn} \hat{\mathtt{y}}^{uv}_{ijn}$$
$$\text{s.t. } (5.1) - (5.10b). \tag{5.11}$$

When we re-optimize the VNE problem several times to adapt it to slightly changing parameters of the VNR or the SN, it is desirable that the embedding changes as little as possible. This avoids the replacement of the embedded logical functions or the reshaping of the communication paths. To achieve this, we define the auxiliary binary variables $\bar{\mathtt{x}}^u_i$ and $\bar{\mathtt{y}}^{uv}_{ijn}$, which indicate whether the current node and link mappings remain the same as in a previous mapping, respectively. Specifically, $\bar{\mathtt{x}}^u_i$ and $\bar{\mathtt{y}}^{uv}_{ijn}$ are defined as follows:

$$\bar{\mathtt{x}}^u_i \geq \hat{\mathtt{x}}^u_i - \hat{\mathtt{x}}^{u\prime}_i, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \tag{5.12}$$

$$\bar{\mathtt{x}}^u_i \geq \hat{\mathtt{x}}^{u\prime}_i - \hat{\mathtt{x}}^u_i, \forall u \in \tilde{\mathcal{V}}, \forall i \in \mathcal{V}, \tag{5.13}$$

$$\bar{\mathtt{y}}^{uv}_{ijn} \geq \hat{\mathtt{y}}^{uv}_{ijn} - \hat{\mathtt{y}}^{uv\,\prime}_{ijn}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}, \tag{5.14}$$

$$\bar{\mathtt{y}}^{uv}_{ijn} \geq \hat{\mathtt{y}}^{uv\,\prime}_{ijn} - \hat{\mathtt{y}}^{uv}_{ijn}, \forall (u,v) \in \tilde{\mathcal{L}}, \forall (i,j,n) \in \mathcal{L}. \tag{5.15}$$

where $\hat{\mathtt{x}}^{u\prime}_i$ and $\hat{\mathtt{y}}^{uv\,\prime}_{ijn}$ are the values of the previous mapping.

Thus, we can modify (5.11) to include the cost of changing embeddings as

$$\eta = \min \sum_{u \in \tilde{\mathcal{V}}} \sum_{i \in \mathcal{V}} \left(c_i \hat{\mathtt{x}}^u_i + \lambda^V \bar{\mathtt{x}}^u_i \right) + \sum_{(u,v) \in \tilde{\mathcal{L}}} \sum_{(i,j,n) \in \mathcal{L}} \left(c_{ijn} \hat{\mathtt{y}}^{uv}_{ijn} + \lambda^L \bar{\mathtt{y}}^{uv}_{ijn} \right)$$
$$\text{s.t. } (5.1) - (5.15), \tag{5.16}$$

where λ^V and λ^L parametrize the adaptation cost for node and link mappings, respectively.

5.5 Experiments

In this section, we demonstrate our SVNE algorithm first on the exemplary communication network shown in Fig. 5.2 and on a Smart Grid voltage control setting, as shown in Fig. 5.1. The MILP model (5.16) is implemented in Python, utilizing the modeling language PuLP [87], and is solved with CPLEX [28].

To generate the set of scenarios Ω for both examples, we consider the failure of any single element $\sigma \in \Sigma$, along the normal operation mode. We calculate the probability π_ω of each scenario $\omega \in \Omega$ using the failure indicators $\xi_{\sigma,\omega}$ of element σ, as shown in following:

$$\pi_\omega = \frac{1}{Z} \prod_{\sigma \in \Sigma} \left[(1 - r_\sigma)\, \xi_{\sigma,\omega} + r_\sigma \left(1 - \xi_{\sigma,\omega}\right) \right]. \tag{5.17}$$

Here, r_σ denotes the reliability of element σ, and Z is a suitable constant for normalizing the probabilities. Note that, due to the normalization step, the marginal reliability of the elements of the SN is increased compared to the specified values.

5.5.1 Demonstrative Communication Network

In this example, we show how the proposed SVNE algorithm embeds VNEs with reliability constraints on nodes, links, both nodes and links, and how adaptation costs affect subsequent embeddings. Unless stated otherwise, we use the capacity and reliability parameters from Fig. 5.2. We set the cost for embedding c_i for all substrate nodes to 1, and the cost for link embedding c_{ijn} to 5, 10, and 20 for links with index 1, 2, and 3, respectively. The number of full VNEs K is set to 2, which is sufficient to satisfy all VNR requirements. The resulting embeddings are shown in Fig. 5.3.

First, we embed the VNR onto the SN considering only the node reliability constraints (5.8a)-(5.8c). The resulting embedding is shown in the left diagram of Fig. 5.3a. The SVNE algorithm uses the cheapest links to embed the request, as the link reliability constraints (5.9a)-(5.9c) are not considered in this case. When we increase the reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$, the program generates another copy of node A, along with its adjacent links, as seen in the right diagram of Fig. 5.3a.

Next, we consider embeddings using only the link failure constraints (5.9a)-(5.9c). The left diagram of Fig. 5.3b shows the resulting embedding for default parameters. The program embeds the virtual links onto substrate links of indices 2 and 3 because the reliability of substrate links with index 1 is below the required reliability level of the VNR, demanding a double embedding at an increased cost. When we increase the reliability

