

An Integrative Approach to Data-Driven Monitoring and Control of Electric Distribution Networks

Roel Dobbe



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2018-45

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-45.html>

May 10, 2018

Copyright © 2018, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

**An Integrative Approach to Data-Driven Monitoring and Control
of Electric Distribution Networks**

by

Roel Ignatius Jacobus Dobbe

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Claire J. Tomlin, Chair
Professor Duncan Callaway
Professor Alexandra von Meier

Spring 2018

**An Integrative Approach to Data-Driven Monitoring and Control
of Electric Distribution Networks**

Copyright 2018
by
Roel Ignatius Jacobus Dobbe

Abstract

An Integrative Approach to Data-Driven Monitoring and Control
of Electric Distribution Networks

by

Roel Ignatius Jacobus Dobbe

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Claire J. Tomlin, Chair

The commodification of computing, sensors, actuators, data storage and algorithms has unleashed a new wave of automation throughout society. Motivated by the promise of new capabilities, quality improvements, or efficiency gains, data-driven technologies have captured the attention and imagination of the public and many domain experts. Though opportunities are ample, the rapid introduction of data-driven functionality also triggers well-founded concerns about safeguarding critical values, such as safety, privacy and justice. In the context of operating electric distribution networks, the need for data-driven monitoring and control is explained by the irreversible transition from fossil to renewable generation and the accompanied electrification of our economy in areas like transportation and heating. The traditional fit-and-forget paradigm of designing networks conservatively for the projected peak loads assumed unidirectional power flow, predictable future demand and monotonic voltage drops, and allowed for operating at near-100% reliability with minimal requirement for sensing and actuation. The intermittent nature of Distributed Generation (DG), its ability to feed power back to the grid and cause bidirectional power flow, and the diversifying and nonlinear behavior of electric loads are all eating away at the robustness of this approach, causing Distribution System Operators (DSOs) to put caps on the allowable DG and revisit their design and operating practice. Rather than making traditional expensive network reinforcements in often aging physical infrastructures, DSOs are trying to increase the observability and controllability of their networks by leveraging new sensing and actuation technologies and exploring the ability to use data-driven algorithms to help with the integration of more DG in a more distributed (in space and time) and cost-effective way. This dissertation works towards this vision by formulating a systematic control-theoretic approach for integrating data-driven monitoring and control in the operation of electric distribution networks. Firstly, a Bayesian approach to state estimation overcomes the constraint of limited available real-time sensors by integrating voltage forecasting. A second class of tools discussed is the use of machine learning to decentralize Optimal Power Flow (OPF) methods, by utilizing inverter-interfaced Distributed Energy Resources (DERs). The Decentralized OPF method lets each

DER learn a policy that contributes to network objectives from its local historical data and measurements alone. This approach is formulated as a compression and reconstruction problem through an information-theoretic lens, providing fundamental limits of reconstruction and a strategy for optimal communication to improve learning-based reconstruction of optimal policies throughout a network. Lastly, the ambition to control networks in a distributed fashion triggers concerns about privacy-sensitive information that may be inferred from an agent's shared data. For a general class of algorithms, a new notion of local differential privacy is integrated that allows each agent to customize the protection of local information captured in constraints and objective functions. The ultimate goal of the work presented in this dissertation is to contribute to a framework for the integral and value-sensitive design and implementation of data-driven methodologies in critical infrastructure. To address the inherent cross-disciplinary nature of this larger goal, the final chapter explains how each automated decision-making tool reflects and affects values important to its stakeholders. The chapter argues that in order to enable beneficial integration of such tools, practitioners need to reflect on their epistemology and situate the design of automated decision-making in its inherently dynamic and human context.

To Lorenzo, his siblings and cousins

Contents

Contents	ii
List of Figures	v
List of Tables	viii
1 Introduction	1
1.1 A Revolution in Automation	1
1.2 Emerging Values and Concerns	2
1.3 Central Question and Organization	4
2 Electricity Distribution	6
2.1 Traditional Power Distribution	6
2.2 Trends and Challenges	7
2.3 Dissertation Scope and Research Questions	9
3 Modeling and Optimal Power Flow	12
3.1 A Control-Theoretic Perspective	12
3.2 First Principles	17
3.3 Single-Phase Power Flow	19
3.4 Optimal Power Flow	24
3.5 A Simpler Optimal Power Flow Problem	28
3.6 Three-Phase Power Flow	29
3.7 Optimal Power Flow in Unbalanced Networks	35
4 Forecasting-Based State Estimation in Single- and Three-Phase Distribution Systems	38
4.1 Introduction	39
4.2 Methodology	42
4.3 Forecasting	44
4.4 Real-time Estimation	46
4.5 Results	49
4.6 Suggestions for Observability Analysis and Sensor Placement	54

4.7	Conclusions	55
5	Decentralized Optimal Power Flow	57
5.1	Introduction	58
5.2	A Control-Theoretic Lens for OPF	61
5.3	Learning Decentralized Feedforward Controllers	63
5.4	Reconstruction and Communication	65
5.5	Results for Single-Phase Optimal Power Flow	66
5.6	Results for Three-Phase Optimal Power Flow	71
5.7	Discussion	71
5.8	Conclusions	75
6	Fully Decentralized Policies for Learning in Multi-Agent Systems: An Information-Theoretic Approach	76
6.1	Introduction	76
6.2	Related Work	78
6.3	General Problem Formulation	79
6.4	Allowing Restricted Communication	83
6.5	Application to Optimal Power Flow	84
6.6	Conclusions and Future Work	87
7	Customized Differential Privacy for Distributed Optimization in Multi-Agent Systems	88
7.1	Introduction	89
7.2	Preliminaries and Problem Statement	92
7.3	Main Results	94
7.4	Application: Simplified Distributed Optimal Power Flow	103
7.5	Sharing or Pricing the Privacy Budget?	107
7.6	Conclusions and Future Work	109
8	The Epistemology and Dynamics of Automated Decision-Making	110
8.1	Introduction	110
8.2	A Broader View On Bias	111
8.3	Technical Bias Is About Epistemology	112
8.4	Emerging Bias Is About Dynamics	114
8.5	Our Positionality Shapes Our Epistemology	116
9	Conclusions	118
9.1	Future Research Directions	120
	Bibliography	123
	A Power System Stability	139

A.1	Classic Notions of Efficacy in Power Systems	139
A.2	Power System Stability in Distribution Networks	140
A.3	Impact of Trends on Power System Stability	142
B	Proofs for Theorems and Lemmas Local Differential Privacy	147

List of Figures

3.1	Open-loop control.	13
3.2	Feedback control.	14
3.3	Feedback control.	14
3.4	Feedforward control.	15
3.5	Combined feedback and feedforward control.	15
3.6	Combined feedback and feedforward control.	16
3.7	From left to right: Depiction of Kirchhoff's Current Law, Kirchhoff's Voltage Law and Ohm's Law.	18
3.8	Left - An example radial feeder circuit based on the IEEE 37 Bus Test Feeder [107]. Right - Depiction of branch flow and bus injection over across two buses in a radial electric network.	19
3.9	Four simple flow scenarios and their relation to voltage magnitude and angle differences	23
3.10	Left - Nine flow scenarios and their relation to voltage magnitude and angle differences, the colors indicate scenarios with similar magnitude differences. Right - The solid lines indicate scenarios for which either magnitude or angle differences are zero, the arrows show how the other variable changes along the line, the dashed lines indicate how the lines rotate as the R/X -ratio decreases.	24
4.1	Overview of the forecasting and state estimation methodology.	42
4.2	Sources of uncertainty affecting the accuracy of the state estimator	43
4.3	Forecast of an aggregate load using a Gaussian Process model with only discrete-valued time features. Only 10% of the loads in the aggregate were recorded in historical data. The other 90% of loads were imputed with the average load profile. Poor forecast performance, such as on March 21st, motivates the use of Bayesian estimation.	46
4.4	IEEE 37 node test feeder model, voltage sensors are indicated with red circles.	51
4.5	Example voltage profile with forecast and estimation update at all the buses, numbered as in Figure 4.4.	51

4.6	ARMSE in p.u. for each non-measured bus across all phases. Buses with higher forecast errors benefit significantly from the estimation procedure. Buses that already have good accuracy on forecasted values of the order < 0.01 p.u. do not necessarily gain much from estimation. This can be attributed to the fact that these errors are in the same order as the modeling errors due to the estimator's linear approximation.	52
4.7	GIS view of Alliander's low voltage network of Rijsenhout. The outlined modeled network is the network that is considered for the DGSE model. The modeled part of the network consists of 34 customers. The unmodeled cables are not physically connected to the modeled network. At the distribution transformer and the community battery both power and voltage are measured. At 12 households, the power was measured. For privacy reasons, their exact location could not be displayed, but they are almost uniformly distributed along the cable.	53
4.8	Comparison of predicted and estimated voltage with real voltage measurement at the Neighborhood Battery bus.	54
5.1	Dynamical system representation with feedback and feedforward control.	62
5.2	Distribution feeder for scenario 1. Substation is located on the far left, locations of loads with PV inverters (squares), and without PV inverters (circles) are included.	67
5.3	Sample real power consumption and PV data at the node of Inverter B on July 4th 2014 from Pecan Street data. The profile is obtained by aggregating data from six individual residences.	68
5.4	All network voltages for approaches a , b , c , and d . Colored planes represent the range between the maximum and minimum voltages at any node in the network. Lower voltage bound is indicated with red dashed line, an additional indicator is included as a black dashed line at 0.98 p.u..	69
5.5	Upper figure - Objective function values of a , b . Lower figure - Difference in objective value for test date between Decentralized OPF control c and centralized OPF. The maximum difference is two orders of magnitude smaller than the optimal objective function, amounting to a suboptimality gap of 1.6%.	70
5.6	Total net load delivered by the substation, in the case of: no control, OPF and Decentralized OPF.	71
5.7	Voltage magnitude balancing. In grey, voltage magnitudes across three phases without control. In black, voltage magnitudes across three phases applying decentralized OPF.	72
5.8	Comparison of OPF communication strategies [54].	73

6.1	(a) Distributed multi-agent problem. The circles denote the local state x_i of an agent (a subset of (5.1)), the dashed arrow denotes its inputs u_i (red) and disturbances d_i (blue), and the double arrows denote physical coupling between the state variables of different agents, (b) Markov Random Field (MRF) graphical model of the dependency structure of all variables in the decentralized learning problem. Note that the state variables x_i and the optimal actions u_i^* form a fully connected undirected network, and the local policy \hat{u}_i only depends on the local state x_i	80
6.2	A flow diagram explaining the key steps of the decentralized regression method, depicted for the example system in Figure 6.1a. We first collect data from a multi-agent system, and then solve the centralized optimization problem using all the data. The data is then split into smaller training and test sets for all agents to develop individual decentralized policies $\hat{\pi}_i(x_i)$ (or $\hat{\pi}_i(d_i)$ as proposed in the context of OPF) that approximate the optimal solution of the centralized problem. These policies are then implemented in the multi-agent system to collectively achieve a common global behavior.	81
6.3	Results for optimal communication strategies on a synthetic Gaussian example. (a) shows squared correlation coefficients between of u^* and all x_i 's. The boxed entries correspond to x_9 , which was found to be optimal for $k = 1$. (b) shows that the optimal communication strategy of Thm. 6.4.1 achieves the lowest average distortion and outperforms the average over random strategies.	85
6.4	Results for decentralized learning on an OPF problem. (a) shows an example result of decentralized learning - the shaded region represents the range of all voltages in a network over a full day. As compared to no control, the fully decentralized regression-based control reduces voltage variation and prevents constraint violation (dashed line). (b) shows that the optimal communication strategy \mathcal{S}_i outperforms the average for random strategies on the mean squared error distortion metric. The regressors used are stepwise linear policies $\hat{\pi}_i$ with linear or quadratic features.	87
7.1	A system with 4 subsystems. Arrows denote message directions in distributed optimization. Here, node 0 denotes a central node that communicates with all other nodes. Note that Problems 7.2.1 and 7.2.2 can be fully distributed, and hence a central node is not necessary. The other 3 nodes represent subsystems. Nodes that have a privacy specification are patterned with dots and send out messaged perturbed by noise, as indicated by dashed arrows.	95
7.2	Achievable tradeoffs between privacy level ϵ and suboptimality \mathcal{S} , with Pareto front. (left) indicates increasing number of iterations and (right) increasing noise variance. We assume $\sigma_i = \sigma, G_i = G, \forall i \in \mathcal{C}$	106
7.3	Feasible parameter sets (ν, K) for varying levels of $\bar{\epsilon}$ (left) and $\bar{\mathcal{S}}$ (right), assuming $\sigma_i = \sigma, G_i = G, \forall i \in \mathcal{C}$	106
8.1	A Simple Feedback Model	115

List of Tables

3.1	Partitioning of nodal variables for all bus types.	17
5.1	Partitioning of nodal variables for all bus types.	62
5.2	Normalized regression coefficients for inverter A and B, as depicted in Figure 5.2.	73

Acknowledgments

Claire, my PhD journey started off in Fall 2013 under your wings, and picked up the collaboration we started when I was a visiting student in 2009-2010. Thank you for your continuing support, the freedom you have given me to explore my passions and combine different perspectives, and the patience and care with which you have helped me understand and cope with the challenges of academic research. Your human touch makes many students feel at home, and I feel grateful to have you as my role model. Duncan, when we first met I quickly understood I would benefit tremendously from your rigor and the multi-disciplinary perspective that you so well combine. Thank you for letting me be part of the EMAC group, and supporting me on my academic path. Sascha, thank you for challenging me to contextualize my research in the practical setting of distribution operation, and for showing me what it means to be a great teacher and communicator in a multi-disciplinary setting. Thank you also for being a great friend, I will cherish your honest advice and the many epic jam parties in Oakland. Dan, it has been an honor to work with you on so many different projects. I have benefited tremendously from following in your footsteps as a control engineer entering the field of energy and power. Thanks for your friendship, your comforting advice and for including me in the wonderful community you have built at LBL. Khalid, my interest in the ethical dimensions of engineering and technology quickly directed me to you. Your abilities to think both critically and constructively about our field and its role in society, and the transformative experiences you provide for students on this campus are unprecedented. Thank you for your inspiration and the many great conversations. Henrik, thanks for taking me on the team at C3 Iot and for being a source of inspiration. It was a treat to learn how technology is affecting real organizations, and to be part of such a smart and ambitious team. Ruzena, thank you for encouraging me from early on to explore and question the societal dimensions of our work. It is a great honor to interact with you. Your honesty and love for people are felt by so many, and your legacy is profound and forms the inspiration for us to persevere despite great challenges. I feel blessed to call you my titular grandmother on this side of the Atlantic. A word of appreciation for all the great staff members that make the PhD journey more enjoyable and less stressful and who are vital in making the many student initiatives in our department a success. A special thanks to Jessica, Shirley, Patrick, Sheila, and Audrey. A big thanks to my Control Freaks peers - Jason, Cathy, Eric and Jaime. It was a blast going through the PhD rollercoaster together, and I am thrilled and proud to see the amazing work you are doing and the lives you are building for yourselves. To all my peers and postdocs in the Hybrid Systems Lab - Jeremy, Insoon, Young-Hwan, Max, Mo, Qie, Ye, Ye, Melanie, Shahab, Frauke, Dariush, Kene, Somil, Palak, Margaret, Casey, Datong, Sylvia, Vicenç, Forrest, David, Paty and Donggun; thank you for the great ride. I have seen the group grow closer over the years, and I am grateful for the many fun times and productive collaborations we have shared, which will extend into the future. Thanks to all the folks in the semi-autonomous, EMAC and CAL-RAE groups and extended EECS community for being a second circle to call home on campus, in particular Daniel, Aaron, Anna, Jonny, Jalel, Ioannis and Varun. I have been blessed with the opportunity to learn how to advise talented

people - Oscar, Stephan, Aparna and Vincent; thanks for putting your trust in me and making vital contributions to this work. It is a treat to see you all grow as researchers and professionals. A big thank you to my collaborator Werner, who has enabled me to implement some of my research ideas on real systems in The Netherlands. I look forward to continuing and extending our exciting collaboration. Thanks to my collaborators Young-Hwan, Michael, Margaret, David, Paty, Ye, Ye, and Jingge for all being instrumental in establishing pieces of my PhD work. Without your great minds, I would not have been able to integrate such different perspectives that I believe are vital to solving the complex data-driven problems we are tasked with. Beyond the research priorities, my time at Berkeley has also been enriched through extracurricular and purely social interactions. It has been an honor to represent EECS Peers and help build a more supportive departmental culture for graduate students. Thanks to Kristin and Gireeja for initiating EECS Peers, to Aaron for passing on the baton, and to Aditya for a wonderful collaboration. A shoutout to Regina and Vasuki for all the energy they poured into making BiasBusters an initiative that has touched many in our department to think about implicit bias and cultural change, and that continues to inspire other students to engage in making our department a more inclusive and healthy place to work. In my last year, I have been silly enough to initiate a new student effort. Without the brilliant ideas and support of my GEESE friends Tom, Nitin and Sarah, this idea would have remained on the drawing table. I am excited to see where you will take the GEESE community to contribute to more opportunities on campus to engage across disciplines on issues of society and technology, and will be rooting for you all from New York. Getting a more formal training in music was a long-held dream, and I have been extremely fortunate to get a chance to sing in the UC Men's Chorale and Cal Jazz Choir, both directed by Bill Ganz. Bill, you are a blessing to all of us getting a chance to sing harmonies under your patient, critical and joyful guidance. Thank you for providing a space where we can explore our musical selves, your energy has contributed to a group I call family at Berkeley. Julie, Kubey, Jisoo, Kirstie, Sam, Jennifer, Valerie, Masami, Gina, Sheila, Jon, Jeremy, Jonathan, Min Ho, Finn, Jacob, Anthony and Scott, thanks for the many glorious harmonies, crazy creative moments and ample joy you have given me. In addition to all my peers and colleagues, my life in Berkeley was further shaped by beautiful people that have been there for me as friends and supporters. Darko, Ahmed, Marten, Rusi & Luka, Machiel, Esther & Sebas, Jethro, Miguel & Christina, Jack & Marissa, Olivia, Jonny & Taryn, thank you for being there for me. I was further blessed to share most of this ride with my sister Remy and brother Brit; guys I love you, you have made my time in Berkeley truly special. Thanks to my Dutch friends rooting for me from across the Atlantic, your messages and checking in with me makes all the difference. My parents and siblings have been there with me all the way through, with unconditional love and support and a always a listening ear. Mams, Paps, Luuk, Willemijn, Lieke, Robin en Lorenzo; I am deeply grateful for having you so close with me regardless of the distance. Leaves me to thank my beautiful wife Su Yin; thank you for your patience, for providing me with a mirror, and for being a source of inspiration to learn how to deal with the uncertainty inherent in a life in academia and abroad.

Chapter 1

Introduction

“Our fear of technology is really a fear of empowerment. We now have the ability to design the reality we live in, and we have to step up to the occasion.”

- Douglas Rushkoff

The dissertation starts with a discussion of the broader rise of data-driven technologies, covering its promises in Section 1.1, and its accompanied societal concerns in Section 1.2. Section 1.3 then introduces the central motivating question for this dissertation and a high-level organization of the following chapters. For a more thorough overview of the chapters and research questions covered in this dissertation, forward to Section 2.3.

1.1 A Revolution in Automation

Ubiquitous computing, algorithms, better devices, connectivity, and the ability to collect, store and probe large amounts of data are all becoming new commodities; commodities that are the key ingredients to enabling new forms of automation and innovation. In recent years, the general public is becoming accustomed to various data-driven functionalities. Most notably in online experience, where companies providing search, social network platforms, news, retail, education, entertainment and other personal services are gathering data about us to predict our preferences and determine how to best serve our needs, triggering a behavioral revolution [132]. We are rapidly adopting automation to assist us in navigation and transportation, to minimize the time, cost and hurdles of commuting and to facilitate a revolution towards vehicle autonomy, sharing and electrification [194]. The access to information about health and wellness combined with the ability to track our daily habits in terms of work, sleep, workout, social life and diet is revolutionizing healthcare to be more predictive, preventive, personalized and participatory [79]. And in our energy use, we are

able to control the use of our appliances and central heating from our browser or mobile app, to gather insights into our daily use and where to save energy, to let our heating system learn by itself how to operate and minimize energy use, to track the health of our appliances and when to maintain or replace them, and to incentivize ourselves towards more energy efficient behavior [168]. We adopt solar installations, shift to driving electric vehicles, and install battery storage to further reduce our dependency on the grid or even help the grid balance through providing ancillary services [170].

In professional domains, the commoditization of ingredients for automation has led to companies and institutions seeing tremendous value in data-driven technology, leading to ubiquitous experimentation and commercialization. In the context of operating critical infrastructure, such as traffic, electricity or cyber networks, challenges to make efficiency gains, to improve the efficacy of the infrastructure, or to integrate new technologies may be helped by leveraging the access to historical and real-time data, new sensing and actuation functions, and the increasing power of computation and algorithms. In traffic, the ability to monitor and communicate with large fleets of vehicles can be used to prevent congestion and alter traffic (either in space or time) in times of high demand. At the micro-level, semi-autonomous and fully autonomous driving holds a promise to significantly reduce the number of accidents and fatalities in traffic. In the context of energy systems, the automation of switching operations by letting distributed energy resources adjust their power injection in real-time may significantly reduce the expenses of reconfiguration procedures [178]. And the use of people's energy resources, such as battery storage, electric vehicles or solar installations, are now used to help balance the grid in real-time [170]. Such efforts in turn facilitate higher penetration levels of renewable generation helping the energy transition towards an infrastructure without fossil fuel generation.