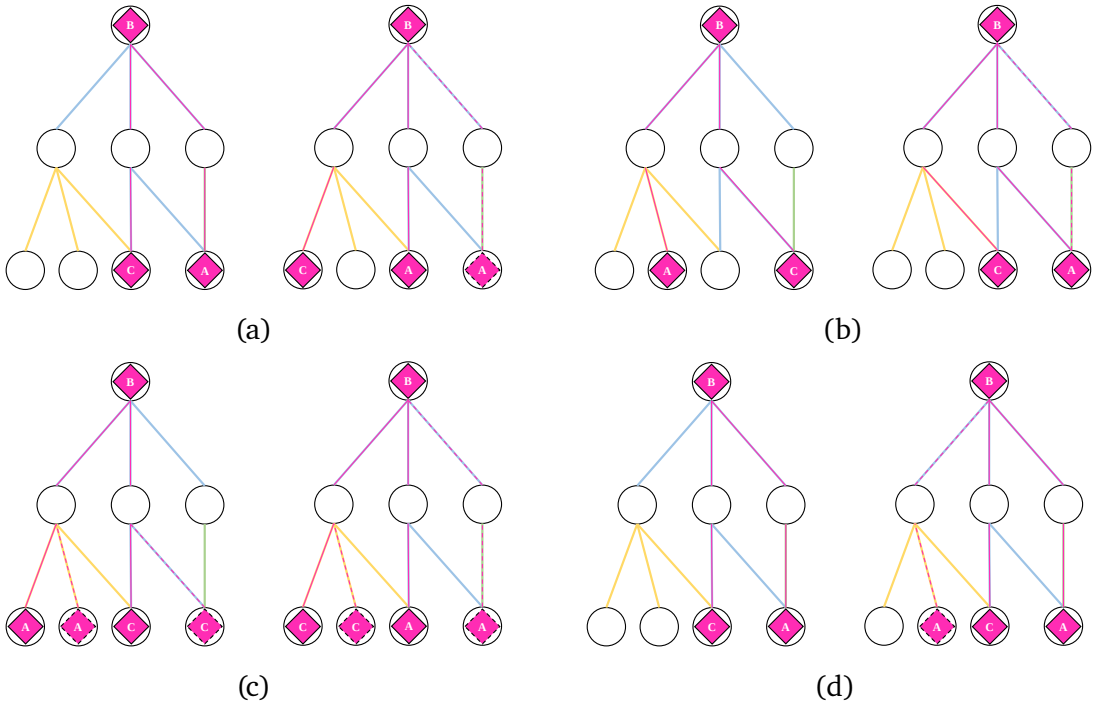

Figure 5.3: Proposed SVNE algorithm applied to the communication network of Fig. 5.2 solved with three different sets of redundancy constraints. The left graph always shows the embedding with the default parameters, while the right one is obtained by increasing the reliability requirements of the VNR. Dashed lines denote redundant link embeddings, while nodes with dashed outlines denote redundant virtual nodes. (a) Reliability constraints (5.8a)-(5.8c) active and increase of reliability requirement for node A from $r^A = 0.8$ to $r^A = 0.95$. (b) Reliability constraints (5.9a)-(5.9c) active and increase of reliability requirement of link (A, B) from $r^{AB} = 0.8$ to $r^{AB} = 0.99$. (c) Reliability constraints (5.8a)-(5.8c), (5.9a)-(5.9c), and (5.10a)-(5.10b) active and increase of reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$ and of link (A, B) from $r^{AB} = 0.8$ to $r^{AB} = 0.9$. (d) Same as (a) but with adaptation costs $\lambda^V = \lambda^L = 100$.

requirement of virtual link (A, B) to $r^{AB} = 0.97$, that virtual link is mapped onto two disjoint paths in the SN, as seen in the right diagram of Fig. 5.3b.

We now consider embeddings with all the conditions (5.8a)-(5.8c), (5.9a)-(5.9c), and (5.10a)-(5.10b). The resulting embedding shown in the left diagram of Fig. 5.3c is computed with default parameters. The VNE program generates redundancy for nodes A and C along with their respective paths to B due to constraints (5.10a)-(5.10b), which relate a node failure to the failure of adjacent links. For instance, the summed probabilities of the scenario with a failure of the leftmost communication link of index 3 and the scenario with a failure of the leftmost intermediate node are larger than the allowed maximum failure level of the virtual link A to B, leading to the generation of another copy of node A with another connecting link. However, entirely disjoint paths between primary and secondary mapping are unnecessary, given the default reliability requirements for the links. When we increase the reliability requirement of link (A, B) to $r^{AB} = 0.9$, the program generates a disjoint redundancy for that link. The resulting embedding is shown in the right diagram of Fig. 5.3c.

In the previous simulations, adaptation costs were not taken into account when updating parameter values, such as changing the reliability requirement of node A from $r^A = 0.8$ to $r^A = 0.95$ in Fig. 5.3a. To address this, we repeat the experiment with $\lambda^V = \lambda^L = 100$, resulting in the right diagram of Fig. 5.3d. Compared to the right diagram of Fig. 5.3a, both mappings of nodes A and C, as well as their respective paths towards node B, remain constant. The only difference is the addition of node redundancy for A and its path towards B due to the increased reliability requirement.

5.5.2 Smart Grid Voltage Control

The growing number of Distributed Energy Resources (DER) in the electric distribution grid can lead to critical voltage rises [38]. To address this issue, OLTC transformers are often used at the connection point of the medium and low voltage level networks. However, for this measure to be effective, a coordinated effort of the MV OLTCs is required [81]. In this section, we use our proposed SVNE algorithm to determine the cost-optimal communication scheme to monitor over-voltages at the MV grid, which can ensure high reliability despite potentially error-prone communication channels.

Smart Grid Setup

Smart Grids are an instance of the general principle of Cyber-Physical Systems (CPS) and is often defined as a set of three interconnected networks. The *physical network* comprises the tangible components of the distribution grid that are part of the control scheme. The

Table 5.2: Parameter values of compute & communication network.

	cost [US$/MB]	**avg. data rate** [bps]	**reliability**
PLC	0	10^4	85%
LoRaWAN	0	5.36	66%
GPRS	5	9×10^3	94%

compute and communication network consists of the compute resources needed to run the required software, both on field devices and in the control center, and the various communication media and infrastructure, such as gateways. This network acts as the SN in our case. The *logic network* holds the interconnected function blocks that implement the voltage control mechanism and acts as the VNR.

In the context voltage control with OLTCs, we define an exemplary Smart Grid with a single MV feeder and 15 OLTC transformers as its physical network. The compute and communication network is made up of three different communication technologies: PLC, LoRaWAN, and GPRS. Table 5.2 lists the relevant properties of the considered technologies. Substrate nodes have three different functions: LCUs make the interface between the OLTCs and the communication gateways. The gateways connect to a control center, which computes the output of the voltage control. We assume that the communication link between the gateways and the control center via a wired backbone network is perfectly reliable and has no data capacity constraints, allowing us to model the path between gateways and control center as a single link with maximum reliability and infinite data capacity.

The logical network consists of a central control function connected to a decentral logic. Note that the node mapping is fixed in this setup, where the central control function is mapped onto the control center and the decentral logic onto the LCUs. Considering a monitoring and/or control message each 5 min and a packet of 25 bytes size, the voltage control function requires at least 0.67 bps for each link between an LCU and the control center. The reliability requirement of the virtual links is determined by the probability of a voltage violation given the current grid state and the desired detection rate. For a graphical overview of the setup, see Fig. 5.1.

Probabilistic Power Flow Model

We use a probabilistic power flow model to determine the reliability requirement of virtual nodes by calculating the probability of voltage violation of each MV bus. In this model, we

only consider the active power flow and disregard any losses. To model uncertain demand and DER power production, we assume a normal distribution.

Let us consider an MV feeder with buses $\{0, 1, \ldots, N\}$, where bus 0 is the slack node. At time t, the active power in-feed at bus i is denoted as $p_i(t)$, which includes the load $d_i(t)$ and the photovoltaic generation $\phi_i(t)$ from the underlying low voltage grids. We model the loads as independent and identically distributed random variables $d_i(t) \sim \mathcal{N}(\mu_i^d(t), (\sigma_i^d)^2)$, while we assume that the photovoltaic production depends on a joint irradiation value for the entire MV feeder area. The photovoltaic production is denoted as $\phi_i(t) = c_i\phi^0(t)$, where $\phi^0(t) \sim N(\mu^\phi(t), (\sigma^\phi)^2(t))$ is the solar irradiation at time t and c_i the installed photovoltaic capacity at bus i.

Denoting $\mathbf{p} = [p_1, \ldots, p_N]^T$ as the vector of power injections at a certain time t, we have that $\mathbf{p} \sim N(\boldsymbol{\mu}^p, \boldsymbol{\Sigma}^p)$, where $\boldsymbol{\mu}^p = -\boldsymbol{\mu}^d(t) + \mu^\phi(t)\mathbf{1}$ and $\boldsymbol{\Sigma}^p = (\sigma^d)^2\mathbb{I} + \mathbf{1}\mathbf{1}^T(\sigma^\phi)^2$, with $\mathbf{c}$, being the vector of photovoltaic capacities.

Using a linearized power flow model, the active power flow f_i from bus i to its upstream neighbor $i-1$ can be calculated as

$$f_i = \sum_{j=i}^{N} p_j,\ i = 1, \ldots, N, \tag{5.18}$$

which can be written in matrix form as $\mathbf{f} = \mathbf{A}\mathbf{p}$, where $\mathbf{A}$ is an appropriately defined matrix and $\mathbf{f} = [f_1, \ldots, f_N]^T$.

Assuming a power factor close to one and denoting the voltage amplitude at bus i in p.u. by u_i, we can also express the power flow from bus i to its upstream neighbor $i-1$ as

$$\tilde{y}_i(u_i - u_{i-1}) = f_i,\ i = 1, \ldots, N, \tag{5.19}$$

where $\tilde{y}_i = \mathcal{R}(Y_i)u^{base}$, $\mathcal{R}(Y_i)$ is the real part of the admittance of the line i $\tilde{y}_i$, and u^{base} is the nominal voltage. We assume that $u_0 = 1$. (5.19) can be rewritten as $\mathbf{B}\mathbf{u} + \mathbf{v} = \mathbf{f}$, with $\mathbf{B}$ and $\mathbf{v}$ being an appropriately defined matrix and vector, respectively, and $\mathbf{u} = [u_1, \ldots, u_N]^T$.

By substituting (5.19) into (5.18), we have the following expression for the bus voltages $\mathbf{u}$:

$$\mathbf{u} = \mathbf{B}^{-1}\mathbf{A}\mathbf{p} - \mathbf{B}^{-1}\mathbf{v} \sim N(\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u), \tag{5.20}$$

where $\boldsymbol{\mu}^u = \mathbf{B}^{-1}\mathbf{A}\boldsymbol{\mu}^p - \mathbf{B}^{-1}\mathbf{v}$ and $\boldsymbol{\Sigma}^u = \mathbf{B}^{-1}\mathbf{A}\boldsymbol{\Sigma}^p\mathbf{A}^T\left(\mathbf{B}^{-1}\right)^T$.

We want to ensure that the probability of an undetected voltage violation at bus i remains below a certain threshold ϵ. This probability is the product of two factors: the

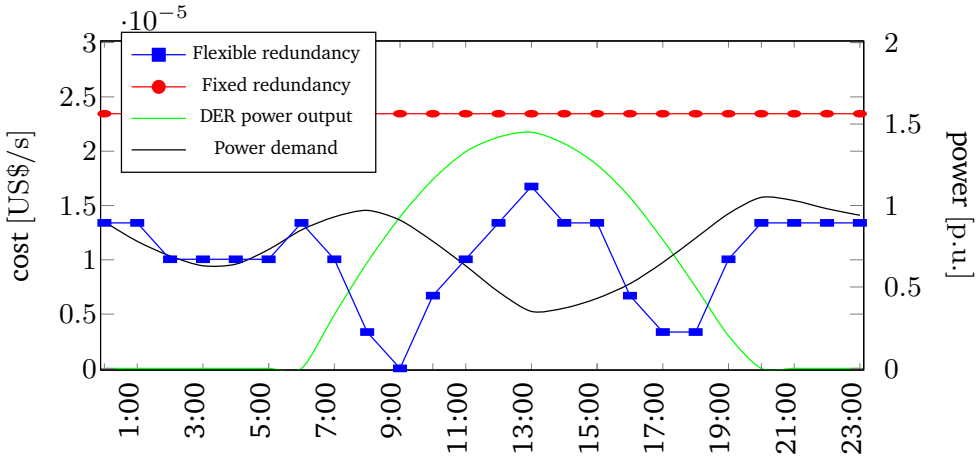

Figure 5.4: Cost comparison of the embeddings generated by proposed SVNE algorithm and fixed redundancy VNE programs as a function of the DER power output and power demand profiles throughout a certain day.

probability of the voltage u_i falling outside the allowed range, $P(u_i < \underline{U} \vee u_i > \overline{U})$, with $\underline{U}$ and $\overline{U}$ being the lower and upper limits of the allowed voltage range, and the probability of a failure in the communication link between the LCU associated to bus i and the control center. The latter is upper-bounded by $1 - r^{0,i}$, where $r^{0,i}$ is the required reliability of the link. This allows us to derive an expression for the reliability requirement of the virtual links in the VNR dependent on the uncertain bus voltages $\mathbf{u}$, which is given by:

$$r^{0,i} \geq 1 - \frac{\epsilon}{P(u_i < \underline{U} \vee u_i > \overline{U})}, \; i = 1, \ldots, N. \tag{5.21}$$

Experimental Results

In our simulation experiment, we use (5.21) to determine the VNR requirements for each bus in the network. We define the nominal voltage as $u^{base} = 33$ kV and set the real part of the admittance between any two buses to $\tilde{y}_i = 0.122$ S. The upper and lower limits of the allowed voltage range are 1.1 and 0.9 p.u., respectively. Powers are determined in p.u. relative to $p^{base} = 5$ MVA. The installed capacity c_i was set to 1 p.u. for every bus i in the feeder. The values for $\mu_i^d(t)$ and $\mu^\phi(t)$ are shown in Fig. 5.4 as the power demand and the DER power output, respectively. The variances of the power outputs are set to $(\sigma_i^d)^2 = 0.1$ and $(\sigma^\phi)^2 = 0.5$. The maximum probability of undetected voltage violations ϵ is set to 0.5%.

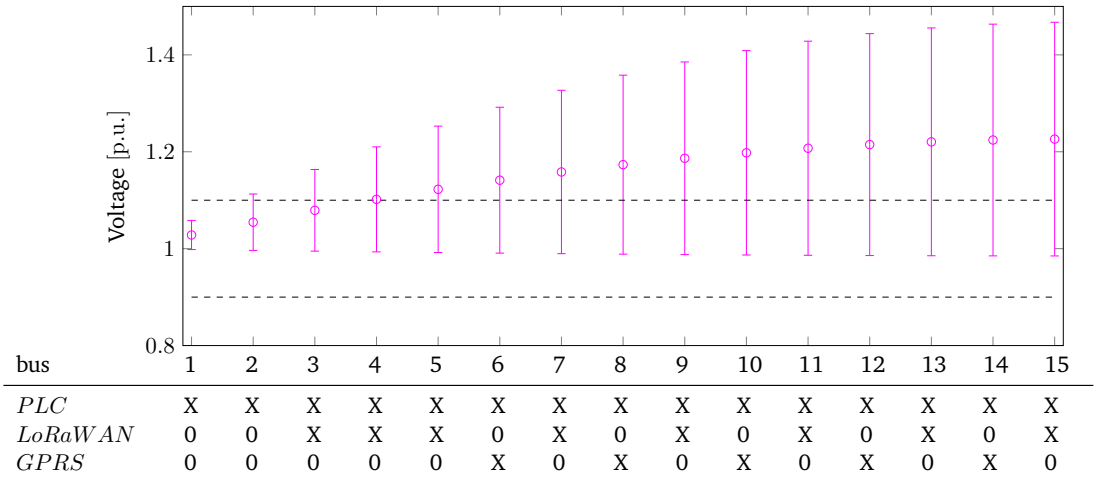

bus	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
PLC	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
$LoRaWAN$	0	0	X	X	X	0	X	0	X	0	X	0	X	0	X
$GPRS$	0	0	0	0	0	X	0	X	0	X	0	X	0	X	0

Figure 5.5: Voltage levels at 13:00 of the 15 buses with 95% (marginal) confidence interval. Below the diagram we show for each bus the communication technology utilized for connecting the associated LCU to the control center.

Fig. 5.5 shows the voltage levels of the 15 buses at 13:00 during maximum photovoltaic in-feed. Since the photovoltaic production exceeds the demand in expectation, the voltage levels rise on average over the length of the feeder. The exact voltage level for each bus depends on the realization of the uncertain demands and the irradiation factor.

It is observed that the voltage uncertainty increases over the length of the feeder. Moreover, the first two buses have a low probability of over-voltage and do not require communication link redundancy. However, as the distance from the slack node increases, a higher degree of redundancy is necessary to ensure reliable communication. Due to bandwidth limitations of the LoRaWAN gateway, not all nodes can use this cheap technology. Therefore, in places where LoRaWAN is not available, some nodes need to use the GPRS service as a backup.