The rise of new data-driven technologies is challenging the existing engineering disciplines to provide sufficient design principles for guaranteeing the safety, reliability and beneficial outcomes of a system. As recently argued by Professor Michael Jordan, "While the building blocks have begun to emerge, the principles for putting these blocks together have not yet emerged, and so the blocks are currently being put together in ad-hoc ways" [113]. This dissertation aims to work towards such principles, hoping to contribute to more comprehensive methodologies to for analyze, design and integrate data-driven technologies in critical infrastructures and human contexts, in ways cognizant of important values and the various interests of directly and indirectly affected stakeholders.

1.2 Emerging Values and Concerns

Despite its many promises and positive impacts, the integration of data-driven methodologies also brings new risks and concerns [55], and "many of our early societal-scale inference-and-decision-making systems are already exposing serious conceptual flaws" [113]. It is becoming increasingly clear that the speed at which new systems are changed or replaced is challenging our abilities to anticipate and mitigate serious side-effects. Here we discuss concerns in the

context of operating and using critical infrastructure.

A first area of risks and concerns in need of consideration is those faced primarily by an infrastructure's operator(s). First and foremost, the technology needs to be integrated in a *legacy system*. As such, data-driven techniques that consider a rethinking of many aspects of the infrastructure may be too costly, or those that forget to make the right assumptions may not be practically implementable. Secondly, the introduction of new functionality brings new types of vulnerability and failure modes. One example is the increasing worry of cybersecurity risks which increase once more sensors and devices are deployed throughout a system, providing more opportunities for hacking and manipulation by adversaries. We will see more automated and increasingly sophisticated social engineering attacks [224]. Another example is the increasing sophistication of algorithms used for control, which may make the behavior of such machines either inscrutable or non-intuitive [181], which can make it harder or impossible to characterize and reason about behaviors that are undesirable. The recent fatal incidents with self-driving vehicles [217] underline the critical importance to understand and characterize these failure modes and to report and communicate about them as new technology is integrated in vital infrastructure. Lastly, the replacement of human control or expertise by algorithmic efforts may make systems more efficient but can also lead to a lack of accountability. If new methods are embraced too quickly and extensively, the expertise needed to understand the system, when failures occur or an intervention is needed, may become harder to come by. And in the event of harm or fatality, it becomes harder to determine who is responsible, especially if different companies and entities were involved in convoluted ways to design and build the system that caused the accident.

A second area of risks and concerns are those faced by the citizens making use of the infrastructure. The most widely discussed concerns fall in the category of privacy, here broadly defined as *the state or condition of being free from being observed or disturbed by other people*. However, no single definition of privacy exists, and the nature of privacy is rather fragmented in how it is interpreted in policies and from culture to culture. Nissenbaum suggests that in the age of information it is more useful to see privacy as *contextual integrity*, which demands that the gathering and dissemination of data be appropriate to the norms of specific context [153]. While data sets are often small in particular contexts and have proven to be useful to determine personalized services to individuals or beneficial to the operation of a broader system or network, there is a growing public awareness that sharing too much may be risky, as these may reveal more information once combined with alternative sources of side information. Interpreting data collected in infrastructure, such as GPS location data or energy usage data from advanced metering infrastructure may reveal a lot about the person, such as where someone is at any given point in time, when someone is not home or what kind of appliances someone is using, especially once combined with other information about a person's social or professional life. This risk motivated the formulation of *differential privacy*, which aims to protect private information in data sets regardless of side information, typically necessitating a randomized mechanism to handle the data [66]. Another concern shared by the public is that of cyber security. With more and more sensors and actuators built in in our homes and close to our bodies connected to the internet, it becomes increasingly more

possible to hack and cause harm. The so-called *Internet of Things* (IoT) revolution is posing many new security challenges among which many might impact the lives of individual people. A third category of concerns is that of fairness and accessibility. Once services are increasingly digitized and personalized, it becomes easier for companies and operators to segment and distinguish between different users, increasing the opportunities for discrimination, unfair treatment or lack of access for certain citizens. In addition, the collection of data from users may yield all kinds of biases that, given the purpose towards the data is used, may lead to unfair outcomes [230]. Lastly, the increasing collection and storage of data, and the trend to trade and combine data sets about individuals are enabling ever improving predications about sometimes very personal and sensitive information about individual human beings. With the recent revelations about Cambridge Analytica's use of such extensive user profiles to design extensive influence campaigns to steer voting behavior in various elections [29], the public is starting to grow wary of being on online platforms and sharing information, as the concerns start to shift from mere privacy to protection of human dignity and democracy. With online platforms being designed to maximize our time online, using persuasion and free features to breed addiction [130], and used to spread violence [203], public unrest is likely to grow and demand revisiting the design of these systems.

It remains to be seen what regulatory actions are taken to protect citizens and how this will affect the use of data in the design and operation of critical infrastructure. The European Union is launching its General Data Protection Regulation on May 25, 2018, which will enact a wide set of rules to guide further digitization and address different privacy concerns [171, 71]. Apart from regulation, in recent years, various new research communities have sprung to understand how values can be incorporated in the design of systems. Examples are the Fairness, Accountability and Transparency community [41, 73], Value-Sensitive Design and Responsible Innovation [82, 103], and various groups around Safety and Artificial Intelligence [4, 108, 38]. It is clear that these concerns should be center stage in using data-driven technologies to maintain, update and extend the functionality of our critical infrastructures. As such, the context for this thesis goes beyond the technical and aims to situate the integration of data-driven methodologies in its appropriate socio-technical context.

1.3 Central Question and Organization

With the many promises and potential risks of data-driven automation in mind, the central question of this dissertation reads:

*How can we integrate data-driven techniques to improve and extend the capabilities of **existing critical infrastructure**, while safeguarding important values such as safety, privacy and social justice?*

While it is the author's intention is to provide a comprehensive and integral answer

to this question, its full scope is beyond the focus of this PhD dissertation. Anticipating that general knowledge may be attained, in this dissertation, the modernization of electric distribution networks to deal with higher levels of renewable generation and electrification will form the area of motivating applications. The energy transition in itself provides ample research problems relevant to the above question. Where possible, techniques developed for the electric grid are generalized to serve problems in other domains.

The word “integrate” refers to various forms of integration that are key in the applied research approach:

1. Integrating new technology into (possibly aging) legacy infrastructure
2. Integrating data-driven solutions with model-based approaches
3. Integrating various concepts and methods: system modeling, control, machine learning, optimization, information theory, differential privacy
4. Integrating the perspectives, interests and values of important stakeholders, such as domain experts and practitioners, citizens, users and other beneficiaries and affected groups or individuals

Chapter 2 introduces the context of electrical distribution systems and relevant challenges in their operations arising from the energy transition. Section 2.3 outlines the scope and concrete research questions for this dissertation in the context of electrical power, covering Chapters 4-7. The final Chapter 8 then takes a bird’s eye view on the practice of automated decision-making addressing the need to think about bias and error from a broader perspective, to reflect on the field’s epistemology, and to take into account the inherent feedback dynamics that occur once decision-making systems are employed in the world. Chapter 9 concludes with directions and questions for future work.

Chapter 2

Electricity Distribution

*“A system that works in practice,
but not in theory.”*

- Alexandra von Meier

This chapter introduces the context of electricity distribution. Section 2.1 provides a brief short history and describes the role of the system operator. Section 2.2 describes the trends and challenges that distribution operation is faced with due to the rapid transition to renewable energy generation. This sets the context for Section 2.3 defining the scope and research questions addressed in this dissertation.

2.1 Traditional Power Distribution

The consumption of electric power has become second nature to most people living in developed societies, and is increasingly becoming a reality in many areas where energy access has historically lagged behind. When we charge our phones or turn on the laundry machine, and we think about what happens at that moment, most of us do not think much further than electrons running through our appliance to make it do what it is supposed to. However, the energy that these electrons are charged with has traveled a long way through a vast infrastructure, which can be called the largest physically connected engineered machine ever built by mankind.

We have to go back to the late 19th century to trace back the origins of our current power system. After a long *War of the Currents* between leading entrepreneurs, inventors and tinkerers working on two different paradigms for power transmission, namely direct current (DC) and alternating current (AC), George Westinghouse and his fellow businessmen pushing for AC drew the longer straw. AC systems were generally less costly due to a smaller need for copper wiring and the ability to effectively scale AC power systems. Early power systems were built to power small communities and neighborhoods. As the AC technology developed,

it became possible to centralize and scale up power generation and connect these generators to existing local power systems.

This is what led to what we now call the *traditional power system*, in which power is generated centrally by large generators, transmitted at high voltage to substations across society, where the voltage is stepped down and the power is distributed to end customers. A central market is set up to match demand and supply at various time scales, from planning energy contracts with suppliers over multi-year periods to week-ahead and day-ahead markets planning the actual production of energy based on shorter-term load forecasts, to a real-time market in which the deviation between forecasted and actual demand is compensated and the stability of the system is managed through various levels of coordination and real-time control of power generators.

The focus of this dissertation is on the distribution of electric power, which is the responsibility of Distribution System Operators (DSOs). In the context of traditional power distribution, the DSO is not concerned with matching supply and demand within a distribution system, but primarily responsible for maintaining the quality of the voltage signal at customer level and ensuring the continuity of supply to each individual customer. To operationalize this responsibility, the DSO can build and control medium and low voltage networks. In many jurisdictions, a DSO is merely able to control the configuration of the network, and is not allowed to trade in energy in order to manage power flow through its networks. In the European context, DSOs with such restrictions are referred to as Distribution Network Operators (DNOs). However, recent legal and political debates are determining whether DNOs should get the right to trade in energy in order to allow for modern forms of network operation in which Distributed Energy Resources (DERs) aid the regulation of power flow and voltage by adjusting their power consumption or generation.

Traditionally, distribution networks are designed and updated in a conservative fashion, through expensive oversizing of all network and control components based on peak load scenarios expected for long time horizons, typically 40+ years. This robust design, established for unidirectional power flow from substation to end customer, allowed for a simple real-time operation and control paradigm. As such, most networks needed little sensing and few control actions by load tap changers, voltage transformers or capacitor banks to operate at very high levels of reliability.

2.2 Trends and Challenges

Most of today's power systems operated in the developed world provide a near 100% reliability, high efficiency, and a very low blackout percentage. With the so-called Renewables Portfolio Standard (RPS), California has enacted rules that require investor-owned utilities (IOUs), publicly owned utilities, electric service providers, and community choice aggregators to increase procurement from eligible renewable energy resources to 33% of total procurement by 2020 and 50% by 2030 [183]. Incentivized by the high price of imported oil, Hawai'i has been California's reference point, at the forefront of high levels of penetration of renewables. Today,

33% of its electricity comes from rooftop solar, 60% of its power comes from renewables on peak days, and it hopes to reach 100% renewables by 2045 [77]. To reach this goal, Hawai'i has had to go through struggles to learn how to cope with the impact on its grid operations [195, 36]. More recently, Scotland has set records in European territory, sourcing 68% of its electricity used in 2017 from green sources, up from 54% in 2016 [49], aiming for 100% by 2020.

The transition towards a power system without fossil fuel [204] and the accompanied electrification of our economy [68, 213] bring about new challenges. The effects of these changes on the reliability, security and stability of the electric power supply are numerous in form. Moreover, these effects are intricately related, and it is often not clear upfront what the technical or economic impact of a grid update will be. As a result, the transmission side of the grid has caught interest in revisiting criteria for power system stability, because of its inherent safety critical operating conditions and economic relevance to many stakeholders [208]. Distribution networks traditionally allow for more variability and comparatively less strict regulation on deviation from prescribed operating conditions. Electricity lines and components are typically over-dimensioned to allow for unforeseen deviations in the absence of monitoring. Furthermore, there is inherently less economic interest in the infrastructure as the number of reliant customers per component decreases. Nevertheless, the evolution of the distribution grid with more distributed generation (roof-top photovoltaic (PV) and wind), more sophisticated load dynamics, progressively more sensitive equipment (computers, servers, etc), and the necessary interfacing power electronics (e.g. inverters, transformers) do lead to new dynamic phenomena in the distribution network that both challenge reliable operation and provide new means of preventing or mitigating stability issues, both for the transmission grid and the distribution network itself [182]. We discuss the necessary impacts on distribution operation here, a more detailed reviewed can be found in Appendix A.

The traditional operation of distribution networks is typified by slowly changing and predictable power flow and simple but robust legacy control schemes using switches, transformers and capacitor banks [215]. This paradigm is challenged by two recent trends. Firstly, the rapid electrification and adoption distributed generation (DG) from wind and photovoltaics (PV) yields high spatio-temporal dynamic variability, at the order of seconds - contrasting a traditionally much more steady steady power flow. The occurrence of bidirectional and intermittent power flow due DG can cause rapid voltage fluctuations that can lead to unintended outages. These are increasingly harder to predict, and can result in protection issues, such as desensitization and unintended islanding or tripping [119]. This can lead to accelerated structural damage and potentially cascading failures, and yields economic burden due to accelerated wear, which decreases the lifetime of legacy equipment [182]. Especially if DG is connected to more sensitive feeders without proper monitoring functionality to assess its effect, cost of integration can easily multiply by a factor 3 to 4 [205, 72]. Secondly, the electrification and automation of many of our systems, such as driving and heating, diversifies load behavior, which is harder to model and predict, and which can lead to new overload scenarios and cascading failures. The rapid adoption of electric vehicles (EVs) will further aggravate this situation [93], especially if charging is optimized for electricity prices [212]. In

addition to challenges in operation, the inability to assess the impact of DG and EVs on the physical network causes utilities to impose conservative caps on the allowable PV capacity and number of EVs, hindering the transition to renewable energy sources.

For transmission grid stability, demand response via controllable loads [33], distributed generation (DG) or distributed energy resources (DERs) is starting to play a significant role in providing frequency regulation (or ancillary services) [155]; a service traditionally provided by fast gas-fueled power plants. Front-running economies like Denmark, which pledge a full transition to renewables and a fully distributed energy supply, understand the dependence on distribution side flexibility for stabilizing frequency (and more often voltage) as the alternative to traditional centralized generating sources [21]. As a result, one also needs to deal with the inherent variability of such distribution networks, which poses new stability problems for the transmission network.

Concerning stability of the distribution grid itself, there are many open problems when considering future scenarios. As distributed renewable generation keeps increasing, it is not unthinkable to assume that eventually most of our generation will be distributed [52] (as a matter of fact, Denmark is not far from such a scenario [21]). Apart from trying to update our current infrastructure and control systems to facilitate high levels of DG, such scenarios also trigger the question whether it is possible for local networks to disconnect or *island* and operate by themselves without the support from the transmission grid. As we know from operating (islanded) *microgrids* for remote operations or electrification of rural and developing communities, there are many new challenges in reliably (and cost-effectively) operating a distribution network without the supply and stabilizing inertia of transmission network. In such scenarios, a microgrid needs to manage the notions of stability on its own, as discussed in Appendix A.

2.3 Dissertation Scope and Research Questions

While the context of the full transition to renewable sources is a source for many relevant research questions, both applied and more theoretical, the scope of this dissertation is how to facilitate the more immediate and pressing needs that DSOs are facing in modernizing the planning and operation of their distribution networks. The major challenges to deal with the increased complexity due to the intermittency and variability of DG as well as the diversification of loads due to electrification and digitization have proved to provide plenty of problems, with both relevance in the application as well as yielding interesting system theoretic insights. Following the preamble of this chapter, the questions asked and answered in this dissertation aim to understand how emerging commodity ingredients for automation, as discussed in Chapter 1, can be *utilized* to address pressing issues of renewable integration, electrification and their effect on power system stability.

We can formulate the overall challenges in terms of the lack of two core competencies needed for safe and beneficial automation. Firstly, there is a lack of *observability*. The DSO has little visibility of what happens in the network and cannot act timely and appropriately.

The lack of sensors employed also makes it hard to assess the impact of new technologies on the grid. The conservative approach over over-dimensioning remains a robust approach to ensure near-100% reliability. However, a lack of understanding of power flow dynamics and how these evolve over time forces DSOs to make even more conservative and costly investments. As such, an investment in observability is an investment in allowing more advanced control strategies as well as smarter investment strategies.

Secondly, traditional distribution networks required little *controllability*. The diversification of dynamics in power flow and voltage occur at shorter timescale and at finer spatial resolutions, that most legacy control devices are not built for. In addition, bidirectional power flow breaks down the monotonicity of voltage profiles. A natural direction to pursue is seeing how controllable inverter-interfaced devices that are appearing across networks can be used to provide controllability at higher spatial and temporal resolutions. With the use of more devices throughout the network comes the necessity to compute and communicate solutions to optimization and control problems in real-time. As it is unlikely that robust communication infrastructures will be installed in many distribution networks, an important consideration is to see what can be done with local measurement. As the grid itself, and the problems of coordination that a DSO is tasked with are all intimately connected, studying ways to decentralize or distribute control or optimization problems over networks is a critical step towards practical implementation.

Lastly, many distribution networks have access to growing data repositories describing system parameters and historical state information, for instance via advanced metering infrastructure (AMI) and supervisory control and data acquisition (SCADA) systems. With the advent of machine learning techniques being increasingly able to find and represent high-dimensional patterns in low-dimensional models, an important area of problems is to see how learning can aid in monitoring and control. Machine learning is already being implemented in various aspects of operations to forecast loads, weather and renewable generation. In this dissertation, we study how such functionality can help make monitoring and control problems more efficient, more accurate, more data-driven and hence easier implement by overcoming excessive tuning of more ad-hoc control strategies.

With the collection of historical data, including customers' private consumption and generation, DSOs are facing increasing scrutiny and regulation from authorities governing what data may be used in what circumstances, and how it should be stored and managed. Conversations with DSOs in an early stage of this dissertation research made it evident that a lack of access to such data may be an important constraint on the solution space for monitoring and control problem. Therefore, we add two aims. Firstly, we see what can be done with purely historical data, without access to customers personal data in real-time. The aim to think about distributed or fully decentralized control strategies also contributes this aim. Secondly, we actively study how privacy may be defined technically and integrated in the very design of distributed optimization problems.

Research questions

Given the scope defined above, and the considerations drawn in Chapter 1, this dissertation addresses the following research questions in the context of Electricity Distribution:

- **Overarching:** How can data, modeling and algorithms be best combined to improve the spatial and temporal resolution of monitoring and control of voltage and power flow in distribution networks?
- **Chapter 4:** How may data-driven learning methods aid the design of state estimation methods with access to limited number of real-time sensors?
- **Chapter 6:** To what extent can local measurement reconstruct the information of a control policy that is optimal to a central optimization problem?
- **Chapter 5:** How may data-driven learning methods aid the design of decentralized control methods for voltage and power flow regulation that incorporate safety constraints and minimize the need of private information in real-time?
- **Chapter 7:** What information useful in optimal power flow may be sensitive in the sense of privacy? And how may be incorporate privacy in the design of distributed optimization problems?

The final Chapter 8 then takes a bird's eye view on the practice of automated decision-making addressing the need to think about bias and error from a broader perspective, to reflect on the field's epistemology, and to take into account the inherent feedback dynamics that occur once decision-making systems are employed in the world. Chapter 9 concludes with directions and questions for future work.

Chapter 3

Modeling and Optimal Power Flow

*“As far as the laws of
mathematics refer to reality,
they are not certain; and as far
as they are certain, they do not
refer to reality.”*

- Albert Einstein

This chapter provides a comprehensive overview of the control, modeling and optimization principles needed to set up the rest of the dissertation. It serves as a resource with more detail, aiming to keep the following chapters concise. Section 3.1 formulates a control-theoretic perspective to describe voltage and power flow regulation in quasi-steady state power systems. Section 3.2 goes into first principles for modeling power flow. Sections 3.3 and 3.6 then discuss power flow modeling in respectively single-phase and three-phase settings. Section 3.4 provides a rigorous overview of the literature on optimal power flow and provides some general optimization formulations. In Section 3.5, we then cover a simplified version of optimal power flow that does not take into account voltage as a state. Lastly, Section 3.7 covers a few formulations of optimal power flow using the modeling formulations from Section 3.6.

3.1 A Control-Theoretic Perspective

Here, we adopt and extend the formulations initially proposed in [219] for treating optimization algorithms from a *control-theoretic* perspective, which were further interpreted for power systems in [101]. This lens will help us interpret the plethora of monitoring and control methods proposed in the literature for electric distribution systems, and it will naturally trigger relevant questions around the analysis, development and integration of various solutions.

The State-Space Representation

In control theory, the state-space representation of a *dynamical system* can be thought of as described by a set of variables including input, state and output variables [81], related by a set of differential and/or algebraic equations. The input variables are also called *exogenous* as these typically enter the system from outside boundary, and can be split up in *controllable* input variables u (in short, “inputs”) and *uncontrollable* input variables d (in short, “disturbances”). The state variables x (or “states”) are all the variables necessary to mathematically describe the system’s full dynamics. The states may contain certain inputs, and as such the remaining states that are not inputs but needed to describe the system dynamics are called *endogenous*. Endogenous variables have values that are determined by other exogenous variables in the system, and are sometimes called *dependent* variables in econometrics [94]. The output variables z (or “outputs”) are all the variables that are measured in some *real-time* capacity, potentially at various sampling-rates.

The purpose of control is to analyze the system’s dynamical behavior and design a *controller* to reshape the dynamical behavior according to a set of specifications. We distinguish between open-loop, feedforward and feedback or closed-loop control.




Figure 3.1: Open-loop control.

In open-loop control, the input u^{ol} to the system is determined *independently* of the states or output, which are the variables that are being controlled. Open-loop control does not use feedback to determine if the output z has achieved a desired goal or reference set point r_z (similarly for the state). Some do include feedforward control to be open-loop as it does not use feedback, but for distinction we treat these as separate categories in this thesis, which means open-loop control in our definition also doesn’t depend on disturbances d . An example in power systems is scheduling the settings of a capacitor bank or load tap changer based on the time-of-day (note that time is not a state or output variable). Another example is a human operator deciding to turn on an extra generator.