Fig. 5.4 depicts the operational communication costs for the developed SVNE algorithm with a flexible redundancy approach and a fixed redundancy embedding plotted over a 24-hour period. Apart from periods of high DER generation and low demand, it is observed that the requirement for a high degree of redundancy also occurs during periods of low photovoltaic generation and increased energy demand to prevent under-voltage contingencies. In the times in between, demand and photovoltaic production are likely to balance each other out, and thus, less redundancy is required. Compared to the

fixed redundancy method, the flexible redundancy SVNE algorithm produces a cheaper embedding for every time step, with average savings of 57%. This value corresponds to a reduction of US$ 27/day or US$ 9,855/year of the operating cost for each MV feeder in the distribution grid.

5.6 Summary

This chapter presented a new SVNE algorithm that generates a flexible degree of redundancy for each VNR element. To this end, we formulated the algorithm as a robust MILP model, in which we considered single-element failure modes to create the uncertainty set. We also penalized large changes in consecutive embeddings by adding an adaptation cost to the objective function.

The developed algorithm was integrated into a voltage control scheme in a Smart Grid to reduce communication costs while guaranteeing high reliability services. The reliability requirement of the virtual links was set accordingly to a probabilistic model of the MV feeder's power flow. The experiment showed that the proposed program yields cheaper mappings than VNE with fixed degree of redundancy, by generating just enough redundancy for the requirements.

The described use case presented a timely application of Smart Grids, as more and more renewables need to be integrated into the distribution grids, and intelligent, communication-enhanced OLTC transformers are becoming more common. This is especially true for rural areas with long feeders, lots of renewable energy production, and often only patchy availability of cellular communication infrastructure. In this context, redundancy may be required for safe grid operation, as there is a high probability of violation of the voltage limits, and the few available communication links are often unreliable.

6 Conclusion and Outlook

Power and communication systems form the backbone of modern society's critical infrastructure, enabling numerous essential commercial and leisure activities. Given their growing interdependence, it is crucial that both systems function efficiently and effectively to provide high-quality services. In addition, optimizing the planning and operation of power grids and communication networks can make our infrastructure more reliable and less costly for system operators and end-users alike. However, as uncertainties in power and communication systems increase with the rise of renewable energy sources and the exponential growth of mobile devices connected to the Internet, respectively, traditional, deterministic optimization algorithms are becoming less effective or even obsolete. To address this challenge, a novel class of optimization algorithms capable of providing optimal solutions and quantifiable guarantees despite the high variance in the system's parameters and inputs is required.

In this thesis, we propose new adaptive algorithms for optimizing power and data flows and finding the optimal placement of flow-controlling devices. The solution of the algorithms is expressed as control laws that are linear-affine functions of the uncertainty in the network system. The proposed algorithms adopt a robust optimization framework, which guarantees system feasibility for any realization of uncertain factors while minimizing the cost of flow-controlling devices installation and network operation. Additionally, the algorithms consider deployability constraints of the underlying system in their solution search.

We prove the effectiveness of our proposed algorithms in different settings and compare their results with established solutions. In one setting, we demonstrate the framework's effectiveness in placing PSTs and optimizing redispatch-control laws in the presence of high uncertainty in the power generators' outputs in a transmission power system. In another scenario, we apply a similar approach to optimize routing in communication networks, showing that the robust approach outperforms two standard routing protocols regarding the average delay, dropping rate, and queue occupancy. For both use cases, we also develop heuristic algorithms that can solve the problem much faster than solvers with optimality guarantees. These methods are especially useful for analyzing large networks or when there are deployment constraints on how fast a solution must be available. Finally,

we apply the robust framework to a Smart Grid setting, where we optimize the power grid's communication arrangement to minimize the probability of undetected voltage violations.

The authors acknowledge that robust optimization in network systems remains an active research area. In the following, we present some approaches that can be used to improve the algorithms presented in this work and potential areas that could benefit from this work's methodology.

To enhance the robustness of network systems against extreme scenarios, both flow control laws and the placement of flow-controlling devices could be optimized while considering $N - k$ contingencies. However, scalability is a significant challenge when dealing with numerous simultaneous contingencies since the number of possible combinations of contingencies grows exponentially with k. A potential solution to improve scalability is to use graph reduction techniques to simplify the optimization problem's complexity. Establishing an equivalence between the original and reduced systems is crucial in this case. Additionally, one could aggregate highly correlated uncertain inputs, thus reducing the complexity of the uncertainty set and the number of active constraints of the optimization problem.

To reduce the conservativeness of the solution, one could consider different scenario-reduction techniques based on the probability of scenario occurrence. Such methods could decrease the uncertainty set size and provide more cost-efficient solutions. For instance, statistical information on the uncertain input could be used to eliminate unlikely scenarios or to extract important correlations between different uncertainty sources.

Moreover, the algorithms presented in this work could be extended to other power and data routing models and the placement of different flow-controlling devices. For instance, a similar approach to Chapter 3 could be developed for the placement of different FACTS devices, such as TCSC or SVC, with a corresponding linear power grid model. Furthermore, different metrics could be used to optimize data routing in communication networks with the scheme presented in Chapter 4, and applications such as outage detection could be mapped onto a Smart Grid using the SVNE approach of Chapter 5.

Lastly, the proposed algorithms could be applied to different sorts of network systems, such as heat, transportation, and logistic networks. Although the physics in these systems differ from power and communication networks, deploying this thesis' algorithms is still possible as long as the constraints on the system's flows can be expressed as linear dependencies on the uncertain inputs.

Curriculum Vitae

08/23 –	**Energy System Architect** HIF EMEA GmbH, Germany
02/20 – 07/23	**Research Associate and Ph.D. Candidate** Energy Information Networks & Systems, Technical University of Darmstadt, Germany
09/17 – 08/19	**M.Sc. in Electrical Engineering and Information Technology (Double-Degree)** Technical University of Darmstadt, Germany
04/18 – 07/19	**Research Assistant** Energy Information Networks & Systems, Technical University of Darmstadt, Germany
08/16 – 08/17	**Research Assistant** Integrated Systems Laboratory, University of São Paulo, Brazil
01/14 – 12/19	**Diploma in Electrical Engineering and Information Technology (Double-Degree)** University of São Paulo, Brazil

Publications

Journal Publications

- Allan Santos and Florian Steinke. “Robust placement and control of phase-shifting transformers considering redispatch measures.” In: *Energies* 16.11 (2023), p. 4438.
- Allan Santos, Amr Rizk, and Florian Steinke. “Adaptive global coordination of local routing policies for communication networks.” In: *Computer Communications* 204 (2023), pp. 101–108.
- Sara Mollaeivaneghi, Allan Santos, and Florian Steinke. “A new approach to quantile estimation for probabilistic economic dispatch.” In: *IEEE Transactions on Power Systems* (submitted).

International Conference Proceedings

- Johannes Sindt, Allan Santos, Marc E. Pfetsch, and Florian Steinke. “Evaluation of multiparametric linear programming for economic dispatch under uncertainty.” In: *IEEE PES Innovative Smart Grid Technologies Europe* (ISGT Europe) (2021), pp. 1-5.
- Allan Santos, Edwin Mora, Jan Peters, and Florian Steinke. “Decentralized data-driven tuning of droop frequency controllers applied to a non-linear grid model.” In: *PESS IEEE Power and Energy Student Summit* (2020), pp. 1-5.
- Allan Santos, Edwin Mora, Jan Peters, and Florian Steinke. “Decentralized data-driven tuning of droop frequency controllers.” *IEEE PES Innovative Smart Grid Technologies Europe* (ISGT Europe) (2020), pp. 141-145.
- Allan Santos, Amr Rizk, and Florian Steinke. “Flexible redundancy generation for virtual network embedding with an application to smart grids.” In: *Proceedings of the Eleventh ACM International Conference on Future Energy Systems* (2020), pp. 97–105.