In feedback control, the output z is used in real-time to inform the input u of the system, typically by comparing z to some defined reference signal or set point r_z . This form in which the output, i.e. measured variables, are used is called *output feedback control*. An example is the use of voltage measurements across a distribution network to automatically adjust the feeder voltage through controlling the load tap changer. Another example in which output feedback is less direct but arguably present, is when a human operator decides to reset control equipment based on observations of measurements on the system. In this case, the human control actions tend to not be permanent (unless the operator is constantly monitoring, such

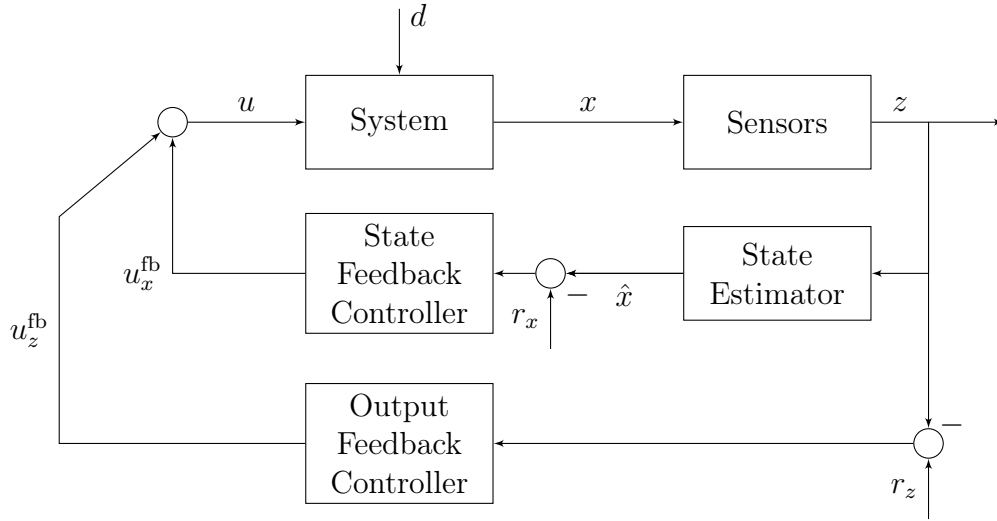


Figure 3.2: Feedback control.

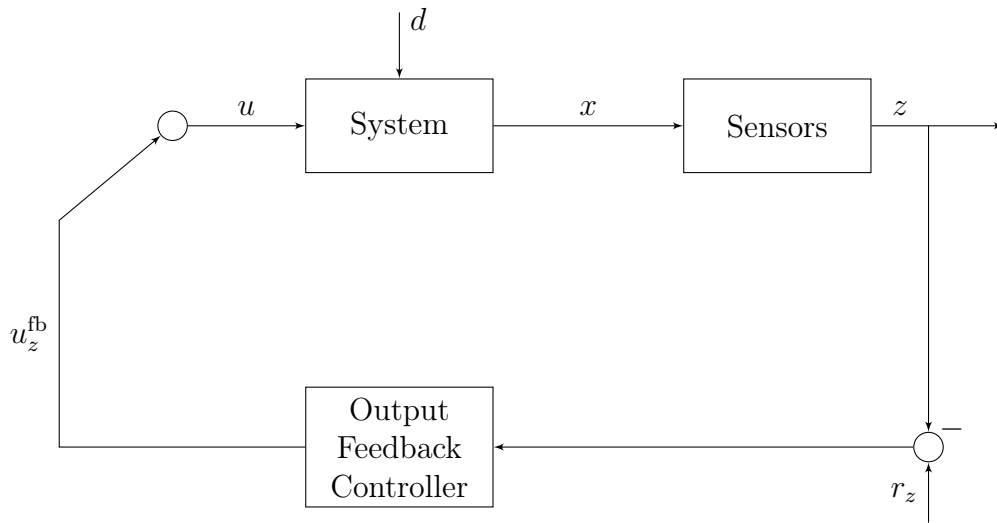


Figure 3.3: Feedback control.

as in air traffic control around airports), which makes it the overall behavior more open-loop. A second class of feedback control is when the the full state vector x is fed back to design the input u , usually called *state feedback control*. In this context, the state vector x may not be fully measured in the output z , and an *state estimator* or *observer* is needed to estimate the full state \hat{x} from the output.

Feedforward control is traditionally known as a simple and powerful technique to complement feedback control. It can be used both to improve the response to reference signals and to reduce the effect of measurable disturbances [10]. As such, there are broadly two forms of feedforward control. Firstly, a reference (either on output r_z or state r_x) can be




Figure 3.4: Feedforward control.

used to determine an input u_r^{ff} that *anticipates* the system’s dynamical response, typically by feeding r_z through an *inverse model* of the system. Secondly, information about the disturbances d , either through real-time measurement or a prediction, may be used to design an input u_d^{ff} that tries to compensate for its harmful effects by *attenuating* it directly or, again, *anticipating* the system’s dynamical response.

By combining feedforward and feedback control in the right way, one has two degrees of freedom to split the control design problem into two parts. The feedback controller can be designed to give good robustness and effective disturbance attenuation, and the feedforward part can be designed independently to give the desired response to command signals or objectives [10].

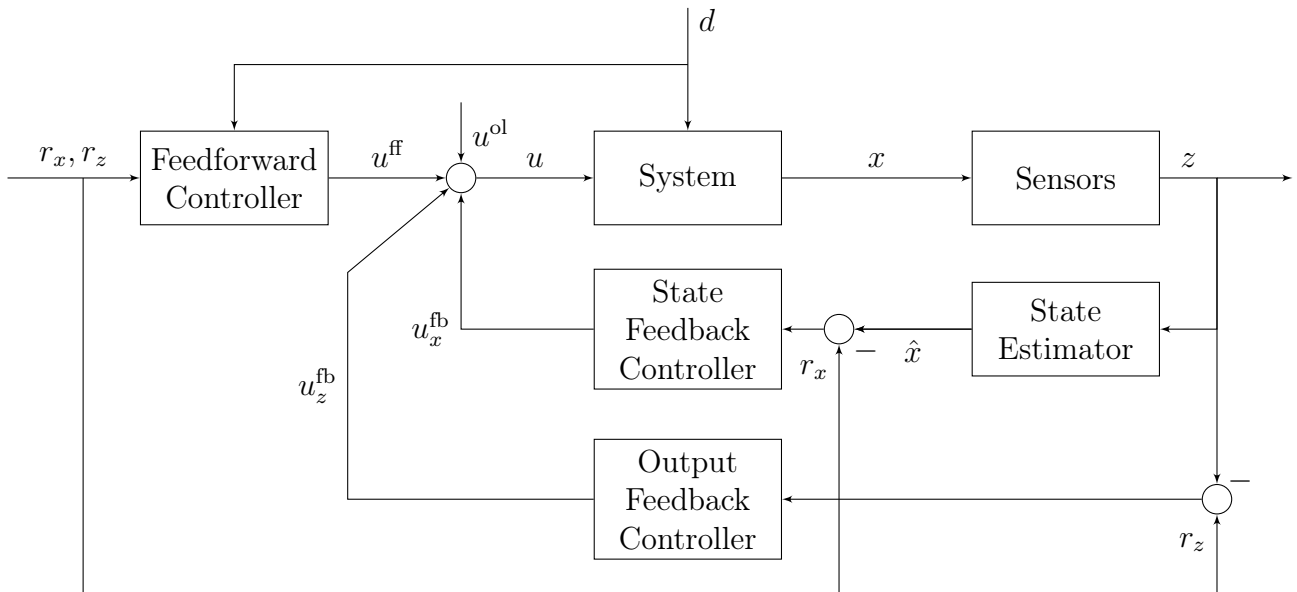


Figure 3.5: Combined feedback and feedforward control.

Note that traditionally, these notions of control are considered for systems with temporal dynamics described by first-order ordinary differential equations (ODEs). In electric distribution, temporal dynamics of voltage and power flow are present, but not captured in (ODEs). Rather, we consider the system to be in a *quasi steady-state* described by algebraic

equations following the first principles described in Chapter 3.2. Hence, we will adjust the control-theoretic framework in the next section.

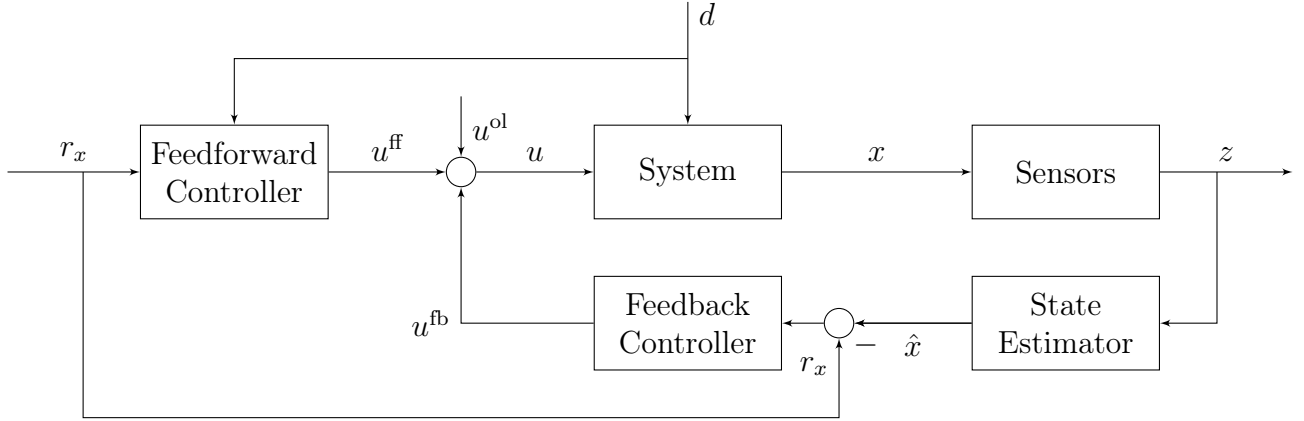


Figure 3.6: Combined feedback and feedforward control.

A State-Space Representation for Electric Distribution Systems

As argued above, the temporal dynamics in electric distribution systems with alternating current (AC) power are typically not described with ODEs, as the system's dynamics are assumed to be in a *steady-state*. This assumption rests on the condition that the distribution network is connected to the transmission grid through a substation, which can be modeled as an *infinite* bus, meaning that the transmission grid *admits* or *compensates* any changes in power across the distribution network to enable a *fixed* voltage phasor (hence the substation admittance is infinite or, equivalently, its impedance is zero).

With this assumption, voltage and power flow across a network can be determined through solving a set of algebraic equations describing fundamental first principles in electric power. Consider a network modeled by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with \mathcal{N} a set of nodes representing all buses with cardinality $\nu := |\mathcal{N}|$ and \mathcal{E} a set of edges representing all branches with cardinality $\eta := |\mathcal{E}|$. At each bus, we have 4 variables; voltage magnitude V_n , voltage angle δ_n , and nodal real and reactive power p_n and q_n , for all $n \in \mathcal{N}$. Vectorized for all buses, we have $V, \delta, p, q \in \mathbb{R}^\nu$ and we define the state of the system as

$$x := \begin{bmatrix} V \\ \delta \\ p \\ q \end{bmatrix} \in \mathbb{R}^{4\nu}. \quad (3.1)$$

As suggested in Section 3.1, we partition the state x into controllable inputs u , uncontrollable inputs or disturbances d and endogenous variables x^{end} . This partitioning is done per bus, based on the bus type, as suggested in [101] and summarized in Table 5.1.

	Exogenous		Endogenous
	Controllable - u	Uncontrollable - d	$x_{\setminus\{u,d\}}$
PQ generation	p_n, q_n		V_n, δ_n
PQ load		p_n, q_n	V_n, δ_n
PV generation	p_n, V_n		q_n, δ_n
slack bus	V_0	δ_0	p_0, q_0

Table 3.1: Partitioning of nodal variables for all bus types.

Note that the exogenous inputs are also treated as state variables: $\{u, d\} \subset x$. In principle, the classic state-representation allows for states to be inputs in an ODE although this is less common (e.g. $\dot{x} = u$, which will cause the state trajectory to be the initial condition plus the integral of the input over a time interval). In our setting, the state x is the solution to either a set of algebraic equations or an optimization problem. In the former case, we have to solve the power flow equations $F(x) = 0$ to find x ,

$$x \in \{\xi \in \mathbb{R}^{4\nu} : F(\xi) = 0\} =: \mathcal{M}. \quad (3.2)$$

We now derive the power flow equations $F(\cdot) : \mathbb{R}^{4\nu} \rightarrow \mathbb{R}^{2\nu}$ and discuss what is needed to find x . \mathcal{M} denotes the set of all power system states that satisfy $F(x) = 0$, and is named the *power flow manifold*, as defined in [25].

3.2 First Principles

We consider modeling the flow, consumption and generation of *alternating current* (AC) power in electric distribution networks with a *radial* topology, as well as the corresponding spatial voltage dynamics. We model these phenomena as steady state power flow, assuming the frequency of all wave forms to be constant, which allows us to describe power, voltage and current as *phasors*. We start by introducing key quantities and fundamental laws that govern electric power flow and voltage dynamics.

Let $\mathcal{T} = (\mathcal{N}, \mathcal{E})$ denote a graph representing a *radial* distribution feeder, as depicted in Figure 3.8 on the left, where \mathcal{N} is the set of buses or nodes of the feeder and \mathcal{E} is the set of line segments. Zooming in on one branch connecting two buses, as shown in Figure 3.8 on the right. Nodes are indexed by m and n , with $m, n \in \mathcal{N}$. Let $\mathcal{N} \triangleq \{\infty, 0, 1, \dots, |\mathcal{N}|\}$, where node 0 denotes the substation (feeder head). Immediately upstream of node 0 is an additional node used to represent the transmission system, indexed by ∞ . We treat node ∞ as an infinite bus, decoupling interactions in the downstream distribution system from the rest of the grid.

The complex voltage is $\mathbb{V}_n = V_n \angle \delta_n$, with V_n the voltage magnitude and δ_n the voltage angle at a node $n \in \mathcal{N}$. The voltage captures the energy per unit of charge and has the unit

volt or joule per coulomb. The complex current is defined as $\mathbb{I}_n = I_n \angle \phi_n$, with I_n the current magnitude and ϕ_n the current angle. The current measures the flow of electric charge across a surface (in our case the cross section of the conductive part of an electric wire or bus) and has the unit ampere or coulomb per second. The angle of both phasors are measured against a reference point or *slack bus* in the network. The complex power phasor for power flowing on a branch *out of* bus m towards bus $n \in \mathcal{N}$ is $\mathbb{S}_{mn} \triangleq \mathbb{V}_m(\mathbb{I}_{mn})^* = P_{mn} + jQ_{mn}$, measured in volt ampere, where P_{mn} is the real power flow corresponding to energy consumption and generation, as measured in watt, and Q_{mn} is the reactive power flow corresponding to the exchange of energy between inductive and capacitive elements in the system, as measured in volt ampere reactive [215]. For modeling purposes, we also introduce the power flowing *into* a node $n \in \mathcal{N}$ (rather than *out of* node m) as $\mathbb{S}_n \triangleq \mathbb{V}_n(\mathbb{I}_{mn})^* = P_n + jQ_n$ (i.e. we drop a subscript), with $\mathbb{S}_{mn} = \mathbb{S}_n + Z_{mn}I_{mn}^2$ being different due to the line loss. See Figure 3.8 for a pictorial explanation.




Figure 3.7: From left to right: Depiction of Kirchhoff's Current Law, Kirchhoff's Voltage Law and Ohm's Law.

Figure 3.7 introduces the first principles, both pictorially and mathematically. Kirchhoff's Current Law (KCL) describes the *conservation of charge* around a bus or node with multiple branches with current flowing in or out. Kirchhoff's Voltage Law (KVL) describes the *conservation of energy* around any circuit or any loop in a circuit. Ohm's Law theorizes that the current through a conductor between two points is directly proportional to the voltage across the two points. Here, $Z_{mn} = r_{mn} + jx_{mn}$ is the impedance value of the branch between nodes m and n , with r_{mn} the resistance and x_{mn} the reactance. The inverse of the impedance is the admittance $Y_{mn} = g_{mn} + jb_{mn}$.

For radial networks, we can simplify the notation slightly, as every branch can be uniquely assigned to a bus. Whenever we choose this notation, we can drop one of the indices on variables related to a branch: impedance, current and power flow.

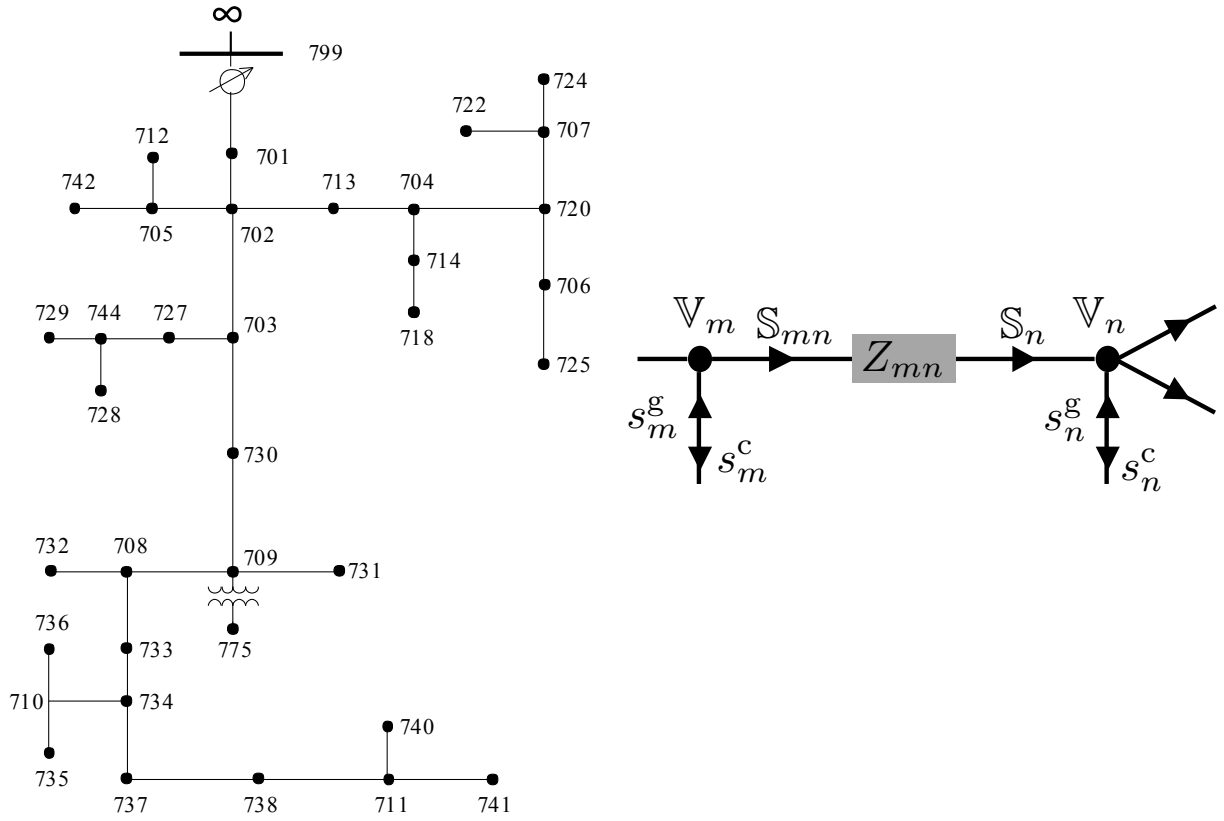


Figure 3.8: Left - An example radial feeder circuit based on the IEEE 37 Bus Test Feeder [107]. Right - Depiction of branch flow and bus injection over across two buses in a radial electric network.

3.3 Single-Phase Power Flow

Real and reactive power flowing on all branches in a *radial* network is dominated by the conservation of energy (KCL) and can be formulated as

$$P_{mn} = r_{mn}\ell_{mn} + p_n^c - p_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} P_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.3a)$$

$$Q_{mn} = x_{mn}\ell_{mn} + q_n^c - q_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} Q_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.3b)$$

where $\ell_{mn} := I_{mn}^2$ denotes the squared current magnitude and $r_{mn}\ell_{mn}$, $x_{mn}\ell_{mn}$ the real and reactive power loss on branch (m, n) .

The local voltage relation between two adjacent nodes in a radial network connected by an impedance, can be described as a function of the branch power flow by rewriting Ohm's

Law,

$$\mathbb{V}_m = \mathbb{V}_n + Z_{mn}\mathbb{I}_{mn}. \quad (3.4)$$

We now take different variants of this simple expression to derive linear approximations for voltage angle and magnitude differences in a network.

Voltage angle differences

Multiplying both sides of Equation (3.4) by the complex conjugate voltage \mathbb{V}_n^* yields, for the LHS and RHS respectively,

$$\begin{aligned} \mathbb{V}_n^* \mathbb{V}_m &= V_n (\cos \delta_n - j \sin \delta_n) V_m (\cos \delta_m + j \sin \delta_m) \\ &= V_n V_m \{ (\cos \delta_n \cos \delta_m + \sin \delta_n \sin \delta_m) + j (\cos \delta_n \sin \delta_m - \sin \delta_n \cos \delta_m) \} \\ &= V_n V_m (\cos(\delta_n - \delta_m) + j \sin(\delta_m - \delta_n)) , \end{aligned}$$

$$\begin{aligned} \mathbb{V}_n^* (\mathbb{V}_n + Z_{mn}\mathbb{I}_{mn}) &= V_n^2 + Z_{mn}\mathbb{S}_n^* \\ &= V_n^2 + (r_{mn} + jx_{mn}) (P_n - jQ_n) \\ &= V_n^2 + (r_{mn}P_n + x_{mn}Q_n) + j (x_{mn}P_n - r_{mn}Q_n) , \end{aligned}$$

where \mathbb{S}_n, P_n, Q_n denote power flowing *into* node n (instead of the convention with subscript $(\cdot)_{mn}$ which refers to the power flowing *out of* node m), see Figure 3.8. This yields the equation

$$V_n V_m (\cos(\delta_n - \delta_m) + j \sin(\delta_m - \delta_n)) = V_n^2 + (r_{mn}P_n + x_{mn}Q_n) + j (x_{mn}P_n - r_{mn}Q_n) ,$$

where by convention P_n, Q_n are the real and reactive power flowing into node n . We now consider the imaginary part of this equation separately, which yields the fundamental power equation

$$V_n V_m \sin(\delta_m - \delta_n) = x_{mn}P_n - r_{mn}Q_n . \quad (3.5)$$

We now consider small angle differences, i.e. $\sin(\delta_m - \delta_n) \approx \delta_m - \delta_n$. This approximation gives errors smaller than 1% for angle differences smaller than 0.244 radians (or 14 °) [190].

$$\delta_m - \delta_n \approx \frac{x_{mn}P_n - r_{mn}Q_n}{V_n V_m} . \quad (3.6)$$

In the context of high-voltage transmission networks, these equation are often further simplified. Since voltage levels are typically stable one can assume $V_n = 1$ pu. Furthermore, as $r \ll x$, the term $r_{mn}Q_n$ is dropped, and the voltage angle difference is a robust indicator for the real power flow between two adjacent nodes: $\delta_m - \delta_n \approx x_{mn}P_n$. This approximation explains the use of measuring synchronous voltage phase angles in transmission systems (through synchrophasor technology) to determine real power flow and balance. In distribution systems, however, these extra assumptions are not universally valid. Distribution circuits can contain line segments in which resistances and reactances are of the same order of magnitude. In addition, one cannot assume that voltage magnitudes are static and stable with respect

to changing power flows, which is an especially poor expectation towards the end of radial distributions feeders where power flow is small and rapidly changing in a relative sense leading to more volatile voltage fluctuation.

Voltage magnitude differences

We now consider Equation (3.4), and multiply both sides by their complex conjugates, which yields

$$\begin{aligned}
\mathbb{V}_m^* \mathbb{V}_m &= (\mathbb{V}_n + Z_{mn} \mathbb{I}_n)^* (\mathbb{V}_n + Z_{mn} \mathbb{I}_n) \\
V_m^2 &= V_n^2 + Z_{mn} \mathbb{V}_n^* \mathbb{I}_n + Z_{mn}^* \mathbb{V}_n \mathbb{I}_n^* + Z_{mn}^2 I_n^2 \\
&= V_n^2 + Z_{mn} \mathbb{S}_n^* + Z_{mn}^* \mathbb{S}_n + |Z_{mn}|^2 \frac{S_n^2}{V_n^2} \\
&= V_n^2 + (r_{mn} + jx_{mn})(P_n - jQ_n) + (r_{mn} - jx_{mn})(P_n + jQ_n) + (r_{mn}^2 + x_{mn}^2) \frac{P_n^2 + Q_n^2}{V_n^2} \\
&= V_n^2 + 2(r_{mn}P_n + x_{mn}Q_n) + (r_{mn}^2 + x_{mn}^2) \frac{P_n^2 + Q_n^2}{V_n^2}.
\end{aligned}$$

This yields a relation for the difference in squared voltage magnitude between two adjacent nodes in the network and the power flow on the intermediate branch,

$$V_m^2 - V_n^2 = 2(r_{mn}P_n + x_{mn}Q_n) + (r_{mn}^2 + x_{mn}^2) \frac{P_n^2 + Q_n^2}{V_n^2}. \quad (3.7)$$

We can assume that the term $(r_{mn}^2 + x_{mn}^2) \frac{P_n^2 + Q_n^2}{V_n^2}$ associated to losses is relatively small, motivating the approximation

$$V_m^2 - V_n^2 \approx 2(r_{mn}P_n + x_{mn}Q_n). \quad (3.8)$$

Note that one can further simplify this approximation by assuming $V_n \approx 1$ p.u., which yields $V_m^2 - V_n^2 \approx 2(V_m - V_n) \implies V_m - V_n \approx r_{mn}P_n + x_{mn}Q_n$.

A Linear Model For Single-Phase Power Flow

In a radial network setting, where we know the power consumption/generation, and assume that the losses are negligible, we can retrieve δ_n and V_n in a closed form solution. First we compute the power flow through all the lines, by starting at the final nodes of all branches and adding/subtracting the absorbed/injected power bottom-up towards the feeder (thus ignoring losses).

$$\begin{aligned}
P_n &\approx p_n + \sum_{k \in \mathcal{C}_n} P_k \approx \sum_{k \in \mathcal{D}_m} p_k, \\
Q_n &\approx q_n + \sum_{k \in \mathcal{C}_n} Q_k \approx \sum_{k \in \mathcal{D}_m} p_k,
\end{aligned} \quad (3.9)$$

where (P_n, Q_n) is the power flowing into node n , (p_n, q_n) is the power consumed/produced at node n , \mathcal{C}_n is the set of child nodes of node n , and \mathcal{D}_m is the set of nodes downstream of m . Together with Equation (3.8), these equations were termed the *Linear Distribution Flow* equations, in short *LinDistFlow*, which were introduced in [17] for optimal capacitor bank placement and have recently gained much attention due to work on convex methods for optimal power flow [136] and voltage regulation [74].

With (P_n, Q_n) for all lines in the network, we can now compute δ_n and V_n , starting at the feeder and working top-down towards the ends of all branches, using the derived equations:

$$\begin{aligned} \delta_m - \delta_n &\approx \frac{x_{mn}P_n - r_{mn}Q_n}{V_n V_m}, \\ V_m^2 - V_n^2 &\approx 2(r_{mn}P_n + x_{mn}Q_n). \end{aligned} \quad (3.10)$$

For ease of reference, we give these equations the name *CompDistFlow* equations, inspired by the fact that with the voltage angle, we now analyze the full voltage phasor \mathbb{V} in the complex plane.

Analysis of *CompDistFlow* Equations

Remember that we are in the load convention, and hence a positive power flow on branch $(m, n) \in \mathcal{E}$ into node n , i.e. $P_n, Q_n > 0$ means that there is a surplus of consumption in the area *downstream* of node m . A first aspect that jumps out when analyzing Equations (3.10) is that the magnitude difference is a positively weighted sum of the terms P_n, Q_n , whereas the angle difference is computed by taking the difference of weighted terms. As voltage magnitudes $V_n \geq 0$, the angle difference hence yields an interesting complement to the typically used magnitude difference. For ease of analysis, we start with setting $r_{mn} = x_{mn} = 1$ and inspect the following equations:

$$\begin{aligned} \Delta\delta_{mn} &\approx \frac{P_n - Q_n}{V_n V_m}, \\ \Delta V_{mn}^2 &\approx 2(P_n + Q_n). \end{aligned} \quad (3.11)$$

Imagine four different scenarios, in which either P or Q is very positive/negative, and the other part is relatively small. Inspecting Equations (3.11), it is straightforward to see that these relate to $\Delta\delta$ and ΔV as depicted in Figure 3.9. This leads to addressing the research question of how this approximation can help us *build intuition* about how real and reactive power flow direction and magnitude relate to changes in the voltage phasor. As the determination of real and reactive power flow is critical for voltage and power regulation, this can then help us build analysis and control techniques that are more sophisticated, then those used in practice (which mainly rely on using voltage magnitude and imposing transmission style assumptions on the distribution system).

We now evaluate a set of 9 different scenarios of power flow, i.e. $\{P > 0, P \approx 0, P < 0\} \times \{Q > 0, Q \approx 0, Q < 0\}$, and inspect the resulting magnitude and angle differences. The

	$\Delta V > 0$	$\Delta V < 0$
$\Delta\delta > 0$	$P > 0$	$Q < 0$
$\Delta\delta < 0$	$Q > 0$	$P < 0$

Figure 3.9: Four simple flow scenarios and their relation to voltage magnitude and angle differences

scenarios are depicted in Figure 3.10. The main take away from analyzing the left diagram is that for a specific value of ΔV there are multiple possible flow scenarios. For instance, observing a small voltage drop $\Delta V \approx 0$, may be because both the real and reactive power flowing between the nodes is low, i.e. $P, Q \approx 0$, or it may be that the voltage rise due to reverse real power flow $P < 0$ is canceled by reactive power consumption $Q > 0$. A similar ambiguity arises for $\Delta V > 0$. A certain value can be induced by either positive real or reactive power flow, or a combination thereof. *An important observation is that the voltage angle difference $\Delta\delta$ allows us to distinguish between the different cases for which ΔV is similar.*

The right diagram in Figure 3.10 provides some further intuition. First, it shows how ΔV and $\Delta\delta$ change in the P, Q -plane. This suggests a monotonic relation. Second, the dashed lines in the diagram indicate what happens to the relationships as the R/X -ratio decreases (remember that we did the initial analysis with $r = x = 1$). Note that as the $r/x \rightarrow 0$, i.e. for reactance dominant circuits, the lines end up as a cross with voltage magnitude explaining reactive power flow (horizontal) and voltage angle explaining real power flow. This is the equivalent of the decoupled power flow equations, ubiquitously used in transmission systems [80, Section 6]. In mathematical terms, we can write

$$\begin{bmatrix} \Delta\delta_{mn} \\ \Delta V_{mn}^2 \end{bmatrix} = \begin{bmatrix} x_{mn} & -r_{mn} \\ 2r_{mn} & 2x_{mn} \end{bmatrix} \begin{bmatrix} P_n \\ Q_n \end{bmatrix}, \quad (3.12)$$

where we simplified the angle difference equation by omitting the division by the voltage magnitude product. Observing the determinant of the 2×2 matrix, we see that

$$\det \begin{bmatrix} x_{mn} & -r_{mn} \\ 2r_{mn} & 2x_{mn} \end{bmatrix} = 2x_{mn}^2 + 2r_{mn}^2 = 2|Z_{mn}| > 0, \forall r_{mn}, x_{mn} \neq 0, \quad (3.13)$$

indicating that the Equation (3.12) is invertible for any impedance Z_{mn} . Hence, we can also write the inversion as

$$\begin{bmatrix} P_n \\ Q_n \end{bmatrix} = \frac{1}{2x_{mn}^2 + 2r_{mn}^2} \begin{bmatrix} 2x_{mn} & r_{mn} \\ -2r_{mn} & x_{mn} \end{bmatrix} \begin{bmatrix} \Delta\delta_{mn} \\ \Delta V_{mn}^2 \end{bmatrix}. \quad (3.14)$$

As such, the Equations (3.10) provide a *bijection* that can be used to infer the power flow on a branch based on knowledge of the voltage phasor difference, regardless of the impedance value.




Figure 3.10: Left - Nine flow scenarios and their relation to voltage magnitude and angle differences, the colors indicate scenarios with similar magnitude differences. Right - The solid lines indicate scenarios for which either magnitude or angle differences are zero, the arrows show how the other variable changes along the line, the dashed lines indicate how the lines rotate as the R/X -ratio decreases.

3.4 Optimal Power Flow

Optimal Power Flow (OPF) refers to solving an optimization problem that minimizes some economic or operational objective subject to power flow constraints (and other relevant constraints). Traditionally, OPF is developed for planning and coordinating overall transmission grids as an extension of economic dispatch [80] meant to ensure that the dispatch of generators minimizes economic cost and results into a power flow scenario that is physically feasible and satisfying relevant operating and safety constraints. As such, solving the OPF problem requires a model of the electric grid describing both topology and impedances, as described in Equations (3.9)-(3.10) for single-phase and Equations (3.60)-(3.63) for three-phase power.

OPF can be used *offline* as a design tool for network upgrades to size and place equipment, as proposed for capacitor planning in [17, 15], or as a planning tool to schedule the dispatch of generators and control equipment [80]. With the advent of real-time sensing, the implementation of OPF in an *online* setting has been a popular research area in recent years. Historically, the non-convex nature of the OPF problem has limited its online integration due to the lack of convergence guarantees [148]. Due to the nonlinear nature of power flow equations, many OPFs are formulated as quadratically constrained quadratic programs (QCQPs) (see [26] for general background). QCQPs are generally non-convex and NP-hard, unless certain strict conditions hold. Progress in optimization theory has since enabled solutions to the OPF problem through a *convex* formulation [12, 128] that exploits Semi Definite Programming (SDP) *relaxations* of QCQPs (see [136] for a holistic introduction in the context of OPF). SDP-based OPF problems can be solved for a broad range of practical conditions in single-phase networks under certain assumptions on the system's structure and operating conditions [192, 136, 137]. For single-phase networks, theoretical results have been derived for both radial [88] and meshed [140] networks, as well as for both the branch

flow model [75] and the bus injection model [128] ([136] discusses the equivalence of these results across the two models). Limitations of the SDP relaxations for single-phase networks have been reported in [129, 146, 147], with negative locational marginal prices [14] and nondegeneracy and inexactness of semidefinite relaxations [135] forming relevant obstacles for broad reliable integration.

Practically, implementing OPF in real-time settings for distribution networks is still hampered by various challenges. Firstly, the power flow dynamics in distribution networks are three-phase and unbalanced, requiring additional modeling efforts [122] and rendering most of the above-listed theory on convex relaxations insufficient or not relevant. A popular method for analyzing such OPFs is relaxation via semidefinite programming (SDP) [45, 46]. It is well documented that relaxation of OPF problems via SDP often fails to achieve a rank-one solution. As an example, in the work of [135] too many binding constraints will preclude convergence to a rank-one solution. The authors of [140] explored the extension of SDP to weakly meshed networks. Their technique was able to achieve a rank-one solution only after incorporating significant penalties on reactive power dispatch, effectively limiting the feasible region of control. Finally, in [46], the authors faced difficulty in obtaining a rank-one solution for certain network configurations. As the inability of relaxations via SDPs to achieve a rank one solution limits the practicality of these approaches, it is necessary to consider alternative approaches for solving OPF problems. One such alternative is the creation of linear approximations for power flow that are sufficiently accurate for control purposes, and that can be incorporated into convex OPF formulations, as proposed in [87, 178].

Secondly, implementing OPF relies on an extensive and robust communication infrastructure and a high-fidelity network model. As most distribution networks consist of far more nodes with a limited communication infrastructure, it is more challenging to use OPF approaches to govern control equipment and Distributed Energy Resources (DERs), as proposed in [48].

Power Flow Constraints and Convex Relaxation

The full nonlinear power flow model for single-phase networks reads

$$P_{mn} = r_{mn}\ell_{mn} + p_n^c - p_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} P_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.15a)$$

$$Q_{mn} = x_{mn}\ell_{mn} + q_n^c - q_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} Q_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.15b)$$

$$y_m - y_n = 2(r_{mn}P_{mn} + x_{mn}Q_{mn}) - (r_{mn}^2 + x_{mn}^2)\ell_{mn}, \quad \forall (m, n) \in \mathcal{E}, \quad (3.15c)$$

$$\ell_{mn} = \frac{P_{mn}^2 + Q_{mn}^2}{y_m}, \quad \forall (m, n) \in \mathcal{E}, \quad (3.16)$$

where $y_n \triangleq V_m^2$, and where we have assumed the nodal injections to be constant power and not sensitive to voltage or current fluctuations. Under fairly general conditions, it can be

shown that the nonlinear equality constraint (3.16) can be relaxed in an optimization setting to the second order cone inequality constraint [136]:

$$y_m \ell_{mn} \geq P_{mn}^2 + Q_{mn}^2 \quad (3.17)$$

This equation can be rewritten as:

$$\begin{aligned} 4y_m \ell_{mn} &\geq 4P_{mn}^2 + 4Q_{mn}^2 \\ 2y_m \ell_{mn} &\geq 4P_{mn}^2 + 4Q_{mn}^2 - 2y_m \ell_{mn} \\ 2y_m \ell_{mn} + \ell_{mn}^2 + y_m^2 &\geq 4P_{mn}^2 + 4Q_{mn}^2 - 2y_m \ell_{mn} + \ell_{mn}^2 + y_m^2 \end{aligned} \quad (3.18)$$

This yields the second order cone:

$$\left\| \begin{array}{c} 2P_{mn} \\ 2Q_{mn} \\ \ell_{mn} - y_m \end{array} \right\|_2 \leq \ell_{mn} + y_m \quad (3.19)$$

We also consider an approximated model in which the power losses are assumed to be small, that is $r_{mn}\ell_{mn}, x_{mn}\ell_{ij} \approx 0$. This yields:

$$P_{mn} = p_n^c - p_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} P_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.20a)$$

$$Q_{mn} = q_n^c - q_n^g + \sum_{(n,k) \in \mathcal{E}, k \neq m} Q_{nk}, \quad \forall n \in \mathcal{N}, \quad (3.20b)$$

$$y_m - y_n = 2(r_{mn}P_{mn} + x_{mn}Q_{ij}), \quad \forall (m, n) \in \mathcal{E}, \quad (3.20c)$$

Operational Objectives and Constraints

In the case studies of this thesis, we aim to optimize various objectives relevant in distribution grid operations. Firstly, we will minimize voltage deviation from a reference voltage or general voltage variability due to the intermittent nature of generation and consumption. This objective can help to keep the voltage in a safely operable regime. Secondly, we will address efficiency and systems losses by minimizing power flow over branches throughout the network. Through this objective, we implicitly try to minimize the power delivered from the transmission grid, injected at the top of a radial network, and maximize the use of distributed energy resources to provide loads nearby. Lastly, in the case that we rely on the convex relaxation in Equation (3.19), the objective requires a convex objective of the square current magnitude ℓ_{mn} , acting as a force to ensure the inequality constraint ends up on the feasibility boundary, yielding the satisfaction of the original equality constraint (3.16). These goals are formulated with the objective function, with α, β, γ denoting trade-off parameters (these could take on the form of a financial price), and y_{ref} the reference voltage throughout the network.

$$f_o := \sum_{n \in \mathcal{N}} \alpha (y_n - y_{\text{ref}})^2 + \beta \sum_{(m,n) \in \mathcal{E}} (P_{mn}^2 + Q_{mn}^2) + \gamma \sum_{(m,n) \in \mathcal{E}} r_{mn} \ell_{mn}. \quad (3.21)$$

Two constraints are included to account for inverter capacity and voltage goals. Assume we are given a subset $\mathcal{C} \subset \mathcal{N}$ of buses that are equipped with a controllable source or load. Each controller $i \in \mathcal{C}$ is ultimately limited by a local capacity on total apparent power capacity \bar{s}_i . We consider a different controller architectures that can deliver *both* real and reactive power. A first version considers that real and reactive power capacity are decoupled and assumed to have been assigned their own capacity

$$\underline{p}_i(t) \leq u_i^p(t) \leq \bar{p}_i(t), \quad \underline{q}_i(t) \leq u_i^q(t) \leq \bar{q}_i(t) \quad (3.22)$$

A second version considers a *four-quadrant* configuration, which yields the disk constraint in the complex plane

$$(u_i^p)^2(t) + (u_i^q)^2(t) \leq \bar{s}_i^2(t) \quad (3.23)$$

Lastly, we consider inverters that also interface a system with local photovoltaic (PV) installation, an electric vehicle (EV) charging system, smart loads or a combination thereof. Due to the nature of these subsystems having some hierarchy of priority, it can happen that the capacity available for the OPF problem is dependent on the scheduling and activity of these other systems. For instance, it is possible the inverter can only deliver/consume reactive power based on the remaining capacity *after* injecting surplus energy from the PV installation into the grid. In this case, the reactive power capacity $\bar{q}_i[t]$ at time t of an inverter is limited by the total apparent power capacity \bar{s} (constant) minus the real power $p_i^g[t]$. As a result, the demand of reactive power does not interfere with real power generation. Therefore, the reactive power capacity can be formulated as

$$|u_i^q[t]| \leq \bar{q}_i[t] = \sqrt{\bar{s}_i^2 - (p_i^g[t])^2}. \quad (3.24)$$

In this scenario, we will assume that each inverter has some overcapacity with respect to the maximum real power output, e.g. $\bar{s} = 1.05\bar{p}$. Another concrete example is when a household or business owning the inverter has a schedule for charging EVs or serving other loads, yielding a similar constraint formulation, possibly for both real and reactive power,

$$(u_i^p + p_i^{\text{prio}})^2(t) + (u_i^q + q_i^{\text{prio}})^2(t) \leq \bar{s}_i^2(t), \forall i \in \mathcal{C}. \quad (3.25)$$

The responsibility of American utilities to maintain service voltage within $\pm 5\%$ of 120V as specified by ANSI Standard C84.1 is expressed as a constraint in the optimization problem

$$\underline{y} \leq y_n \leq \bar{y}, \forall n \in \mathcal{N}. \quad (3.26)$$

Economic Objective

We consider a scenario in which the utility negotiates different prices for different capacities, potentially at different points in time, with different third party DER owners. The resulting business exchange can be formulated as

$$f_e := \sum_{i \in \mathcal{C}} \lambda_i ((u_i^p)^2 + (u_i^q)^2) + \mu_i (u_i^p + u_i^q). \quad (3.27)$$

Here, u_i refers to real or reactive power used for the optimal power flow scheme from agent i , and λ_i, μ_i denote the price levels for procuring a kVAr or kWatt from agent i .