Supervised Works

Master's Theses

- João Kröger. "Zero-shot learning optimal power flow with graph neural networks." 2023.
- Tim Frieß. "Data-based model identification for the nonlinear optimization of industrial cooling systems." 2023.
- Jiali Huang. "Modeling uncertainty for flexibility management in power distribution systems." 2022.
- Johannes Sindt. "Uncertainty quantification in energy system design." 2020.

Bachelor's Theses

- Konstantin Preusser. "Robust placement of phase-shifting transformers considering N-1 contingencies." 2023.

Project Seminars

- Konstantin Preußer. "Robust placement of phase-shifting transformers considering N-1 contingencies." 2022.
- Isabella Grieser. "Virtual power plant formation algorithm." 2021.
- René Prinz. "Simulation environment of smart grid communication networks." 2020.

Bibliography

[1] Farzaneh Abdollahi and Kash Khorasani. "A novel H_∞ control strategy for design of a robust dynamic routing algorithm in traffic networks." In: *IEEE Journal on Selected Areas in Communications* 26.4 (2008), pp. 706–718.

[2] Bernardetta Addis, Antonio Capone, Giuliana Carello, Luca G. Gianoli, and Brunilde Sansò. "Energy management through optimized routing and device powering for greener communication networks." In: *IEEE/ACM Transactions on Networking* 22.1 (2013), pp. 313–325.

[3] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[4] Ravindra K. Ahuja, Thomas L. Magnanti, James B. Orlin, and M. R. Reddy. "Applications of network optimization." In: *Handbooks in Operations Research and Management Science* 7 (1995), pp. 1–83.

[5] Teodoro Alamo, Roberto Tempo, and Amalia Luque. "On the sample complexity of randomized approaches to the analysis and design under uncertainty." In: *Proceedings of American Control Conference* (2010), pp. 4671–4676.

[6] Mohammad Alizadeh, Tom Edsall, Sarang Dharmapurikar, Ramanan Vaidyanathan, Kevin Chu, Andy Fingerhut, Vinh The Lam, Francis Matus, Rong Pan, Navindra Yadav, et al. "CONGA: Distributed congestion-aware load balancing for datacenters." In: *Proceedings of ACM Conference on SIGCOMM* (2014), pp. 503–514.

[7] Dawei Bai, Tamra Carpenter, and John Mulvey. "Making a case for robust optimization models." In: *Management science* 43.7 (1997), pp. 895–907.

[8] The World Bank. *World Development Indicators*. 2022. URL: `https://datacatalog.worldbank.org/search/dataset/0037712/World-Development-Indicators` (visited on 09/28/2022).

[9] Thomas Bauschert, Christina Büsing, Fabio D'Andreagiovanni, Arie M. C. A. Koster, Manuel Kutschka, and Uwe Steglich. "Network planning under demand uncertainty with robust optimization." In: *IEEE Communications Magazine* 52.2 (2014), pp. 178–185.

[10] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Vol. 28. Princeton university press, 2009.

[11] Aharon Ben-Tal and Arkadi Nemirovski. "Robust convex optimization." In: *Mathematics of Operations Research* 23.4 (1998), pp. 769–805.

[12] Aharon Ben-Tal and Arkadi Nemirovski. "Robust solutions of uncertain linear programs." In: *Operations Research Letters* 25.1 (1999), pp. 1–13.

[13] Aharon Ben-Tal, Arkadi Nemirovski, and Cees Roos. "Robust solutions of uncertain quadratic and conic-quadratic problems." In: *SIAM Journal on Optimization* 13.2 (2002), pp. 535–560.

[14] Dimitris Bertsimas, David B. Brown, and Constantine Caramanis. "Theory and applications of robust optimization." In: *SIAM Review* 53.3 (2011), pp. 464–501.

[15] Dimitris Bertsimas, Eugene Litvinov, Xu Andy Sun, Jinye Zhao, and Tongxin Zheng. "Adaptive robust optimization for the security constrained unit commitment problem." In: *IEEE Transactions on Power Systems* 28.1 (2012), pp. 52–63.

[16] Dimitris Bertsimas and Melvyn Sim. "The price of robustness." In: *Operations Research* 52.1 (2004), pp. 35–53.

[17] Daniel Bienstock, Michael Chertkov, and Sean Harnett. "Chance-constrained optimal power flow: Risk-aware network control under uncertainty." In: *Siam Review* 56.3 (2014), pp. 461–495.

[18] John R. Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

[19] Salim Bitam and Abdelhamid Mellouk. "Bee life-based multi constraints multicast routing optimization for vehicular ad hoc networks." In: *Journal of Network and Computer Applications* 36.3 (2013), pp. 981–991.

[20] Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick McKeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, George Varghese, et al. "P4: Programming protocol-independent packet processors." In: *ACM SIGCOMM Computer Communication Review* 44.3 (2014), pp. 87–95.

[21] Juan Felipe Botero, Xavier Hesselbach, Michael Duelli, Daniel Schlosser, Andreas Fischer, and Hermann De Meer. “Energy efficient virtual network embedding.” In: *IEEE Communications Letters* 16.5 (2012), pp. 756–759.

[22] Zhiping Cai, Fang Liu, Nong Xiao, Qiang Liu, and Zhiying Wang. “Virtual network embedding for evolving networks.” In: *IEEE Global Telecommunications Conference GLOBECOM* (2010), pp. 1–5.

[23] Rajat Chaudhary, Gagangeet Singh Aujla, Sahil Garg, Neeraj Kumar, and Joel J. P. C. Rodrigues. “SDN-enabled multi-attribute-based secure communication for smart grid in IIoT environment.” In: *IEEE Transactions on Industrial Informatics* 14.6 (2018), pp. 2629–2640.

[24] Mosharaf Chowdhury, Muntasir Raihan Rahman, and Raouf Boutaba. “Vineyard: Virtual network embedding algorithms with coordinated node and link mapping.” In: *IEEE/ACM Transactions on Networking* 20.1 (2011), pp. 206–219.

[25] Shihabur Rahman Chowdhury, Reaz Ahmed, Md Mashrur Alam Khan, Nashid Shahriar, Raouf Boutaba, Jeebak Mitra, and Feng Zeng. “Dedicated protection for survivable virtual network embedding.” In: *IEEE Transactions on Network and Service Management* 13.4 (2016), pp. 913–926.

[26] Cisco. *White paper: Cisco annual internet report, 2018–2023*. Tech. rep. 2019.

[27] Cisco. *White paper: Cisco visual networking index: global mobile data traffic forecast update, 2017–2022*. Tech. rep. 2019.

[28] IBM ILOG Cplex. “V12. 1: User’s Manual for CPLEX.” In: *International Business Machines Corporation* 46.53 (2009), p. 157.

[29] John D’Ambrosia, David Law, and Mark Nowell. “40 Gigabit Ethernet and 100 Gigabit Ethernet Technology Overview.” In: *Ethernet Alliance White Paper* (2008).

[30] Advait Dixit, Pawan Prakash, Y. Charlie Hu, and Ramana Rao Kompella. “On the impact of packet spraying in data center networks.” In: *Proceedings of IEEE INFOCOM* (2013), pp. 2130–2138.

[31] Hermann W. Dommel and William F. Tinney. “Optimal power flow solutions.” In: *IEEE Transactions on Power Apparatus and Systems* 10 (1968), pp. 1866–1876.

[32] Xinshu Dong, Hui Lin, Rui Tan, Ravishankar K. Iyer, and Zbigniew Kalbarczyk. “Software-defined networking for smart grid resilience: Opportunities and challenges.” In: *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security* (2015), pp. 61–68.