Optimization Problem

The goal of optimal power flow is to determine the control setpoints that minimize the global operational and economic objectives subject to the operational constraints. The linear OPF problem reads

$$\begin{aligned} \min_{\mathbf{z}} \quad & f_o + f_e & (3.28) \\ \text{s.t.} \quad & (3.3), (3.22), (3.26) \\ & \mathbf{z} = (y_n, P_{mn}, Q_{mn}, u_i^p, u_i^q) \quad \forall n \in \mathcal{N}, \quad \forall (m, n) \in \mathcal{E}, \quad \forall i \in \mathcal{C}. \end{aligned}$$

3.5 A Simpler Optimal Power Flow Problem

Operational Constraints

In this case study, we solely aim to prevent overload of real power flow over certain critical branches in an electric network, and we do not model the voltage levels. We assume a radial network topology and that line losses are negligible, which allows us to model the branch flow on branch $(m, n) \in \mathcal{E}$ as the nodal power of all nodes downstream of m , denoted by \mathcal{D}_m , as formulated in Equation (3.9). This aim is formulated through constraints

$$\begin{aligned} \sum_{k \in \mathcal{D}_m} \{p_k^c - p_k^g + u_k^p\} - \overline{P_{mn}} & \leq 0, \\ \underline{P_{mn}} - \sum_{k \in \mathcal{D}_m} \{p_k^c - p_k^g + u_k^p\} & \leq 0, \quad \forall (m, n) \in \mathcal{E}_{\text{safe}}, \end{aligned} \quad (3.29)$$

with $\mathcal{E}_{\text{safe}} \subset \mathcal{E}$ a subset of branches for which power flow limitations are defined, $\overline{P_{mn}}, \underline{P_{mn}}$ denoting the upper and lower power flow bounds on branch $(m, n) \in \mathcal{E}_{\text{safe}}$. Constraints are included to account for inverter capacity. Each node i is ultimately limited by the local capacity on total apparent power capacity \bar{s}_i . We consider a simple controller architectures that can deliver real power, which we assume are both given their own capacity

$$\underline{u}_i \leq u_i^p \leq \bar{u}_i, \quad \forall i \in \mathcal{C}. \quad (3.30)$$

Economic Objective

We consider a scenario in which the utility negotiates different prices for different capacities, potentially at different points in time, with different third party DER owners. The resulting business exchange can be formulated as:

$$f_e := \sum_{i \in \mathcal{C}} \lambda_i (u_i^p)^2 \quad (3.31)$$

Here, u_i^p refers to real power used for the optimization scheme from agent i , and λ_i denotes the price for procuring a kWatt from agent i .

Optimization Problem

The goal of optimal power flow is to determine the control setpoints that minimize the global operational and economic objectives subject to the operational constraints. The linear OPF problem reads

$$\begin{aligned} \min_{\mathbf{z}} \quad & \sum_{i \in \mathcal{C}} \lambda_i (u_i^p)^2 & (3.32) \\ \text{s.t.} \quad & (3.29), (7.28), \\ & \mathbf{z} = (u_i^p) \quad \forall i \in \mathcal{C}. \end{aligned}$$

3.6 Three-Phase Power Flow

In this section, which is adopted from joint work [178] and included for sake of completeness, we derive a linear model for three-phase unbalanced power flow. Power flow linearizations have been studied more broadly recently in [24, 25]

Preliminaries

Let $\mathcal{T} = (\mathcal{N}, \mathcal{E})$ denote a graph representing an unbalanced distribution feeder, where \mathcal{N} is the set of nodes of the feeder and \mathcal{E} is the set of line segments. Nodes are indexed by m and n , with $m, n \in \mathcal{N}$. Let $\mathcal{N} \triangleq \{\infty, 0, 1, \dots, N\}$, where node 0 denotes the substation (feeder head). Immediately upstream of node 0 is an additional node used to represent the transmission system, indexed by ∞ . We treat node ∞ as an infinite bus, decoupling interactions in the downstream distribution system from the rest of the grid. While the substation voltage may evolve over time, we assume this evolution takes place independently of DER control actions in \mathcal{T} .

Each node and line segment in \mathcal{T} can have up to three phases, labeled a , b , and c . Phases are referred to by $\phi \in \{a, b, c\}$ and $\psi \in \{a, b, c\}$. We define \mathcal{P}_m and \mathcal{P}_n as the set of phases at nodes m and n , respectively, and \mathcal{P}_{mn} as set of phases of line segment (m, n) . If phase ϕ is present at node m , then at least one line connected to m must contain phase ϕ . If line (m, n) exists, its phases are a subset of the phases present at both node m and node n , such that $(m, n) \in \mathcal{E} \Rightarrow \mathcal{P}_{mn} \subseteq \mathcal{P}_m \cap \mathcal{P}_n$.

The current/voltage relationship for a three phase line (m, n) between adjacent nodes m

and n is captured by Kirchhoff's Voltage Law (KVL) in its full (3.33), and vector form (3.34):

$$\begin{bmatrix} \mathbb{V}_m^a \\ \mathbb{V}_m^b \\ \mathbb{V}_m^c \end{bmatrix} = \begin{bmatrix} \mathbb{V}_n^a \\ \mathbb{V}_n^b \\ \mathbb{V}_n^c \end{bmatrix} + \begin{bmatrix} Z_{mn}^{aa} & Z_{mn}^{ab} & Z_{mn}^{ac} \\ Z_{mn}^{ba} & Z_{mn}^{bb} & Z_{mn}^{bc} \\ Z_{mn}^{ca} & Z_{mn}^{cb} & Z_{mn}^{cc} \end{bmatrix} \begin{bmatrix} \mathbb{I}_{mn}^a \\ \mathbb{I}_{mn}^b \\ \mathbb{I}_{mn}^c \end{bmatrix}, \quad (3.33)$$

$$\mathbf{V}_m = \mathbf{V}_n + \mathbf{Z}_{mn} \mathbf{I}_{mn}. \quad (3.34)$$

Here, $Z_{mn}^{\phi\psi} = r_{mn}^{\phi\psi} + jx_{mn}^{\phi\psi}$ denotes the complex impedance of line (m, n) across phases ϕ and ψ . We have presented (3.33) and (3.34) where $\mathcal{P}_{mn} = \{a, b, c\}$. For lines with fewer than three phases ($|\mathcal{P}_{mn}| \leq 2$) (3.34) becomes:

$$[\mathbf{V}_m = \mathbf{V}_n + \mathbf{Z}_{mn} \mathbf{I}_{mn}]_{\mathcal{P}_{mn}}, \quad (3.35)$$

by indexing by the set of line phases \mathcal{P}_{mn} , where the rows associated with phases $\psi \notin \mathcal{P}_{mn}$ of (3.34) are removed, as are the appropriate columns of \mathbf{Z}_{mn} . To give two examples, if $\mathcal{P}_{mn} = \{a\}$, then (3.35) is $[\mathbf{V}_m]_{\{a\}} \equiv \mathbb{V}_m^a = \mathbb{V}_n^a + Z_{mn}^{aa} \mathbb{I}_{mn}^a$, and if $\mathcal{P}_{mn} = \{a, c\}$ then (3.35) is:

$$[\mathbf{V}_m]_{\{a,c\}} \equiv \begin{bmatrix} \mathbb{V}_m^a \\ \mathbb{V}_m^c \end{bmatrix} = \begin{bmatrix} \mathbb{V}_n^a \\ \mathbb{V}_n^c \end{bmatrix} + \begin{bmatrix} Z_{mn}^{aa} & Z_{mn}^{ac} \\ Z_{mn}^{ca} & Z_{mn}^{cc} \end{bmatrix} \begin{bmatrix} \mathbb{I}_{mn}^a \\ \mathbb{I}_{mn}^c \end{bmatrix}.$$

Kirchhoff's Current law at node m is given in its full (3.36) and vector (3.37) forms:

$$\sum_{l:(l,m) \in \mathcal{E}} \begin{bmatrix} \mathbb{I}_{lm}^a \\ \mathbb{I}_{lm}^b \\ \mathbb{I}_{lm}^c \end{bmatrix} = \begin{bmatrix} i_m^a \\ i_m^b \\ i_m^c \end{bmatrix} + \sum_{n:(m,n) \in \mathcal{E}} \begin{bmatrix} \mathbb{I}_{mn}^a \\ \mathbb{I}_{mn}^b \\ \mathbb{I}_{mn}^c \end{bmatrix}, \quad (3.36)$$

$$\sum_{l:(l,m) \in \mathcal{E}} \mathbf{I}_{lm} = \mathbf{i}_m + \sum_{n:(m,n) \in \mathcal{E}} \mathbf{I}_{mn}. \quad (3.37)$$

We assume a complex load, $s_n^\phi \forall \phi \in \mathcal{P}_n, \forall n \in \mathcal{N} \setminus \infty$, is served on all existing phases at each node except the transmission line, defined as:

$$s_n^\phi (V_n^\phi) = \left(\beta_{S,n}^\phi + \beta_{Z,n}^\phi (V_n^\phi)^2 \right) d_n^\phi + u_n^\phi - j c_n^\phi, \quad (3.38)$$

where d_n^ϕ is the complex demand, with constant power and impedance terms $\beta_{S,n}^\phi + \beta_{Z,n}^\phi = 1$, with $u_n^\phi = u_n^{p,\phi} + j u_n^{q,\phi}$ represents complex power available for control (e.g. DER), and c_n^ϕ denotes capacitance, all for $\phi \in \mathcal{P}_n$. Note, if $\phi \notin \mathcal{P}_n$ (i.e. phase ϕ does not exist at node n), we define $V_n^\phi = i_n^\phi = s_n^\phi = d_n^\phi = u_n^\phi = c_n^\phi = 0$. If $\phi \notin \mathcal{P}_{mn}$ (i.e. phase ϕ does not exist on line segment (m, n)), we define $\mathbb{I}_{mn}^\phi = 0$.

Throughout this work, we use the symbol \circ to represent the Hadamard Product (HP) of two matrices of the same dimension, also known as the element-wise product, which can be written as:

$$C = A \circ B = B \circ A \Rightarrow C_{ij} = A_{ij} B_{ij} = B_{ij} A_{ij}.$$

Power and Losses

We now derive complex power and loss terms at a node $m \in \mathcal{N}$. This analysis, and the derivation of 3.6, we do not claim as novel contributions (see [87, 173]). Full derivation of these results are necessary to support one of the main contributions of this work, which is presented in 3.6. To start, we take the Hadamard Product of \mathbf{V}_m and the complex conjugate (non-transposed) of (3.37):

$$\sum_{l:(l,m) \in \mathcal{E}} \mathbf{V}_m \circ \mathbf{I}_{lm}^* = \mathbf{V}_m \circ \mathbf{i}_m^* + \sum_{n:(m,n) \in \mathcal{E}} \mathbf{V}_m \circ \mathbf{I}_{mn}^*. \quad (3.39)$$

The \mathbf{V}_m term inside the summation on the RHS is substituted using (3.34):

$$\begin{aligned} \sum_{l:(l,m) \in \mathcal{E}} \mathbf{V}_m \circ \mathbf{I}_{lm}^* &= \mathbf{V}_m \circ \mathbf{i}_m^* + \dots \\ &\sum_{n:(m,n) \in \mathcal{E}} \mathbf{V}_n \circ \mathbf{I}_{mn}^* + (\mathbf{Z}_{mn} \mathbf{I}_{mn}) \circ \mathbf{I}_{mn}^*. \end{aligned} \quad (3.40)$$

Here, we define the complex power phasor on phase ϕ entering node n on line (m, n) as $\mathbf{S}_{mn}^\phi = \mathbb{V}_n^\phi (\mathbb{I}_{mn}^\phi)^*$, and the 3×1 vector of complex power phasors entering node n on line (m, n) as $\mathbf{S}_{mn} = [\mathbb{S}_{mn}^a, \mathbb{S}_{mn}^b, \mathbb{S}_{mn}^c]^T = \mathbf{V}_n \circ \mathbf{I}_{mn}^*$. The complex load at node m on phase ϕ is defined as $s_m^\phi = \mathbb{V}_m^\phi (i_m^\phi)^*$, and the 3×1 vector of complex load phasors at node m is $\mathbf{s}_m = [s_m^a, s_m^b, s_m^c]^T = \mathbf{V}_m \circ \mathbf{i}_m^*$. We now rewrite (3.40):

$$\sum_{l:(l,m) \in \mathcal{E}} \mathbf{S}_{lm} = \mathbf{s}_m + \sum_{n:(m,n) \in \mathcal{E}} \mathbf{S}_{mn} + \mathbf{L}_{mn}. \quad (3.41)$$

The term $\mathbf{L}_{mn} \in \mathbf{C}^{3 \times 1}$ represents nonlinear losses on the line. As in [87, 173, 17], we assume that losses are negligible compared to line flows, so that $|\mathbf{L}_{mn}^\phi| \ll |\mathbf{S}_{mn}^\phi| \forall (m, n) \in \mathcal{E}$. Thus, we neglect line losses, linearizing (3.41) into:

$$\sum_{l:(l,m) \in \mathcal{E}} \mathbf{S}_{lm} \approx \mathbf{s}_m + \sum_{n:(m,n) \in \mathcal{E}} \mathbf{S}_{mn}. \quad (3.42)$$

Voltage Magnitude Equations

In this section, we derive a relation between squared voltage magnitudes and complex multiphase power for unbalanced systems. The reader should note that here we present the derivation for a line with three phases, where $\mathcal{P}_{mn} = \{a, b, c\}$. For lines with less than three phases ($|\mathcal{P}_{mn}| \leq 2$), (3.43) - (3.54) should be indexed by \mathcal{P}_{mn} as (3.35) is.

To start, we consider a line $(m, n) \in \mathcal{E}$, and take the Hadamard Product of (3.34) and its (non-transposed) complex conjugate:

$$\mathbf{V}_m \circ \mathbf{V}_m^* = (\mathbf{V}_n + \mathbf{Z}_{mn} \mathbf{I}_{mn}) \circ (\mathbf{V}_n + \mathbf{Z}_{mn} \mathbf{I}_{mn})^*. \quad (3.43)$$

This can be rewritten by distributing the terms on the RHS:

$$\begin{aligned} \mathbf{V}_m \circ \mathbf{V}_m^* &= \mathbf{V}_n \circ \mathbf{V}_n^* + \mathbf{V}_n \circ (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* + \dots \\ &\quad (\mathbf{Z}_{mn} \mathbf{I}_{mn}) \circ \mathbf{V}_n^* + (\mathbf{Z}_{mn} \mathbf{I}_{mn}) \circ (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* . \end{aligned} \quad (3.44)$$

Here we define the real scalar $y_n^\phi = |\mathbb{V}_n^\phi|^2 = \mathbb{V}_n^\phi (\mathbb{V}_n^\phi)^*$, the 3×1 real vector $\mathbf{y}_n = [y_n^a, y_n^b, y_n^c]^T = \mathbf{V}_n \circ \mathbf{V}_n^*$, and the 3×1 real vector $\mathbf{H}_{mn} = (\mathbf{Z}_{mn} \mathbf{I}_{mn}) \circ (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* = (\mathbf{V}_m - \mathbf{V}_n) \circ (\mathbf{V}_m - \mathbf{V}_n)^*$. With these definitions, we also take advantage of the commutative property of the HP and group the second and third terms of the RHS of (3.44) inside the real operator:

$$\mathbf{y}_m = \mathbf{y}_n + 2 \operatorname{Re} \{ (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n \} + \mathbf{H}_{mn} . \quad (3.45)$$

At this point, we focus on the terms inside the real operator for clarity of presentation, and rewrite them as:

$$\begin{aligned} (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n &= \dots \\ &\begin{bmatrix} \mathbb{V}_n^a (Z_{mn}^{aa} \mathbb{I}_{mn}^a + Z_{mn}^{ab} \mathbb{I}_{mn}^b + Z_{mn}^{ac} \mathbb{I}_{mn}^c)^* \\ \mathbb{V}_n^b (Z_{mn}^{ba} \mathbb{I}_{mn}^a + Z_{mn}^{bb} \mathbb{I}_{mn}^b + Z_{mn}^{bc} \mathbb{I}_{mn}^c)^* \\ \mathbb{V}_n^c (Z_{mn}^{ca} \mathbb{I}_{mn}^a + Z_{mn}^{cb} \mathbb{I}_{mn}^b + Z_{mn}^{cc} \mathbb{I}_{mn}^c)^* \end{bmatrix} . \end{aligned} \quad (3.46)$$

With the definition of complex current on a line, $\mathbb{I}_{mn}^\phi = (\mathbb{S}_{mn}^\phi / \mathbb{V}_n^\phi)^*$, and defining the *voltage ratio* as $\gamma_n^{\phi\psi} = \mathbb{V}_n^\phi / \mathbb{V}_n^\psi$, we rewrite (3.46):

$$\begin{aligned} (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n &= \dots \\ &\begin{bmatrix} (Z_{mn}^{aa})^* \mathbb{S}_{mn}^a + \gamma_n^{ab} (Z_{mn}^{ab})^* \mathbb{S}_{mn}^b + \gamma_n^{ac} (Z_{mn}^{ac})^* \mathbb{S}_{mn}^c \\ \gamma_n^{ba} (Z_{mn}^{ba})^* \mathbb{S}_{mn}^a + (Z_{mn}^{bb})^* \mathbb{S}_{mn}^b + \gamma_n^{bc} (Z_{mn}^{bc})^* \mathbb{S}_{mn}^c \\ \gamma_n^{ca} (Z_{mn}^{ca})^* \mathbb{S}_{mn}^a + \gamma_n^{cb} (Z_{mn}^{cb})^* \mathbb{S}_{mn}^b + (Z_{mn}^{cc})^* \mathbb{S}_{mn}^c \end{bmatrix} . \end{aligned} \quad (3.47)$$

The 3×1 vector on the RHS of (3.47) can be separated into a 3×3 matrix multiplying the 3×1 vector \mathbf{S}_{mn} :

$$\begin{aligned} (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n &= \dots \\ &\begin{bmatrix} (Z_{mn}^{aa})^* & \gamma_n^{ab} (Z_{mn}^{ab})^* & \gamma_n^{ac} (Z_{mn}^{ac})^* \\ \gamma_n^{ba} (Z_{mn}^{ba})^* & (Z_{mn}^{bb})^* & \gamma_n^{bc} (Z_{mn}^{bc})^* \\ \gamma_n^{ca} (Z_{mn}^{ca})^* & \gamma_n^{cb} (Z_{mn}^{cb})^* & (Z_{mn}^{cc})^* \end{bmatrix} \begin{bmatrix} \mathbb{S}_{mn}^a \\ \mathbb{S}_{mn}^b \\ \mathbb{S}_{mn}^c \end{bmatrix} . \end{aligned} \quad (3.48)$$

We use the definition of the HP to factor the 3×3 matrix into two 3×3 matrices as in (3.49), where Γ_n is the 3×3 matrix to the left of the Hadamard Product symbol (\circ) within the parentheses on the RHS:

$$\begin{aligned} (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n &= (\Gamma_n \circ \mathbf{Z}_{mn}^*) \mathbf{S}_{mn} = \dots \\ &\left(\begin{bmatrix} 1 & \gamma_n^{ab} & \gamma_n^{ac} \\ \gamma_n^{ba} & 1 & \gamma_n^{bc} \\ \gamma_n^{ca} & \gamma_n^{cb} & 1 \end{bmatrix} \circ \begin{bmatrix} Z_{mn}^{aa} & Z_{mn}^{ab} & Z_{mn}^{ac} \\ Z_{mn}^{ba} & Z_{mn}^{bb} & Z_{mn}^{bc} \\ Z_{mn}^{ca} & Z_{mn}^{cb} & Z_{mn}^{cc} \end{bmatrix}^* \right) \begin{bmatrix} \mathbb{S}_{mn}^a \\ \mathbb{S}_{mn}^b \\ \mathbb{S}_{mn}^c \end{bmatrix} . \end{aligned} \quad (3.49)$$

Placing (3.49) back into (3.45) gives:

$$\mathbf{y}_m = \mathbf{y}_n + 2 \operatorname{Re} \{ (\Gamma_n \circ \mathbf{Z}_{mn}^*) \mathbf{S}_{mn} \} + \mathbf{H}_{mn}. \quad (3.50)$$

Finally, we separate the complex power vector into its active and reactive components, $\mathbf{S}_{mn} = \mathbf{P}_{mn} + j\mathbf{Q}_{mn}$, and apply the real operator on the RHS to obtain

$$\begin{aligned} \mathbf{y}_m &= \mathbf{y}_n + 2\mathbf{M}_{mn}\mathbf{P}_{mn} - 2\mathbf{N}_{mn}\mathbf{Q}_{mn} + \mathbf{H}_{mn} \\ \mathbf{M}_{mn} &= \operatorname{Re} \{ \Gamma_n \circ \mathbf{Z}_{mn}^* \}, \quad \mathbf{N}_{mn} = \operatorname{Im} \{ \Gamma_n \circ \mathbf{Z}_{mn}^* \}. \end{aligned} \quad (3.51)$$

We have derived equations that govern the relationship between squared voltage magnitudes and complex power flow across line (m, n) . This nonlinear and nonconvex system is difficult to directly incorporate into an OPF formulation without the use of convex relaxations. Following the analysis in [87], we apply two approximations. The first is that the higher order term \mathbf{H}_{mn} , which is the change in voltage associated with losses, is negligible, such that $\mathbf{H}_{mn} \approx [0, 0, 0]^T \quad \forall (m, n) \in \mathcal{E}$. The second assumes that node voltages are ‘‘nearly balanced’’ (i.e. approximately equal in magnitude and 120° apart). This is only applied to Γ_n in the RHS of (3.51), such that $\gamma_n^{ab} = \gamma_n^{bc} = \gamma_n^{ca} \approx \alpha$, and $\gamma_n^{ac} = \gamma_n^{ba} = \gamma_n^{cb} \approx \alpha^2$ for all $n \in \mathcal{N}$. Under these assumptions, Γ_n becomes:

$$\Gamma_n \approx A = \begin{bmatrix} 1 & \alpha & \alpha^2 \\ \alpha^2 & 1 & \alpha \\ \alpha & \alpha^2 & 1 \end{bmatrix} \quad \forall n \in \mathcal{N}, \quad (3.52)$$

where $\alpha = 1 \angle 120^\circ = \frac{1}{2}(-1 + j\sqrt{3})$ and $\alpha^2 = \alpha^{-1} = \alpha^* = 1 \angle 240^\circ = \frac{1}{2}(-1 - j\sqrt{3})$. Note that we make the ‘‘nearly balanced’’ assumption in the process of the formal derivation, but that does not imply that the voltages need to actually be balanced for the linearization to be valid.

Applying these approximations for \mathbf{H}_{mn} and Γ_n to (3.51), we arrive at a linear system of equations:

$$\mathbf{y}_m \approx \mathbf{y}_n + 2\mathbf{M}_{mn}\mathbf{P}_{mn} - 2\mathbf{N}_{mn}\mathbf{Q}_{mn}, \quad (3.53)$$

$$\mathbf{M}_{mn} = \operatorname{Re} \{ A\mathbf{Z}_{mn}^* \}, \quad \mathbf{N}_{mn} = \operatorname{Im} \{ A\mathbf{Z}_{mn}^* \}. \quad (3.54)$$

The matrices \mathbf{M}_{mn} and \mathbf{N}_{mn} are modified impedance matrices, where the off-diagonal elements are rotated by $\pm 120^\circ$ (see (3.52)). The diagonal entries of \mathbf{M}_{mn} are $r_{mn}^{\phi\phi}$. Off-diagonal entries of \mathbf{M}_{mn} are $\frac{1}{2}(-r_{mn}^{\phi\psi} + \sqrt{3}x_{mn}^{\phi\psi})$ for $(\phi, \psi) \in \{ab, bc, ca\}$, and $\frac{1}{2}(-r_{mn}^{\phi\psi} - \sqrt{3}x_{mn}^{\phi\psi})$ for $(\phi, \psi) \in \{ac, ba, cb\}$. Diagonal entries of \mathbf{N}_{mn} are $-x_{mn}^{\phi\phi}$. Off-diagonal entries of \mathbf{N}_{mn} are $\frac{1}{2}(x_{mn}^{\phi\psi} + \sqrt{3}r_{mn}^{\phi\psi})$ for $(\phi, \psi) \in \{ab, bc, ca\}$, and $\frac{1}{2}(x_{mn}^{\phi\psi} - \sqrt{3}r_{mn}^{\phi\psi})$ for $(\phi, \psi) \in \{ac, ba, cb\}$.

Motivating Intermezzo

Now that we have set up the necessary equations, it is possible to further demonstrate the need of controlling phase angle for switching operations. Consider the following illustrative

example, where we assume a single phase line and therefore omit superscripts denoting phase. For two nodes m and n that are not connected by a line, no current flows between the nodes. However should a switch between the two nodes be closed, the power at node n is $\mathbf{S}_{mn} = \mathbb{V}_n (\mathbb{V}_m - \mathbb{V}_n)^* Y_{mn}^*$, where $Y_{mn} = g_{mn} + jb_{mn}$ is the admittance of the line. Assuming the voltage magnitudes at nodes m and n are equal, the line power at node n can be written as:

$$\begin{aligned} \mathbf{S}_{mn} &= |\mathbb{V}_n|^2 (g_{mn}(\cos(\theta_{mn}) - 1) - b_{mn} \sin(\theta_{mn})) \dots \\ &\quad + j |\mathbb{V}_n|^2 (b_{mn}(1 - \cos(\theta_{mn})) - g_{mn} \sin(\theta_{mn})) , \end{aligned}$$

where $\theta_{mn} = \theta_m - \theta_n$. It is clear that even with equal voltage magnitudes at nodes m and n , larger voltage angles differences will cause increased real and reactive power flows. This highlights the importance of the ability to control voltage angle, which present OPF formulations lack.

Voltage Phase Angle Equations

We now derive an extension of the power and voltage magnitude system that relates differences in voltage angles between adjacent nodes to complex power flows. This derivation builds heavily upon the analysis of Section 3.6.

The derivation presented here represents a three phase line, $\mathcal{P}_{mn} = \{a, b, c\}$. For lines with less than three phases ($|\mathcal{P}_{mn}| \leq 2$), all equations should be indexed by \mathcal{P}_{mn} as (3.35) is.

We begin with the Hadamard Product of \mathbf{V}_n and the complex conjugate of (3.34):

$$\mathbf{V}_m^* \circ \mathbf{V}_n = \mathbf{V}_n^* \circ \mathbf{V}_n + (\mathbf{Z}_{mn} \mathbf{I}_{mn})^* \circ \mathbf{V}_n . \quad (3.55)$$

From the analysis in Section 3.6, we substitute both terms on the RHS, and expand the LHS with the polar representations of voltage phasors:

$$\begin{bmatrix} |\mathbb{V}_m^a| |\mathbb{V}_n^a| \angle (-\theta_m^a + \theta_n^a) \\ |\mathbb{V}_m^b| |\mathbb{V}_n^b| \angle (-\theta_m^b + \theta_n^b) \\ |\mathbb{V}_m^c| |\mathbb{V}_n^c| \angle (-\theta_m^c + \theta_n^c) \end{bmatrix} = \mathbf{y}_n + (\Gamma_n \circ \mathbf{Z}_{mn}^*) \mathbf{S}_{mn} . \quad (3.56)$$

We negate (3.56) and take the imaginary component of both sides:

$$\begin{aligned} \begin{bmatrix} |\mathbb{V}_m^a| |\mathbb{V}_n^a| \sin(\theta_m^a - \theta_n^a) \\ |\mathbb{V}_m^b| |\mathbb{V}_n^b| \sin(\theta_m^b - \theta_n^b) \\ |\mathbb{V}_m^c| |\mathbb{V}_n^c| \sin(\theta_m^c - \theta_n^c) \end{bmatrix} &= -\mathbf{Im} \{ (\Gamma_n \circ \mathbf{Z}_{mn}^*) \mathbf{S}_{mn} \} \\ \dots &= -\mathbf{N}_{mn} \mathbf{P}_{mn} - \mathbf{M}_{mn} \mathbf{Q}_{mn} . \end{aligned} \quad (3.57)$$

where \mathbf{M}_{mn} and \mathbf{N}_{mn} are defined as in (3.51).

Inspection of the voltage angle equation reveals some interesting similarities compared to the voltage magnitude equations (3.51). The RHS of (3.51) and (3.57) are the real and imaginary parts of the same argument (except for a scaling factor of one-half).

To simplify (3.57), we apply the same assumptions as considered for the magnitude equations. As in the previous section, we assume voltages are “nearly balanced”, with $\Gamma_n \approx A$ as in (3.52). Second we assume that $\theta_m^\phi - \theta_n^\phi$ is sufficiently small such that the small angle approximation holds, so that $\sin(\theta_m^\phi - \theta_n^\phi) \approx \theta_m^\phi - \theta_n^\phi \forall \phi \in \mathcal{P}_{mn}, \forall (m, n) \in \mathcal{E}$. Lastly, we fix all voltage magnitudes on the LHS of (3.57) to unity, so that $|\mathbb{V}_m^\phi| = |\mathbb{V}_n^\phi| = 1 \forall \phi \in \mathcal{P}_{mn}, \forall (m, n) \in \mathcal{E}$. With these three assumptions applied to (3.57), we arrive at:

$$\Theta_m \approx \Theta_n - \mathbf{N}_{mn} \mathbf{P}_{mn} - \mathbf{M}_{mn} \mathbf{Q}_{mn}, \quad (3.58)$$

with $\Theta_m = [\theta_m^a, \theta_m^b, \theta_m^c]^T$, and \mathbf{M}_{mn} and \mathbf{N}_{mn} defined by (3.54). Note that we make the “nearly balanced” and unity voltage magnitude assumptions in the process of formal derivation, but that does not imply that voltages actually be perfectly balanced or have a magnitude of exactly 1 in practice. The accuracy of these approximations in modeling system power flows and voltages will be explored in [178].

Linearized Unbalanced Power Flow Model

We now present the full set of equations that comprise a linearized model for unbalanced power flow. We named these the *LinDist3Flow* equations [178]. Equations for lines $(m, n) \in \mathcal{E}$, (3.61) - (3.63), should be indexed by line phases \mathcal{P}_{mn} as in (3.35), as defined in Section 3.6.

Per phase node complex load:

$$\mathbf{s}_m^\phi (y_m^\phi) = \left(\beta_{S,m}^\phi + \beta_{Z,m}^\phi y_m^\phi \right) d_m^\phi + u_m^\phi - j c_m^\phi, \forall \phi \in \mathcal{P}_m, \forall m \in \mathcal{N} \quad (3.59)$$

Branch power flow:

$$\sum_{l:(l,m) \in \mathcal{E}} \mathbf{S}_{lm} \approx \mathbf{s}_m + \sum_{n:(m,n) \in \mathcal{E}} \mathbf{S}_{mn}, \forall m \in \mathcal{N} \quad (3.60)$$

Magnitude and angle equations:

$$[\mathbf{y}_m \approx \mathbf{y}_n + 2\mathbf{M}_{mn} \mathbf{P}_{mn} - 2\mathbf{N}_{mn} \mathbf{Q}_{mn}]_{\mathcal{P}_{mn}}, \forall (m, n) \in \mathcal{E}, \quad (3.61)$$

$$[\Theta_m \approx \Theta_n - \mathbf{N}_{mn} \mathbf{P}_{mn} - \mathbf{M}_{mn} \mathbf{Q}_{mn}]_{\mathcal{P}_{mn}}, \forall (m, n) \in \mathcal{E}, \quad (3.62)$$

$$\text{with } [\mathbf{M}_{mn} = \mathbf{Re}\{AZ_{mn}^*\}, \mathbf{N}_{mn} = \mathbf{Im}\{AZ_{mn}^*\}]_{\mathcal{P}_{mn}} \quad (3.63)$$

The accuracy of the approximations in the power and voltage magnitude equations has been investigated in [87] and [173]. In [178], we perform a Monte Carlo analysis to explore the level of error introduced by the voltage angle equation assumptions.

3.7 Optimal Power Flow in Unbalanced Networks

In addition to single-phase experiments, we will also consider some experiments on three-phase unbalanced networks.

Voltage Balancing

Although the *LinDist3Flow* equations allow for a flexible OPF framework comparable to the one presented for single-phase networks above, here we focus on one instance to demonstrate the generalization of decentralized OPF to three-phase unbalanced systems. We present an instance first presented in [177] that demonstrated OPF for balancing voltage magnitudes on an unbalanced radial distribution feeder. Voltage balancing is an important objective as many three phase loads (induction motors for instance) are sensitive to high levels of imbalance. Furthermore, many 3-phase voltage regulation equipment actuates based solely on single phase measurements. Therefore, significant levels of imbalance can lead to improper operation of these devices. This simulation incorporated both constant power and constant impedance loads as defined in [107] for the 13 node case. The problem was defined as

$$\begin{aligned} \min_{y_n^\phi, P_{mn}^\phi, Q_{mn}^\phi, u_n^{p,\phi}, u_n^{q,\phi}} \quad & \sum_{n \in \mathcal{N}} \left[\sum_{\phi \neq \psi} (y_n^\phi - y_n^\psi)^2 + \rho \sum_{\phi} |u_n^\phi|^2 \right] \\ \text{subject to} \quad & (3.59) - (3.63), \end{aligned} \tag{3.64}$$

$$\begin{aligned} \underline{y} &\leq y_n^\phi \leq \bar{y}, \\ |u_n^\phi| &\leq \bar{u}, \quad \forall \phi \in \mathcal{P}_n, \forall n \in \mathcal{N}. \end{aligned} \tag{3.65}$$

where $\phi, \psi \in \{a, b, c\}$. The parameter ρ was chosen as 0.01 and is a penalty on control action. [177] showed the effectiveness of the *LinDist3Flow* model in drastically reducing voltage imbalance across phases.

Network Reconfiguration

We adapt the OPF formulation from [178]. In this paper, the objective was to close a switch in order to connect two networks on line. To minimize large instantaneous power transfers across the switch upon closing, one needs to match the voltage phasors at the ends of the open switch. To this end, [178] proposed the following OPF to minimize the voltage phasor difference between one or more nodes and the respective reference at each node, while providing feeder voltage support:

$$\begin{aligned} \min_{y_n^\phi, \theta_n^\phi, P_n^\phi, Q_n^\phi, u_n^{p,\phi}, u_n^{q,\phi}} \quad & \rho_y C_y + \rho_\theta C_\theta + \rho_u C_u \\ \text{subject to} \quad & (3.59) - (3.63), \\ & \underline{y} \leq y_n^\phi \leq \bar{y}, \quad \forall \phi \in \mathcal{P}_n, \forall n \in \mathcal{N}, \\ & |u_n^\phi| \leq \bar{u}_n^\phi, \quad \forall \phi \in \mathcal{P}_n, \forall n \in \mathcal{N}, \\ & \mathbf{y}_\infty = [1, 1, 1]^T, \\ & \Theta_\infty = [0, -2\pi/3, 2\pi/3]^T. \end{aligned} \tag{3.66}$$

$$C_y = \sum_{\phi \in \mathcal{P}_{1680,2680}} \left(y_{1680}^\phi - y_{2680}^\phi \right)^2, \quad (3.67)$$

$$C_\theta = \sum_{\phi \in \mathcal{P}_{1680,2680}} \left(\theta_{1680}^\phi - \theta_{2680}^\phi \right)^2, \quad (3.68)$$

$$C_u = \sum_{n \in \mathcal{N}} \sum_{\phi \in \mathcal{P}_n} |u_n^\phi|^2. \quad (3.69)$$

The OPF objective function is a weighted sum of three terms: C_y is the sum of squared voltage magnitude differences squared, C_θ is the sum of voltage angle differences squared, and C_u is the sum of the squared magnitudes of all DER dispatch, to avoid applying excessive amounts of control. Constraints of lower and upper voltage magnitude bounds were imposed as $0.95 \leq |V_n^\phi| \leq 1.05 \forall \phi \in \mathcal{P}_n, \forall n \in \mathcal{N}$ such that $\underline{y} = 0.9025$ and $\bar{y} = 1.1025$. Additionally, DER dispatch is constrained by its apparent power capacity, \bar{u}_n^ϕ .

Chapter 4

Forecasting-Based State Estimation in Single- and Three-Phase Distribution Systems

“Under Bayes’ theorem, no theory is perfect. Rather, it is a work in progress, always subject to further refinement and testing.”

- Nate Silver

Submitted to the IEEE Transactions on Power Systems [58]. Earlier work presented at the 7th International Conference on Cyber-Physical Systems (ICCPS) in Vienna, Austria, April 2016 [59].

State Estimation is an essential technique to provide observability in power systems. Traditionally developed for high-voltage transmission networks, state estimation requires equipping networks with many real-time sensors, which remains a challenge at the scale of distribution networks. This chapter proposes a method to complement a limited set of real-time measurements with voltage predictions from forecast models. The method differs from the classical weighted least-squares approach, and instead relies on Bayesian estimation formulated as a linear least squares estimation problem. We integrate recently developed linear models for unbalanced 3-phase power flow to construct voltage predictions as a linear mapping of load predictions. The estimation step is a linear computation allowing high resolution state estimate updates, for instance by exploiting a small set of phasor measurement units. Uncertainties can be determined a priori and smoothed a posteriori making the method useful for both planning, operation and post hoc analysis. The method is applied to an IEEE benchmark and on a real network testbed at the Dutch utility Alliander. An observability analysis suggests strategies for optimal sensor placement.

4.1 Introduction

Motivation

The operation of electric distribution networks is faced with new challenges due to the rapid adoption of distributed generation (DG) and electrification of our society, such as in driving and heating. The inherent intermittency of renewable generation combined with the diversification of demand make power flow more variable and harder to predict, leading to new protection issues, such as unintended islanding or tripping [119], and economic burden due to accelerated wear [182]. When DG is connected to feeders without proper monitoring functionality to assess its effect, cost of integration can easily multiply by a factor 3 to 4 [205, 72], and the rapid adoption of EVs will further aggravate this problem [93]. Traditional operating paradigms with little sensing make it hard to measure and track the impact on the physical network, causing utilities to impose conservative caps on the allowable DG capacity and number of electric vehicles (EVs), hindering the transition to more renewable energy sources.

To understand and mitigate these risks, many Distribution System Operators (DSOs) are building a stronger information layer on top of their physical infrastructure that exploits recent advances in sensing and communication to enable forecasting and real-time state estimation. Power system state estimation (SE) is the process of leveraging measurement from a subset of states in an electric network to estimate states that are not measured in real-time. In transmission systems, the need for system reliability and economies of scale have long motivated the development of state estimation methods [2, 89]. In the traditional setting, a state estimator relies on an *overdetermined* formulation for the unknown/unmeasured variables to be *observable*. This means that the number of available measurements must be greater than or equal to the number of unknowns (to be estimated). Most methods then rely on weighted least-squares or least absolute value approaches [2].

DSOs are experimenting with two capabilities to improve observability. Firstly, by applying traditional SE methods to distribution systems. This is conceptually similar, but technically and economically a more challenging task. Ensuring observability requires DSOs to equip most buses in a network with real-time sensors and communication infrastructure, leading to steep investments that are hard to scale across all territories. Secondly, by gathering historical data from SCADA and AMI systems to enable forecasting of demand, flow and voltage information. Unfortunately, the increasing variability of power yields probability distributions with long tails, which cause forecasting methods to do poorly in situations when observability is most needed; when extreme and potentially dangerous events happen.

To overcome issues of scalability and prediction error, we combine forecasting based on historical load information with estimation using a limited number of real-time sensors. Our method focuses on estimating the voltage phasor at all buses with no sensor. Voltage forecasting will be done over a slower timescale using load data collected historically and modeling three-phase power flow with novel approximations [178]. We then use Bayesian estimation to update the forecasts in real-time using a small number of sensors with high

temporal resolution. This is motivated by the recent introduction of synchrophasors for distribution systems, which allow the real-time assessment of dynamics [222, 142]. The updated forecasts will then constitute estimated voltage quantities at nodes without sensors at the same high temporal resolution. The method provides DSOs with a critical functionality to correct forecasted values when the system is deviating from the expected scenario, allowing detection and mitigation of risky scenarios.

Previous Work

In conventional state estimation, the measurements $z \in \mathbb{R}^{N_m}$ are expressed as a function of the quantities that are estimated $x \in \mathbb{R}^{N_n}$, by using power flow modeling:

$$z = h(x). \quad (4.1)$$

The state estimation problem is then solved using a weighted least squares (WLS) problem:

$$x^* = \arg \min_x (z - h(x))^T W (z - h(x)). \quad (4.2)$$

For the WLS problem to yield a meaningful result, Equation (4.1) needs to be *overdetermined*, which means that the number of measured variables needs to be greater than the number of estimated variables; $N_m > N_n$. A key challenge in distribution grids is to estimate an N_n -dimensional state vector in scenarios where only a limited set of sensors is available, i.e. $N_m < N_n$, which does not satisfy the requirements for conventional state estimation. As a result, the standard estimation problem is *underdetermined* and hence ill-posed from a computational point of view. In practice, this means that the state vector x is not *observable*. To overcome this inherent challenge, *pseudo-measurements* are typically used to augment real-time measurements in a weighted least squares (WLS) estimation algorithm. Pseudo-measurements are often calculated using load forecasts or historical data that tend to be less accurate than real-time measurements. Initial efforts considered augmenting an already fully observed measurement vector with extra load forecasts [175, 22]. Later efforts tried to use a more limited number of real-time measurements with forecasts from Gaussian Mixture Models [189] or Artificial Neural Networks [141]. There have been many contributions made to enable Distribution System State Estimation based on traditional WLS, and we refer the reader to [162] for a rigorous overview.

Unfortunately, the WLS method requires extensive tuning and is rather sensitive to errors and bad data [90]. Göl and Abur address this challenge through combining WLS with a least absolute value method that is more robust to error, yielding a *hybrid* estimator that combines a limited number of phasor measurement units (PMUs) with a high refresh rate (at the order of 30 Hz), and a fully observed ($M \geq n$) set of SCADA measurements at slower refresh rate (order of 5 to 15 min). The weighted least absolute value method used helps to robustify the the estimate between each SCADA update, but does not address a scenario where a limited set of measurements is available. Furthermore, the estimator is designed for transmission systems that can rely on robust communication networks. SCADA in distribution systems

often lack a reliable communication infrastructure, which can lead to packet failure and unreliable state estimates. Other work by Schenato [179] and Weng [221] study the use of Bayesian estimation for SE, using load statistics to determine a *prior* probabilistic forecast \hat{x} of state variables, which can be updated based on a limited set of real-time measurements. These papers shows their accuracy is comparable to that of conventional WLS estimators, and estimation error confidence intervals can be computed off-line, allowing for engineering trade-offs between number of sensors and estimation accuracy.