[33] Chao Duan, Wanliang Fang, Lin Jiang, and Shuanbao Niu. “FACTS devices allocation via sparse optimization.” In: *IEEE Transactions on Power Systems* 31.2 (2015), pp. 1308–1319.

[34] Iain Dunning, Joey Huchette, and Miles Lubin. “JuMP: A modeling language for mathematical optimization.” In: *SIAM Review* 59.2 (2017), pp. 295–320. DOI: 10.1137/15M1020575.

[35] Abdel-Aty Edris. “Proposed terms and definitions for flexible AC transmission system (FACTS).” In: *IEEE Transactions on Power Delivery* 12.4 (1997), pp. 1848–1853.

[36] Laurent El Ghaoui and Hervé Lebret. “Robust solutions to least-squares problems with uncertain data.” In: *SIAM Journal on Matrix Analysis and Applications* 18.4 (1997), pp. 1035–1064.

[37] Energinet.dk. *Production and Consumption Data*. 2021. URL: https://www.energidataservice.dk/collections/production-and-consumption (visited on 09/06/2021).

[38] Nicholas Etherden and Math H. J. Bollen. “Overload and overvoltage in low-voltage and medium-voltage networks due to renewable energy–some illustrative case studies.” In: *Electric Power Systems Research* 114 (2014), pp. 39–48.

[39] SolarPower Europe. *Global market outlook for solar power, 2021–2025*. Tech. rep. 2021.

[40] Shimon Even, Alon Itai, and Adi Shamir. “On the complexity of time table and multi-commodity flow problems.” In: *16th Annual Symposium on Foundations of Computer Science* (1975), pp. 184–193.

[41] Frank J. Fabozzi, Sergio M. Focardi, Petter N. Kolm, and Dessislava A. Pachamanova. *Robust portfolio optimization and management*. John Wiley & Sons, 2007.

[42] Jamal Faraji, Hamed Hashemi-Dezaki, and Abbas Ketabi. “Stochastic operation and scheduling of energy hub considering renewable energy sources’ uncertainty and N-1 contingency.” In: *Sustainable Cities and Society* 65 (2021), p. 102578.

[43] Mohammad Al-Fares, Sivasankar Radhakrishnan, Barath Raghavan, Nelson Huang, Amin Vahdat, et al. “Hedera: Dynamic flow scheduling for data center networks.” In: *Proceedings of USENIX NSDI* (2010), pp. 89–92.

[44] Maryam Fazel and Mung Chiang. “Network utility maximization with nonconcave utilities using sum-of-squares method.” In: *Proceedings of the 44th IEEE Conference on Decision and Control* (2005), pp. 1867–1874.

[45] Andreas Fischer, Juan Felipe Botero, Michael Till Beck, Hermann De Meer, and Xavier Hesselbach. “Virtual network embedding: A survey.” In: *IEEE Communications Surveys & Tutorials* 15.4 (2013), pp. 1888–1906.

[46] Vladimir Frolov, Priyanko Guha Thakurta, Scott Backhaus, Janusz Bialek, and Michael Chertkov. “Operations- and uncertainty-aware installation of FACTS devices in a large transmission system.” In: *IEEE Transactions on Control of Network Systems* 6.3 (2019), pp. 961–970.

[47] Sam Gao, Mark Handley, and Stefano Vissicchio. “Stats 101 in P4: Towards in-switch anomaly detection.” In: *Proceedings of the 20th ACM Workshop on Hot Topics in Networks* (2021), pp. 84–90.

[48] Sijia Geng, Maria Vrakopoulou, and Ian A. Hiskens. “Chance-constrained optimal capacity design for a renewable-only islanded microgrid.” In: *Electric Power Systems Research* 189 (2020), p. 106564.

[49] Stéphane Gerbex, Rachid Cherkaoui, and Alain J. Germond. “Optimal location of multi-type FACTS devices in a power system by means of genetic algorithms.” In: *IEEE Transactions on Power Systems* 16.3 (2001), pp. 537–544.

[50] Soudeh Ghorbani, Zibin Yang, P. Brighten Godfrey, Yashar Ganjali, and Amin Firoozshahian. “Drill: Micro load balancing for low-latency data center networks.” In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication* (2017), pp. 225–238.

[51] Long Gong and Zuqing Zhu. “Virtual optical network embedding (VONE) over elastic optical networks.” In: *Journal of Lightwave Technology* 32.3 (2013), pp. 450–460.

[52] Natasha Gude, Teemu Koponen, Justin Pettit, Ben Pfaff, Martın Casado, Nick McKeown, and Scott Shenker. “NOX: Towards an operating system for networks.” In: *ACM SIGCOMM Computer Communication Review* 38.3 (July 2008), pp. 105–110.

[53] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2021. URL: `https://www.gurobi.com` (visited on 05/30/2021).

[54] Bo Han, Vijay Gopalakrishnan, Lusheng Ji, and Seungjoon Lee. “Network function virtualization: Challenges and opportunities for innovations.” In: *IEEE Communications Magazine* 53.2 (2015), pp. 90–97.

[55] Rhaban Hark, Nieke Aerts, David Hock, Nils Richerzhagen, Amr Rizk, and Ralf Steinmetz. “Reducing the monitoring footprint on controllers in software-defined networks.” In: *IEEE Transactions on Network and Service Management* 15.4 (2018), pp. 1264–1276.

[56] Frederik Hauser, Marco Häberle, Daniel Merling, Steffen Lindner, Vladimir Gurevich, Florian Zeiger, Reinhard Frank, and Michael Menth. “A survey on data plane programming with p4: Fundamentals, advances, and applied research.” In: *Journal of Network and Computer Applications* 212 (2023), p. 103561.

[57] Miao He and Junshan Zhang. “A dependency graph approach for fault detection and localization towards secure smart grid.” In: *IEEE Transactions on Smart Grid* 2.2 (2011), pp. 342–351.

[58] Mu He, Amir Varasteh, and Wolfgang Kellerer. “Toward a flexible design of SDN dynamic control plane: An online optimization approach.” In: *IEEE Transactions on Network and Service Management* 16.4 (2019), pp. 1694–1708.

[59] Sandra Herker, Ashiq Khan, and Xueli An. “Survey on survivable virtual network embedding problem and solutions.” In: *International Conference on Networking and Services, ICNS* (2013), pp. 99–104.

[60] Rabih A. Jabr. “Robust transmission network expansion planning with uncertain renewable generation and loads.” In: *IEEE Transactions on Power Systems* 28.4 (2013), pp. 4558–4567.

[61] Rabih A. Jabr, Izudin Džafić, and Bikash C. Pal. “Robust optimization of storage investment on transmission networks.” In: *IEEE Transactions on Power Systems* 30.1 (2014), pp. 531–539.

[62] Raj Jain and Subharthi Paul. “Network virtualization and software defined networking for cloud computing: A survey.” In: *IEEE Communications Magazine* 51.11 (2013), pp. 24–31.

[63] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, et al. “B4: Experience with a globally-deployed software defined WAN.” In: *ACM SIGCOMM Computer Communication Review* 43.4 (2013), pp. 3–14.

[64] Stacy L. Janak, Xiaoxia Lin, and Christodoulos A. Floudas. “A new robust optimization approach for scheduling under uncertainty: II. Uncertainty with known probability distribution.” In: *Computers & Chemical Engineering* 31.3 (2007), pp. 171–195.

[65] Dong Jin, Zhiyi Li, Christopher Hannon, Chen Chen, Jianhui Wang, Mohammad Shahidehpour, and Cheol Won Lee. “Toward a cyber resilient and secure microgrid using software-defined networking.” In: *IEEE Transactions on Smart Grid* 8.5 (2017), pp. 2494–2504.