Contributions

This work forms a bridge between forecasting and full state estimation in three-phase distribution systems by embracing a Bayesian approach. Inspired by the development of linear approximations for unbalanced three-phase power flow [178], we derive a *closed-form analytic state estimator* that takes as its inputs load forecasting information, a network model and real-time measurements from a limited set of sensors. Our method estimates *voltage phasor differences* rather than on the absolute phasor. This reduces modeling and forecast errors and solves the issue of having to stop a numerical algorithm to solve power flow equations as done in [179]. This also circumvents the need for a reference voltage (typically the feeder head as in [179]) and allows the algorithm to be implemented in a distributed fashion for different parts of the network. In addition, the estimator in [179] models the voltage and current phasor in rectangular form, i.e. the complex voltage at node i is formulated as $\mathbb{V}_i = \text{Re}(\mathbb{V}) + j\text{Im}(\mathbb{V})$, where $j := \sqrt{-1}$. Since available synchrophasors typically report in polar form, i.e. $\mathbb{V}_i = V_i \angle \delta_i$, a nonlinear transformation is needed to adjust sensor readings before feeding these to the estimator: $\text{Re}(\mathbb{V}) = V_i \cos \delta_i$ and $\text{Im}(\mathbb{V}) = V_i \sin \delta_i$. This can lead to undesirable magnification of measurement errors in the voltage angle δ_i , which can be problematic for applications where angle estimates are used, for instance to close switches [178]. Our method formulates the problem in polar form, bypassing this source of error. Applied and assessed on a specific IEEE test feeder, we show that the method reduces average the error of forecasts by an average 60%, with more dramatic improvements for specific buses where forecasts are not able to perform appropriately. Lastly, we implement the full method on a utility testbed showcasing its applicability in real world circumstances.

Notation

We use $\|\cdot\|$ to denote the ℓ_2 -norm, $(\cdot)^*$ to denote an optimal value, and $^\top$ stands for the transpose operator. $\mathfrak{N}(\mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 . Throughout this work, we use the symbol \circ to represent the Hadamard Product of two matrices (or vectors) of the same dimension, also known of the element-wise product, such that:

$$C = A \circ B = B \circ A \Rightarrow C_{ij} = A_{ij}B_{ij} = B_{ij}A_{ij}$$

where i indicates the row and j indicates the column of the vector or matrix.

4.2 Methodology




Figure 4.1: Overview of the forecasting and state estimation methodology.

This paper proposes a data-driven approach to do state estimation, that relies on minimum mean squares estimation (MMSE) [218]. MMSE is related to weighted least squares, but grounded in Bayesian principles and does not require an overdetermined measurement equation. Instead, our method relies on linear power flow models that enable us to express *voltage differences* (both magnitude and angle) throughout a network as a function of nodal load and generation. By expressing both the measured differences and the estimated differences as a function of the load, we are able to set up a linear least squares estimation (LLSE) problem, the linear version of MMSE. The LLSE has an analytical solution that can update the forecast of non-measured voltage differences by comparing the measured voltage differences against their forecasted values. As such, the method reminds of the Kalman Filter [115], which is a repeated execution of LLSE problems taking into account the dynamic evolution of the state variables. The approach comes with a trade-off, as the quality of the updates depends on the number of sensors and their placement in the network. This is discussed in Section 4.6.

The MMSE approach enables an end-to-end pipeline from historical load and network data to voltage forecasting to updating these forecasts in real-time using a limited set of sensors. The methodology is depicted in Figure 4.1. The three main steps of forecasting, modeling and real-time estimation are developed in Section 4.3, Chapter 3 and Section 4.4 respectively. Before we dissect these steps we first cover the sources of uncertainty in state estimation and we introduce MMSE.

Sources of information and uncertainty for state estimation

The practical reality of DSOs is that they face many sources of uncertainty in trying to construct a state estimator, outlined in Figure 4.2. Following the proposed construction of the state estimator as depicted in Figure 4.1, the overall accuracy of the available information for state estimation depends on three sources of uncertainty: accuracy of the forecasted quantities, of the modeling procedure and of the quantities measured in real-time.

The forecast may be based on a DSO’s historical data, which can include SCADA data of network variables, advanced metering infrastructure (AMI) readings of household consumption (or an anonymized/aggregated version of these), data of distributed generation and storage,

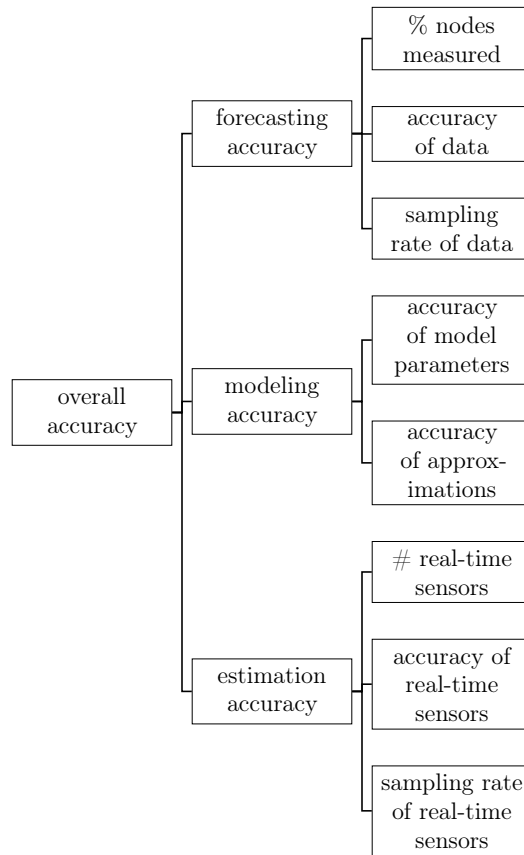


Figure 4.2: Sources of uncertainty affecting the accuracy of the state estimator

public weather data (temperature, humidity, solar irradiance). These data sources typically do not form a perfect representation to forecast all the necessary quantities in the network. Certain nodes may not be recorded, the recordings may be noisy or miss certain data points, and the sampling rate of the recordings may be lower than the anticipated rate for updating the state estimator. In the modeling step, inaccuracy arises from parameter data that is outdated or measured with noise or the use of approximations (such as linear power flow) to enable efficient computation. Lastly, for the actual estimation, we rely on a limited number of sensors, which may be subject to measurement noise and could have various sampling rates.

Introduction to Minimum Mean Square Estimation

Consider the context of having a set of voltage phasor measurements $Z \in \mathbb{R}^{N_m}$ and an unobserved random variable $X \in \mathbb{R}^{N_n}$, representing all non-measured voltage phasors. We aim to determine an estimate of X based on Z that is close to X in some sense. Assume we are given a joint distribution of (X, Z) . We want to find an estimator $\hat{X} = g(Z)$ that minimizes the mean square error $E[\|X - \hat{X}\|^2]$. One can show that the minimum mean

squares estimate (MMSE) of X given Z is equivalent to the conditional expectation, i.e. $\hat{X} = E[X|Z]$ [218].

We consider the case in which both the estimator and the measurements are linear in a shared set of variables for which distributions are available, in our case in the form of load statistics. Let (X, Z) be vectors of random variables on some probability space. It turns out that the estimator minimizing the mean square error is also linear in the measurements, i.e. the *linear least squares estimator* (LLSE) denoted $L[X|Z] = a + BZ$ is a linear function of Z , with $a \in \mathbb{C}^{N_n}$ and $B \in \mathbb{C}^{N_n \times N_m}$ attained by the following optimization problem:

$$\min_{a, B} E [\|X - a - BZ\|^2] . \quad (4.3)$$

Following [218], it can be shown that

$$a^* = E[X] - \Sigma_{X,Z} \Sigma_Z^{-1} E[Z] , \quad B^* = \Sigma_{X,Z} \Sigma_Z^{-1} , \quad (4.4)$$

where $\Sigma_{X,Z} \in \mathbb{R}^{N_n \times N_m}$ and $\Sigma_Z \in \mathbb{R}^{N_m \times N_m}$ denote the cross-covariance matrix of X and Z and covariance matrix of Z . Hence, if Σ_Z is nonsingular, the LLSE is given by

$$L[X|Z] = E[X] + \Sigma_{X,Z} \Sigma_Z^{-1} (Z - E[Z]) . \quad (4.5)$$

Lastly, if Σ_Z is singular, then Σ_Z^{-1} can be replaced by the left pseudo-inverse Σ_Z^\dagger . Interpreting (4.5), $(Z - E[Z])$ represents a deviation of the actual measurement Z from its expected value $E[Z]$, which is called an *innovation*. This innovation triggers the Bayesian estimator $L[X|Z]$ to propose an update of the forecast $E[X]$ by a linear scaling through the covariance matrices. Alternatively, $L[X|Z] = g(Z)$ can be interpreted as a projection of X onto the set of affine functions of Z .

The LLSE has a number of important benefits. Firstly, it has an analytical closed-form solution that can be used to neatly integrate real-time measurements Z and forecast information (as we will see in Section 4.3). Secondly, it is not necessary to explicitly calculate the Bayesian posterior probability density function over X , because $L[X|Z]$ only depends on the first two moments of X and Z . Thirdly, it works for many distributions $(X, Z) \sim \mathfrak{D}$, as long as \mathfrak{D} has well defined first and second moments [218]. Lastly, the number of measurements N_m does not need to be larger than the number of to-be-estimated states N_n , which is the most significant difference with other ubiquitous estimation schemes such as weighted least squares and Gauss-Markov estimation that do not work for $N_m < N_n$. The main challenge of any MMSE approach is understanding what information is lost in the projection that happens in Equation (4.5) through the mapping $\Sigma_{X,Z} \Sigma_Z^{-1}$. For our state estimation method this requires revisiting the notion of network observability, typically defined for situations where $N_m > N_n$, which we will do in Section 4.6, based on the full formulation of the estimator.

4.3 Forecasting

We consider the design of a machine learning model to forecast the mean μ_s and covariance matrix Σ_s of the load \mathbf{s} , which are then used to forecast the mean and covariance of the

voltage magnitude and phase. In practical contexts, DSOs may not have access to voltage or load readings from AMI in real-time, but it is possible that historical readings are used, in combination with other predictive covariates, to predict load values for a future time.

Machine learning models have been used in a variety of ways to predict load values [143]. Two relevant examples are autoregressive moving average (ARMA) models for short term load forecasting and data-driven modeling of physical systems that utilizes regression trees to predict loads, with notable benefits to both. An ARMA model is able to capture trends in previous datapoints [105]. ARMA models are often not practical in distribution operation, since the AMI data is mostly not available in real-time, preventing the use of recent load values. A regression tree model is able to cluster data based on certain characteristics, such as day of the week, temperature, and humidity [19]. Its interpretability makes it useful in contexts where operators need to make decisions based on a model's predictions.

In our setting, the MMSE estimator defined in Section 4.2 necessitates the input of a point estimate of the load and its covariance matrix. This requirement motivates the use of Gaussian Processes (GPs), which offer both mean and variance information [167]. A GP is also flexible in that they can have continuous ARMA features as well as discrete features as its inputs. GPs have previously been used in similar applications for short term load forecasting to predict maximum daily loads [150]. Using GPs does introduce some bias, as load distributions tend to be non-Gaussian, though typically near-unimodal. In our setting, this bias can be compensated by the estimation step. Using a more sophisticated method to retrieve first and second moment information from historical data is left as future work.

Let \mathcal{N} denote a set of buses indexed by $n = 0, 1, \dots, N - 1$, where N is the order (number of buses) of the distribution feeder, and bus 0 denotes the feeder head (or substation). For each node $n \in \mathcal{N}$, we start with a data set of historical readings of inputs $X_n = \{x_n[t] \in \mathcal{X}_n\}_{t=1}^T$ and load values $S_n = \{\mathbf{s}_n[t] \in \mathcal{Y}_n\}_{t=1}^T$. The inputs consist of real-valued and discrete-valued features. We consider the following real-valued features at time t :

$$[l_n[t-1] \cdots l_n[t-k] \quad d_n[t-1] \cdots d_n[t-k+1] \quad \theta[t] \quad \eta[t]] , \quad (4.6)$$

where $l_n[t]$ denotes the load value for bus n at time t , $d_n[t]$ the difference in load between time $t - 1$ and t , and θ_t and η_t are the temperature and humidity at time t . Note that a typical distribution feeder SCADA system often does not have access to load measurement, and hence the features l and d may only be available historically or in real-time for only a subset of the buses. Hence, we also consider discrete-valued features representing date and time:

$$[DST \quad MOY \quad BD \quad DOW \quad HOD \quad MOH] , \quad (4.7)$$

which respectively denote an indicator for daylight saving time, month of year, an indicator for business day, day of week, hour of day and minute of hour.

We now want to train a function $f_n : \mathcal{X}_n \rightarrow \mathcal{Y}_n$ with data that best predicts $\mathbf{s}_n[t]$ at some time t based on an input with accessible inputs $x_n[t]$. A GP defined on an input space \mathcal{X}_n can be formulated as

$$f_n(x) = \boldsymbol{\phi}_n(x_n)^\top \beta_n + g_n(x) , \quad (4.8)$$

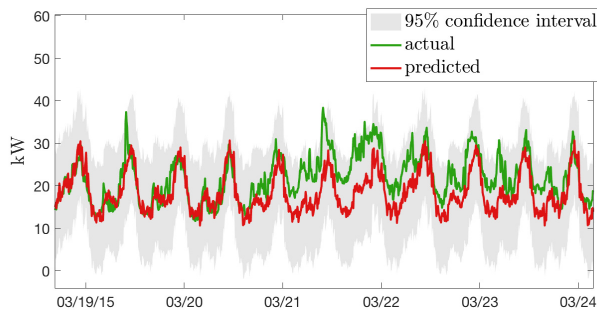


Figure 4.3: Forecast of an aggregate load using a Gaussian Process model with only discrete-valued time features. Only 10% of the loads in the aggregate were recorded in historical data. The other 90% of loads were imputed with the average load profile. Poor forecast performance, such as on March 21st, motivates the use of Bayesian estimation.

where $g_n(x)$ is a zero-mean GP represented as $\mathcal{GP}(0, k_n(x_n, x_n))$, with kernel $k_n(x_n, x_n)$ modeling the covariance across the input space \mathcal{X}_n . $\phi_n(x_n)^\top \beta_n$ determines the translation of the GP from the origin, with $\phi_n(x_n)$ a feature basis for the output given the input vector x_n , β_n are learned coefficients or weights for the basis features [167, Section 2.2]. Given this framework, we can model the distribution of an output at a certain input x_n^* :

$$f(x_n^*) \mid x_n^*, X_n, S_n \sim \mathfrak{N}(\phi_n(x_n^*)^\top \beta_n + g_n(x_n^*), \sigma^2), \quad (4.9)$$

The primary assumption under a GP is that it models a collection of random variables, any finite number of which have a joint Gaussian distribution. Notice that there are two different variances in the system – $k_n(x_n, x_n)$ and σ^2 . The first variance, $k_n(x_n, x_n)$ is the variance on the estimate induced by the covariance of the input features as defined by a covariance function. σ^2 is the noise variance of the data as a whole. To challenge the method, in Section 4.5, we consider a GP model that is based on a poor historical data set and no access to real-valued features. Figure 4.3 exemplifies the resulting forecast accuracy, motivating the use of Bayesian estimation to account for forecast errors such as those experienced on March 21st.

4.4 Real-time Estimation

In this Section, we construct the state estimator based on linear least squares estimation. This method takes in a prior distribution on measured and unmeasured voltage variables, and updates this in real-time with a limited set of measurements. To do so, we require the prior statistics of the voltage based on load forecasts (Section 4.3) and power flow modeling (Chapter 3). We first express measured and unmeasured voltage variables as a linear function of the net load. We can then construct the necessary matrices to express the voltage forecast as function of load statistics.

Voltage as a Function of Net Load

Consider the vector with all the differences in squared voltage magnitude stacked with the differences in voltage angles over all the branches (i.e. for every set of adjacent nodes) in the network,

$$\Delta\mathbf{Y} \triangleq \begin{bmatrix} \mathbf{Y}_0 - \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_{N-1} - \mathbf{Y}_N \end{bmatrix} \in \mathbb{R}^{3N}. \quad (4.10)$$

With Equation (3.61), we can build a model for all the voltage differences over wires throughout the network

$$\Delta\mathbf{Y} = 2 [\text{blkdiag}(\mathbf{M}_{ij}) \text{blkdiag}(\mathbf{N}_{ij})] \begin{bmatrix} \text{vec}(\mathbf{P}_{ij}) \\ \text{vec}(\mathbf{Q}_{ij}) \end{bmatrix} = \mathbf{Z}_b \mathbf{S}, \quad (4.11)$$

where $\mathbf{S} \in \mathbb{R}^{6N}$ is the vector with real and reactive branch flows stacked vertically, and $\mathbf{Z}_b \in \mathbb{R}^{3N \times 6N}$ is a horizontal stack of two block diagonal matrices with the corresponding 3-by-3 matrices from (3.63). With Equation (3.60), we can express the branch flows \mathbf{S} in terms of the nodal net loads, which yields $\mathbf{S} = \mathcal{P}_b \mathbf{s}$, with $\mathbf{s} \in \mathbb{R}^{6N}$ a vector with the nodal net loads, real and reactive power p_n, q_n , $n \in \mathcal{N}$ stacked vertically, and $\mathcal{P}_b \in \mathbb{R}^{6N \times 6N}$ a binary matrix in which a row represents a branch with 1s selecting the nodes downstream of the branch. We have now expressed the differences in voltage magnitude over all N lines in terms of the nodal load vector,

$$\Delta\mathbf{Y} = \mathbf{Z}_b \mathcal{P}_b \mathbf{s} \triangleq \mathbf{Z}_n \mathbf{s}, \quad (4.12)$$

where $\mathbf{Z}_n \triangleq \mathbf{Z}_b \mathcal{P}_b \in \mathbb{R}^{3N \times 6N}$.

Measured quantities

In our actual setting, we do not directly measure voltage differences over all individual wires. Instead, we place the sensors over a distance spanning multiple branches and buses. The voltage difference over the path can be rewritten as the sum of the individual differences of the branches lying on the path. For example, for a later network with 3 buses, the squared voltage magnitude difference over the path from bus 0 through bus 1 to bus 2 can be expressed as

$$\begin{aligned} \mathbf{Y}_0 - \mathbf{Y}_2 &= (\mathbf{Y}_0 - \mathbf{Y}_1) + (\mathbf{Y}_1 - \mathbf{Y}_2) \\ &= 2(\mathcal{M}_{01} \mathbf{P}_1 + \mathbf{N}_{01} \mathbf{Q}_1) + 2(\mathcal{M}_{12} \mathbf{P}_2 + \mathbf{N}_{12} \mathbf{Q}_2) \\ &= 2\mathcal{M}_{01} \mathbf{p}_1 + 2\mathbf{N}_{01} \mathbf{q}_1 + 2(\mathcal{M}_{01} + \\ &\quad \mathcal{M}_{12}) \mathbf{p}_2 + 2(\mathbf{N}_{01} + \mathbf{N}_{12}) \mathbf{q}_2 \end{aligned}$$

Notice that we can form this equation by doing a row operation on Equation (4.12), i.e. adding up rows 1 and 2 of \mathbf{Z}_n . Imagine that we have a set of differences that refer to the

placement of our sensors,

$$\Delta \mathbf{Y}_m \triangleq \begin{bmatrix} \mathbf{Y}_{m_2} - \mathbf{Y}_{m_1} \\ \vdots \\ \mathbf{Y}_{m_M} - \mathbf{Y}_{m_{M-1}} \end{bmatrix} \in \mathbb{R}^{3(M-1)}, \quad (4.13)$$

with $m_1, \dots, m_M \in \mathcal{M}$. We can now formulate the equations, by adding up the differences of all individual lines in between the sensors. I.e., we can formulate a permutation matrix such that $\Delta \mathbf{Y}_m = \mathcal{P}_m \Delta \mathbf{Y}$, and hence

$$\Delta \mathbf{Y}_m = \mathcal{P}_m \mathbf{Z}_n \mathbf{s} = \mathbf{Z}_m \mathbf{s}, \quad (4.14)$$

where $\mathbf{Z}_m \triangleq \mathcal{P}_m \mathbf{Z}_n = \mathcal{P}_m \mathbf{Z}_b \mathcal{P}_b \in \mathbb{R}^{3(M-1) \times 6N}$. This gives us an expression for the measured quantities as a function of the nodal load vector.

Non-measured quantities - Voltage Estimation

We are interested to estimate voltage magnitude and angle at all the $N - M$ buses in the network that are not equipped with a sensor. We aim to do this given a measurement of the voltage phasor at a limited number of M buses in the network, and forecast statistics on the nodal load vector \mathbf{s} . We consider the differences in voltage between a location we want to estimate and a nearby sensor location. These differences are collected in a vector $\Delta \mathbf{Y}_e$ to be estimated as a function of the load vector \mathbf{s} , similar to the construction of the measurement equation:

$$\Delta \mathbf{Y}_e = \mathbf{Z}_e \mathbf{s} \in \mathbb{R}^{3(N+1-M)}, \quad (4.15)$$

where $\mathbf{Z}_e \triangleq \mathcal{P}_m \mathbf{Z}_n = \mathcal{P}_e \mathbf{Z}_b \mathcal{P}_b \in \mathbb{R}^{3(N+1-M) \times 6N}$ is constructed in the same way as \mathbf{Z}_m in (4.14). In order to retrieve an estimate of the absolute voltage value, we can simply take the nearest sensor reading and add/subtract the estimated difference between the location and that sensor location.

Voltage Forecast Statistics

We now have that our measurements are voltage phasor differences, i.e. $z = \Delta \mathbf{Y}_m$ and the estimation quantities are other voltage phasor difference, i.e. $x = \Delta \mathbf{Y}_e$. Given the linear relationships with the load vector \mathbf{s} , we can now derive the statistics on z . The mean of z is

$$\begin{aligned} \mu_z(t) &= E(\Delta \mathbf{Y}_m), \\ &= E(\mathbf{Z}_m \mathbf{s}), \\ &= \mathbf{Z}_m E(\mathbf{s}), \\ &= \mathbf{Z}_m \mu_s(t). \end{aligned} \quad (4.16)$$

Similarly, we have that $\mu_x(t) = E(\Delta \mathbf{Y}_e) = \mathbf{Z}_e \mu_s(t)$. The covariance of z is

$$\begin{aligned}
 \Sigma_z(t) &= E((z - \mu_z)(z - \mu_z)^\top), \\
 &= E((\mathbf{Z}_m \mathbf{s} - \mathbf{Z}_m \mu_s(t))(\mathbf{Z}_m \mathbf{s} - \mathbf{Z}_m \mu_s(t))^\top), \\
 &= \mathbf{Z}_m E((\mathbf{s} - \mu_s(t))(\mathbf{s} - \mu_s(t))^\top) \mathbf{Z}_m^\top, \\
 &= \mathbf{Z}_m \Sigma_s(t) \mathbf{Z}_m^\top.
 \end{aligned} \tag{4.17}$$

Similarly, we have that the cross-covariance of x and z is $\Sigma_{x,z}(t) = \mathbf{Z}_e \Sigma_s(t) \mathbf{Z}_m^\top$. This yields all the statistics we need to construct the distribution grid state estimator.