[66] T. Karagiannis, M. Molle, M. Faloutsos, and A. Broido. “A nonstationary Poisson view of Internet traffic.” In: *Proceedings of IEEE INFOCOM* (2004), pp. 1558–1569.

[67] Naga Katta, Mukesh Hira, Changhoon Kim, Anirudh Sivaraman, and Jennifer Rexford. “HULA: Scalable load balancing using programmable data planes.” In: *Proceedings of the Symposium on SDN Research* (2016), pp. 1–12.

[68] Leonid G. Khachiyan. “Rounding of polytopes in the real number model of computation.” In: *Mathematics of Operations Research* 21.2 (1996), pp. 307–320.

[69] Md Mashrur Alam Khan, Nashid Shahriar, Reaz Ahmed, and Raouf Boutaba. “Multi-path link embedding for survivability in virtual networks.” In: *IEEE Transactions on Network and Service Management* 13.2 (2016), pp. 253–266.

[70] Abdallah Khreishah, Haythem Bany Salameh, Issa Khalil, and Ammar Gharaibeh. “Renewable energy-aware joint caching and routing for green communication networks.” In: *IEEE Systems Journal* 12.1 (2016), pp. 768–777.

[71] Håkon Kile, Kjetil Uhlen, Leif Warland, and Gerd Kjølle. “A comparison of AC and DC power flow models for contingency and reliability analysis.” In: *Power Systems Computation Conference* (2014), pp. 1–7.

[72] Péter Kovács. “Minimum-cost flow algorithms: an experimental evaluation.” In: *Optimization Methods and Software* 30.1 (2015), pp. 94–127.

[73] Mikhail K. Kozlov, Sergei P. Tarasov, and Leonid G. Khachiyan. “Polynomial solvability of convex quadratic programming.” In: *Doklady Akademii Nauk* 248.5 (1979), pp. 1049–1051.

[74] Prabha S. Kundur and Om P. Malik. *Power system stability and control*. McGraw-Hill Education, 2022.

[75] Dara Kusic, Jeffrey O. Kephart, James E Hanson, Nagarajan Kandasamy, and Guofei Jiang. “Power and performance management of virtualized computing environments via lookahead control.” In: *Cluster computing* 12 (2009), pp. 1–15.

[76] Stephen C. H. Leung, Sally O. S. Tsang, Wan-Lung Ng, and Yue Wu. “A robust optimization model for multi-site production planning problem in an uncertain environment.” In: *European Journal of Operational Research* 181.1 (2007), pp. 224–238.

[77] Husheng Li, Lifeng Lai, and H. Vincent Poor. “Multicast routing for decentralized control of cyber physical systems with an application in smart grid.” In: *IEEE Journal on Selected Areas in Communications* 30.6 (2012), pp. 1097–1107.

[78] Tjing T. Lie and Wanhong Deng. “Optimal flexible AC transmission systems (FACTS) devices allocation.” In: *International Journal of Electrical Power & Energy Systems* 19.2 (1997), pp. 125–134.

[79] Flavio G. M. Lima, Francisco D. Galiana, Ivana Kockar, and Jorge Munoz. “Phase shifter placement in large-scale systems via mixed integer linear programming.” In: *IEEE Transactions on Power Systems* 18.3 (2003), pp. 1029–1034.

[80] Haoming Liu, Xingying Chen, Kun Yu, and Yunhe Hou. “The control and analysis of self-healing urban power grid.” In: *IEEE Transactions on Smart Grid* 3.3 (2012), pp. 1119–1129.

[81] Xiaohu Liu, Andreas Aichhorn, Liming Liu, and Hui Li. “Coordinated control of distributed energy storage system with tap changer transformers for voltage rise mitigation under high photovoltaic penetration.” In: *IEEE Transactions on Smart Grid* 3.2 (2012), pp. 897–906.

[82] Chenyang Lu, Abusayeed Saifullah, Bo Li, Mo Sha, Humberto Gonzalez, Dolvara Gunatilaka, Chengjie Wu, Lanshun Nie, and Yixin Chen. “Real-time wireless sensor-actuator networks for industrial cyber-physical systems.” In: *Proceedings of the IEEE* 104.5 (2015), pp. 1013–1024.

[83] Sara Lumbreras and Andrés Ramos. “The new challenges to transmission expansion planning. Survey of recent practice and literature review.” In: *Electric Power Systems Research* 134 (2016), pp. 19–29.

[84] Pin Lv, Xudong Wang, Yang Yang, and Ming Xu. “Network virtualization for smart grid communications.” In: *IEEE Systems Journal* 8.2 (2013), pp. 471–482.

[85] Kostas Margellos, Paul Goulart, and John Lygeros. “On the road between robust optimization and the scenario approach for chance constrained optimization problems.” In: *IEEE Transactions on Automatic Control* 59.8 (2014), pp. 2258–2263.

[86] Felipe Mata, José Luis Garcıa-Dorado, and Javier Aracil. “Detection of traffic changes in large-scale backbone networks: The case of the Spanish academic network.” In: *Computer Networks* 56.2 (2012), pp. 686–702.

[87] Stuart Mitchell, Michael O’Sullivan, and Iain Dunning. *PuLP: A linear programming toolkit for python*. 2011. URL: `http://www.optimization-online.org/DB_FILE/2011/09/3178.pdf` (visited on 01/15/2020).

[88] Roberto Mınguez, Federico Milano, Rafael Zárate-Miñano, and Antonio J. Conejo. “Optimal network placement of SVC devices.” In: *IEEE Transactions on Power Systems* 22.4 (2007), pp. 1851–1860.

[89] Elias Molina, Eduardo Jacob, Jon Matias, Naiara Moreira, and Armando Astarloa. “Using software defined networking to manage and control IEC 61850-based systems.” In: *Computers & Electrical Engineering* 43 (2015), pp. 142–154.

[90] Edwin Mora and Florian Steinke. “On the minimal set of controllers and sensors for linear power flow.” In: *Electric Power Systems Research* 190 (2021), p. 106647.

[91] Tsan Sheng Ng, Yang Sun, and John Fowler. “Semiconductor lot allocation using robust optimization.” In: *European Journal of Operational Research* 205.3 (2010), pp. 557–570.

[92] Elis Nycander, Lennart Söder, Jon Olauson, and Robert Eriksson. “Curtailment analysis for the Nordic power system considering transmission capacity, inertia limits and generation flexibility.” In: *Renewable Energy* 152 (2020), pp. 942–960.

[93] Open Networking Fundation. “Software-defined networking: The new norm for networks.” In: *ONF White Paper* (2012).

[94] Thomas J. Overbye, Xu Cheng, and Yan Sun. “A comparison of the AC and DC power flow models for LMP calculations.” In: *Proceedings of the 37th Annual Hawaii International Conference on System Sciences* (2004), pp. 1–9.

[95] BP p.l.c. *Statistical Review of World Energy*. 2022. URL: `https://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.html` (visited on 09/27/2022).

[96] Panagiotis Papadimitriou, Olaf Maennel, Adam Greenhalgh, Anja Feldmann, and Laurent Mathy. “Implementing network virtualization for a future internet.” In: *20th ITC Specialist Seminar on Network Virtualization Proceedings* (2009).

[97] Jonathan Perry, Amy Ousterhout, Hari Balakrishnan, Devavrat Shah, and Hans Fugal. “Fastpass: A centralized “zero-queue” datacenter network.” In: *SIGCOMM Computer Communication Review* 44.4 (2014), pp. 307–318.

[98] Michal Pióro and Deep Medhi. *Routing, flow, and capacity design in communication and computer networks*. Elsevier, 2004.

[99] Lynn Powell. *Power system load flow analysis*. McGraw Hill Professional, 2004.

[100] András Prékopa. *Stochastic programming*. Vol. 324. Springer Science & Business Media, 2013.

[101] J. Ptacek, P. Modlitba, S. Vnoucek, and J. Cermak. “Possibilities of applying phase shifting transformers in the electric power system of the Czech Republic.” In: *CIGRE Session* (2006), pp. C2–203.

[102] Christopher Rackauckas and Qing Nie. “Differentialequations.jl–a performant and feature-rich ecosystem for solving differential equations in Julia.” In: *Journal of Open Research Software* 5.1 (2017).

[103] Line Roald, Sidhant Misra, Thilo Krause, and Göran Andersson. “Corrective control to handle forecast uncertainty: A chance constrained optimal power flow.” In: *IEEE Transactions on Power Systems* 32.2 (2016), pp. 1626–1637.

[104] Matthew Roughan, Albert Greenberg, Charles Kalmanek, Michael Rumsewicz, Jennifer Yates, and Yin Zhang. “Experience in measuring internet backbone traffic variability: Models metrics, measurements and meaning.” In: *Proceedings of the 18th International Teletraffic Congress* 5 (2003), pp. 379–388.

[105] Allan Santos, Amr Rizk, and Florian Steinke. “Adaptive global coordination of local routing policies for communication networks.” In: *Computer Communications* 204 (2023), pp. 101–108.

[106] Allan Santos, Amr Rizk, and Florian Steinke. “Flexible redundancy generation for virtual network embedding with an application to smart grids.” In: *Proceedings of the Eleventh ACM International Conference on Future Energy Systems* (2020), pp. 97–105.

[107] Allan Santos and Florian Steinke. “Robust placement and control of phase-shifting transformers considering redispatch measures.” In: *Energies* 16.11 (2023), p. 4438.

[108] M. Saravanan, S. Mary Raja Slochanal, P. Venkatesh, and Prince Stephen Abraham. “Application of PSO technique for optimal location of FACTS devices considering system loadability and cost of installation.” In: *2005 International Power Engineering Conference* (2005), pp. 716–721.

[109] Adrian Segall. “The modeling of adaptive routing in data-communication networks.” In: *IEEE Transactions on Communications* 25.1 (1977), pp. 85–95.

[110] Nashid Shahriar, Reaz Ahmed, Shihabur Rahman Chowdhury, Aimal Khan, Raouf Boutaba, and Jeebak Mitra. “Generalized recovery from node failure in virtual network embedding.” In: *IEEE Transactions on Network and Service Management* 14.2 (2017), pp. 261–274.

[111] Ashwani Sharma, Saurabh Chanana, and Sanjoy Parida. “Combined optimal location of FACTS controllers and loadability enhancement in competitive electricity markets using MILP.” In: *IEEE Power Engineering Society General Meeting* (2005), pp. 670–677.

[112] Shan Sinha, Srikanth Kandula, and Dina Katabi. “Harnessing TCP’s burstiness with flowlet switching.” In: *Proceedings of 3rd ACM Workshop on Hot Topics in Networks (Hotnets-III)* (2004).

[113] Suvrit Sra, Sebastian Nowozin, and Stephen J. Wright. *Optimization for machine learning*. Mit Press, 2012.

[114] Goran Strbac. “Demand side management: Benefits and challenges.” In: *Energy policy* 36.12 (2008), pp. 4419–4426.

[115] Wenjing Su, Chunyu Liu, Constantino M. Lagoa, Hao Che, Ke Xu, and Yong Cui. “Integrated, distributed traffic control in multidomain networks.” In: *IEEE Transactions on Control Systems Technology* 23.4 (2014), pp. 1373–1386.

[116] Julija Tastu, Pierre Pinson, Pierre-Julien Trombe, and Henrik Madsen. “Probabilistic forecasts of wind power generation accounting for geographically dispersed information.” In: *IEEE Transactions on Smart Grid* 5.1 (2013), pp. 480–489.

[117] Maria Vrakopoulou, Kostas Margellos, John Lygeros, and Göran Andersson. “A probabilistic framework for reserve scheduling and N-1 security assessment of systems with high wind power penetration.” In: *IEEE Transactions on Power Systems* 28.4 (2013), pp. 3885–3896.

[118] Jingyao Wang, Mahmoud Ashour, Constantino M. Lagoa, Necdet Serhat Aybat, and Hao Che. “A fully distributed traffic allocation algorithm for nonconcave utility maximization in connectionless communication networks.” In: *Automatica* 109 (2019), p. 108506.

[119] Ning Wang, Kin Hon Ho, George Pavlou, and Michael Howarth. “An overview of routing optimization for internet traffic engineering.” In: *IEEE Communications Surveys & Tutorials* 10.1 (2008), pp. 36–56.

[120] WindEurope. *Wind energy in Europe: 2021 Statistics and the outlook for 2022–2026*. Tech. rep. 2022.

[121] Svante Wold, Kim Esbensen, and Paul Geladi. “Principal component analysis.” In: *Chemometrics and Intelligent Laboratory Systems* 2.1-3 (1987), pp. 37–52.

[122] Xu Xu, Jian Zhao, Zhao Xu, Songjian Chai, Jiayong Li, and Yi Yu. “Stochastic optimal TCSC placement in power system considering high wind power penetration.” In: *IET Generation, Transmission & Distribution* 12.12 (2018), pp. 3052–3060.

[123] Zhongxia Yan, Jingguo Ge, Yulei Wu, Liangxiong Li, and Tong Li. “Automatic virtual network embedding: A deep reinforcement learning approach with graph convolutional networks.” In: *IEEE Journal on Selected Areas in Communications* 38.6 (2020), pp. 1040–1057.

[124] Kai Yang, Yihong Wu, Jianwei Huang, Xiaodong Wang, and Sergio Verdú. “Distributed robust optimization for communication networks.” In: *Proceedings of IEEE INFOCOM* (2008), pp. 1157–1165.

[125] Jianchao Zhang, Boon-Chong Seet, Tek-Tjing Lie, and Chuan Heng Foh. “Opportunities for software-defined networking in smart grid.” In: *9th International Conference on Information, Communications & Signal Processing* (2013), pp. 1–5.

[126] Muhong Zhang. “Two-stage minimax regret robust uncapacitated lot-sizing problems with demand uncertainty.” In: *Operations Research Letters* 39.5 (2011), pp. 342–345.

[127] Xiao-Ping Zhang, Christian Rehtanz, and Bikash Pal. *Flexible AC transmission systems: modelling and control*. Springer Science & Business Media, 2012.

[128] Xiaohu Zhang, Di Shi, Zhiwei Wang, Bo Zeng, Xinan Wang, Kevin Tomsovic, and Yanming Jin. “Optimal allocation of series FACTS devices under high penetration of wind power within a market environment.” In: *IEEE Transactions on Power Systems* 33.6 (2018), pp. 6206–6217.

[129] Zhongbao Zhang, Xiang Cheng, Sen Su, Yiwen Wang, Kai Shuang, and Yan Luo. “A unified enhanced particle swarm optimization-based virtual network embedding algorithm.” In: *International Journal of Communication Systems* 26.8 (2013), pp. 1054–1073.

[130] Kemin Zhou, J. C. Doyle, and Keither Glover. *Robust and optimal control*. Prentice hall, 1996.

[131] Ray Daniel Zimmerman, Carlos Edmundo Murillo-Sánchez, and Robert John Thomas. “MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education.” In: *IEEE Transactions on Power Systems* 26.1 (2010), pp. 12–19.