Constructing the State Estimator

We can now analytically derive the LLSE of \mathbf{Y}_e given measurements \mathbf{Y}_m , as a specific form of Equation (4.5) presented in Section 4.2. For our voltage estimation setting this yields

$$\begin{aligned}
 L[\Delta \mathbf{Y}_e | \Delta \mathbf{Y}_m] &= E(\Delta \mathbf{Y}_e) + \dots \\
 &\quad \Sigma_{\Delta \mathbf{Y}_e, \Delta \mathbf{Y}_m} \Sigma_{\Delta \mathbf{Y}_m}^{-1} (\Delta \mathbf{Y}_m - E(\Delta \mathbf{Y}_m)), \\
 &= \mathbf{Z}_e \mu_s + \dots \\
 &\quad \mathbf{Z}_e \Sigma_s \mathbf{Z}_m^\top (\mathbf{Z}_m \Sigma_s \mathbf{Z}_m^\top)^{-1} (\Delta \mathbf{Y}_m - \mathbf{Z}_m \mu_s),
 \end{aligned} \tag{4.18}$$

where we dropped the time index for brevity. With $L[\Delta \mathbf{Y}_e | \Delta \mathbf{Y}_m]$, the voltage estimates can be retrieved as

$$\hat{\mathbf{V}}_e = \sqrt{\mathbf{Y}_{\text{near}} + L[\Delta \mathbf{Y}_e | \Delta \mathbf{Y}_m]}, \tag{4.19}$$

where \mathbf{V}_e denotes a stacked vector with voltages for all buses without measurement, and \mathbf{Y}_{near} are the squared voltages at the nearest measured bus for each estimated bus.

Notice that (4.18) is written in the form $\Delta \hat{\mathbf{Y}}_e = f(\Delta \mathbf{Y}_m)$, as all of the other information needed to evaluate the estimator are forecast statistics, which are known a priori. As such, (4.18) uses the statistical information of the net loads \mathbf{s} , in combination with the topology and impedance information of the network, to present a closed-form analytical estimator of voltage differences throughout the network. This function is linear in the measurements and efficiently computed in real-time.

4.5 Results

Synthetic experiments

Earlier work implemented the distribution grid state estimator on a single-phase radial network [59]. To validate the estimator on a three-phase network, we used a modified version of the IEEE 37 bus distribution feeder model [107], depicted in Figure 4.4. The feeder voltage and power ratings were left unchanged (4.8 kV and 2.5 MVA), as were line segment configuration assignments. We ignored the transformer at node 775 and the voltage regulator at the feeder head. We assumed all loads were constant power. The data used in this

experiment are from datasets provided by Pecan Street for educational use [50]. The raw data contained 15-minute-interval data sampled from July 1, 2013 to September 26, 2016. We aggregated different household time series from the Pecan Street data set such that the aggregated time series data had a spot load marginally less than the 3-phase real and reactive spot loads defined by the IEEE feeder model [107]. The aggregated time series were then used to build a Gaussian Process forecast model for real and reactive power at each bus, as outlined in Section 4.3. Voltage sensors were placed at nine different buses, indicated by red hexagons in Figure 4.4. Figure 4.5 shows the result for one time instance. To assess the overall performance, we compute the Average Root Mean Square Error (ARMSE) on the voltages \mathbf{V}_e that are not measured,

$$\text{ARMSE}(\{\hat{\mathbf{V}}_e[t]\}_{t=1}^T) = \sqrt{\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{V}}_e[t] - \mathbf{V}_e[t]\|^2}. \quad (4.20)$$

Figure 4.6 shows the ARMSE metric for all buses. It is bounded by 0.2 p.u. for the forecasted values and 0.02 p.u. for the estimated values. Notice that buses with higher forecast errors benefit significantly from the estimation procedure. Buses that already have a proper accuracy on forecasted values of the order < 0.01 p.u. do not necessarily gain much from estimation. This can be attributed to the fact that these errors are in the same order as the modeling errors due to linear approximation, which are carefully studied in a separate paper [178].

Validation experiment on a utility testbed

Here, we apply the method to a network in the territory of *Alliander*, the largest Distribution Network Operator (DNO) of the Netherlands serving over three million customers. Alliander is experimenting with community electricity storage in Rijsenhout a suburban village close to Amsterdam, the Netherlands. The project is called “BuurtBatterij” which translates to “Neighborhood or Community Battery”. Figure 4.7 depicts the Rijsenhout feeder that houses the battery project. One of the goals of the community battery experiments is to assess and improve the accuracy of the available network simulation models. As a part of the community battery experiments, the local low voltage power grid is modeled and measurement data is gathered.

We apply the state estimation procedure to the network, relying on a feeder model and real load and voltage measurements. The feeder contains 142 buses, of which 34 are regular household customers, one is the distribution transformer and one is the community battery. The other buses are network cable joints. The source of the network data is the Alliander GIS database, which contains the exact location and properties of the electricity cables. However, the GIS database does not contain on which phase each customer is connected, therefore the estimator is constructed using a balanced single phase model, using the formulation in Section 3.3. The distribution transformer is located at the top of the feeder, and the Neighborhood Battery is installed at the end of the feeder. Both the transformer and battery contain SCADA equipment for measuring power and voltage at a 1-second rate. Of the 34

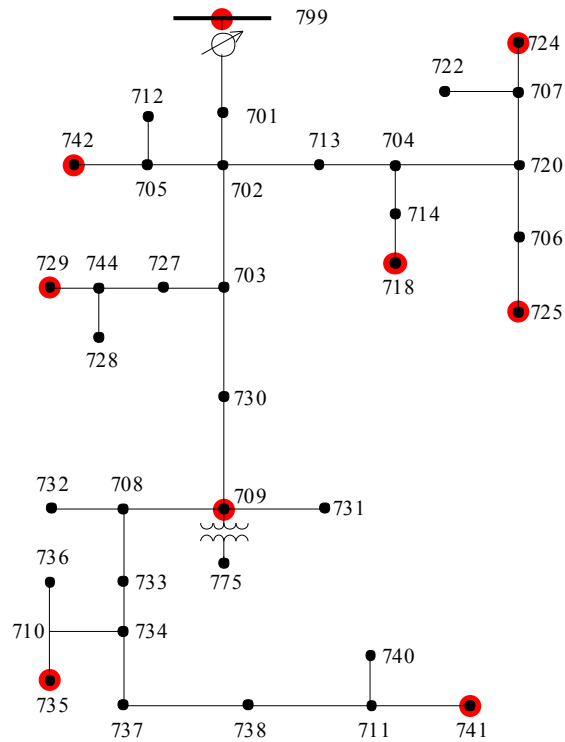


Figure 4.4: IEEE 37 node test feeder model, voltage sensors are indicated with red circles.




Figure 4.5: Example voltage profile with forecast and estimation update at all the buses, numbered as in Figure 4.4.

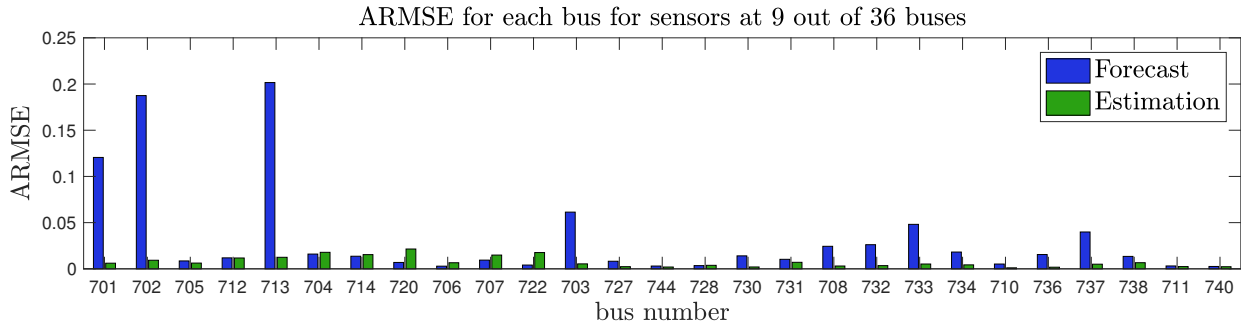


Figure 4.6: ARMSE in p.u. for each non-measured bus across all phases. Buses with higher forecast errors benefit significantly from the estimation procedure. Buses that already have good accuracy on forecasted values of the order < 0.01 p.u. do not necessarily gain much from estimation. This can be attributed to the fact that these errors are in the same order as the modeling errors due to the estimator’s linear approximation.

households connected to this feeder, 12 customers share their power consumption data with Alliander as part of the community battery project. All measurements have been collected with a 1-minute resolution. Customers with no direct measurement were assigned the residual power load, which was defined as the total transformer load minus the sum of all measured loads. Each unmeasured customer was assigned an equal proportional share of the residual load. Note that this introduces some error in the forecast procedure.

Figure 4.8 shows the results of applying SE and comparing the predicted and estimated voltage drop at a particular bus with real voltage measurements. Observe that the estimated values provide a significant improvement over the forecasted values, showing agreement with the actual values. The improvements are particularly strong for larger voltage deviations, providing critical information for safety procedures. At certain times the estimation does not improve accuracy, which has two explanations. Firstly, for smaller voltage deviations, modeling errors due to linearization of power flow are more dominant, as mentioned above. Secondly, the effect of limited real-time voltage sensors (in this case only 2 out of 140 buses) provides significant but limited improvement due to limited observability of all load flow scenarios in the network. This challenge requires revisiting the notion of network observability, which is discussed in Section 4.6. Similar to the IEEE synthetic experiment, SE significantly reduces the ARMSE across all buses in the network, on average by 60%. Given the difficulty of predicting the power consumption of individual households due to their variability, this result is useful for DSOs in improving the fidelity of their forecasting data with limited sensing capabilities, which is a likely context in most networks for the foreseeable future. As such, Alliander is implementing SE algorithms in their critical calculations, and aim to use the data for optimal sensor placement, cable health monitoring, real time overload predictions, and control of voltage and power flow.

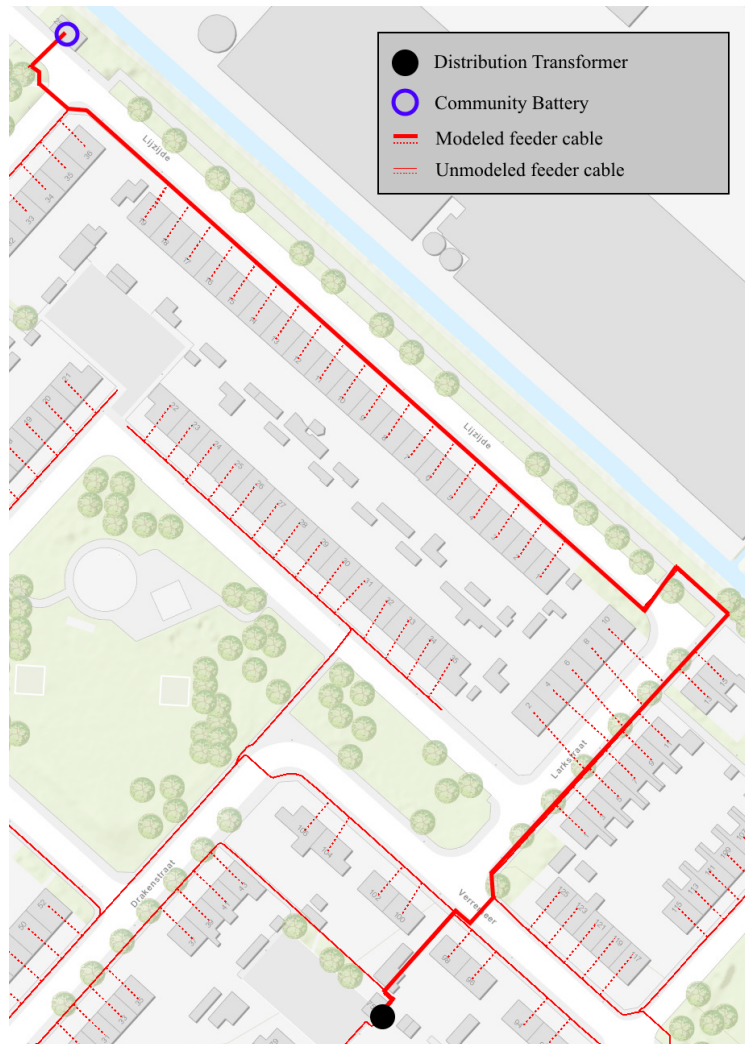


Figure 4.7: GIS view of Alliander’s low voltage network of Rijsenhout. The outlined modeled network is the network that is considered for the DGSE model. The modeled part of the network consists of 34 customers. The unmodeled cables are not physically connected to the modeled network. At the distribution transformer and the community battery both power and voltage are measured. At 12 households, the power was measured. For privacy reasons, their exact location could not be displayed, but they are almost uniformly distributed along the cable.




Figure 4.8: Comparison of predicted and estimated voltage with real voltage measurement at the Neighborhood Battery bus.

4.6 Suggestions for Observability Analysis and Sensor Placement

Network observability characterization for state estimation using the weighted least squares (WLS) approach (as in (4.2)) was derived by Monticelli and Wu for settings that assume the DC approximation [149]. Gómez Expósito and Abur proposed an approach for general nonlinear measurement equations, which involves taking the first order Taylor approximation and can be used to include current magnitude measurements [92].

In contrast to the conventional methods, our state estimator does not require solving a WLS problem. Instead of using the Taylor approximation for a fully nonlinear power flow model, we have expressed both our measurements $\Delta \mathbf{Y}_m$ and our estimation variables $\Delta \mathbf{Y}_e$ as a linear function of the load vector \mathbf{s} . Assuming we have access to a load forecast μ_s for all nodes in the network, we can argue that the voltage forecast $\mathbf{Z}_e \mu_s$ itself is well-defined and provides full *prior observability*; given statistical information for all loads, the mapping from load forecast to voltage forecast is well-posed.

As covered in Section 4.2, our estimator is a LLSE which is equivalent to *projecting* the estimation variable $\Delta \mathbf{Y}_e$ onto the set of linear functions of the measurement $\Delta \mathbf{Y}_m$, which can be interpreted as the best linear unbiased estimator, assuming the linear power flow model is unbiased. The projection is a result of the assumption that the estimation step considers a limited number of sensors $M < N$, which in the context of network theory means this step will never be able to capture all changes in the estimation variables. That said, it is possible to determine a sensor placement that allows the measurements $\Delta \mathbf{Y}_m$ to capture a maximum amount of information about the estimation variables $\Delta \mathbf{Y}_e$.

Definition 4.6.1. A load profile $\mathbf{s} \in \mathbb{R}^N$ is observable if $\mathbf{s} \in \mathcal{R}(\mathbf{Z}_m^\top)$ (row space of the measurement matrix). For any $\mathbf{s} = \mathbf{s}_o + \mathbf{s}_u$, with $\mathbf{Z}_m \mathbf{s}_u = 0$ and $\mathbf{s}_o \in \mathcal{R}(\mathbf{Z}_m^\top)$, such that

$\mathbf{Z}_m \mathbf{s} = \mathbf{Z}_m \mathbf{s}_o$, we say that $\mathbf{s}_o, \mathbf{s}_u$ are the observable and unobservable parts of the load profile.

Our aim now is to design an estimator that minimizes

$$\begin{aligned} \Delta \mathbf{Y}_e - L[\Delta \mathbf{Y}_e | \Delta \mathbf{Y}_m] &\approx \mathbf{Z}_e \mathbf{s} - (\mathbf{Z}_e \mu_s + \mathbf{Z}_e (\mathbf{s}_o^m - \mu_{s,o}^m)) \\ &= \mathbf{Z}_e (\mathbf{s} - \mathbf{s}_o^m) + \mathbf{Z}_e (\mu_{s,o}^m - \mu_s) \\ &= \mathbf{Z}_e (\mathbf{s}_u^m - \mu_{s,u}^m), \end{aligned} \quad (4.21)$$

with respect to some metric over all relevant load scenarios. Here, \mathbf{s}_o^m and \mathbf{s}_u^m are the observable and unobservable parts of the load profile \mathbf{s} , with respect to the mapping $\Sigma_s \mathbf{Z}_m^\top (\mathbf{Z}_m \Sigma_s \mathbf{Z}_m^\top)^{-1} \mathbf{Z}_m$. Notice that as Σ_s and $(\mathbf{Z}_m \Sigma_s \mathbf{Z}_m^\top)^{-1}$ are both full rank square matrices, the null space of the mapping is fully characterized by the null space $\mathcal{N}(\mathbf{Z}_m)$. A desired property is for the unobservable part of the load profile to be insignificant or, ideally, also unobservable with respect to estimation matrix \mathbf{Z}_e . This means that whatever information is lost by the projection by measurement matrix \mathbf{Z}_m does not contribute to changes in the *actual* values of the estimation variables $\Delta \mathbf{Y}_e$. In mathematical terms, this means that we want the null spaces of \mathbf{Z}_m and \mathbf{Z}_e to intersect as much as possible. This can be formulated as the following optimization problem:

$$\min_{\mathbf{Z}_e, \mathbf{Z}_m} \dim \mathcal{N} \left(\begin{bmatrix} \mathbf{Z}_e \\ \mathbf{Z}_m \end{bmatrix} \right) - \dim \mathcal{N}(\mathbf{Z}_m). \quad (4.22)$$

Note that $\mathbf{Z}_e, \mathbf{Z}_m$ are both determined by the sensor placement. Alternatively, given a data set Ξ of historical load profiles, we can formulate a data-driven sensor placement approach which minimizes

$$\min_{\mathbf{Z}_e, \mathbf{Z}_m} \sum_{\xi \in \Xi} \|\mathbf{Z}_e (I - \mathcal{P}_{\mathbf{Z}_m}) \xi\|^2, \quad (4.23)$$

where $\mathcal{P}_{\mathbf{Z}_m} = \mathbf{Z}_m^\top (\mathbf{Z}_m \mathbf{Z}_m^\top)^{-1} \mathbf{Z}_m$ is a projection matrix. Equation (4.23) should be read as trying to minimize the extent to which the parts of all historical load profiles that are unobservable with respect to \mathbf{Z}_m affect the value of \mathbf{Y}_e . This approach allows the DSO to prioritize important load flow scenarios that are more safety-critical, by weighting these differently, yielding

$$\min_{\mathbf{Z}_e, \mathbf{Z}_m} \|\mathbf{Z}_e (I - \mathcal{P}_{\mathbf{Z}_m}) \Xi W\|_F^2, \quad (4.24)$$

where W is a diagonal weight matrix. The above sensor placement strategies are here presented as suggestions. We implement and assess these in a separate paper.

4.7 Conclusions

This chapter addressed the challenge of formulating a distribution grid state estimator, for scenarios where fully observed sensor arrangements are not yet feasible, and load forecasts are subject to large uncertainties due to lack of access to data. We derived an algorithm that

exploits the information in load forecasting and feeder models to construct prior statistics of relevant voltage variables. We then used a Bayesian approach, in the form of the linear least squares estimator, to update prior voltage statistics in real-time based on measured deviations at a limited set of voltage sensors. We applied the method to a benchmark IEEE network and on a real testbed in the Netherlands and showed its ability to provide accurate voltages estimates using limited historical data and real-time sensors. We then discussed suggestions for assessing the methodology in terms of its implications for observability and sensor placement. As such, the method is highly applicable in the typical distribution network setting in which data and sensing will remain limited for the foreseeable future.

Chapter 5

Decentralized Optimal Power Flow

“But in either the technical or political sphere (or both) any significant move to decentralize would amount to retro-fitting our whole society, since centralized institutions have become the norm.”

- Langdon Winner

Submitted to IEEE Transactions on Smart Grid [57]. Earlier work presented at the IEEE Power & Energy Society General Meeting in Boston, USA, July 2016 [193].

The rapid integration of new technologies in electric networks, such as solar generation, electric vehicles and battery storage, is both challenging traditional paradigms and providing new opportunities to operate the grid. A large body of work has addressed the abilities of inverter-interfaced Distributed Energy Resources (DERs) to aid the regulation of voltage and power flow, either through providing decentralized control schemes based on local information or optimization-based approaches using system-wide information or distributed implementations. These methods suffer, either from suboptimal and unstable behavior and extensive tuning or from the need for an extensive communication infrastructure. In this paper, we consider a data-driven approach to determine control policies for a system with multiple DERs that collectively reconstruct the solution of an optimal power flow (OPF) problem by using solely local information. We frame the problem by integrating the learning problem with control-theoretic and information-theoretic tools. A partitioning of control variables guides the selection of variables that can be used safely and adequately for our learning problem, and a rate distortion framework facilitates the analysis of how well the resulting fully decentralized control policies are able to reconstruct the optimal solution. Our methodology also provides a natural extension to decide what nodes an agent should communicate with to improve the performance of its individual policy.

5.1 Introduction

Historically, low and medium voltage distribution networks facilitated uni-directional, predictable power flow from substations to end customers. By conservatively oversizing cables and control equipment, Distribution System Operators (DSOs) could deal with peak and contingency scenarios safely, while requiring limited sensing and control actions. Gradual changes in the network were typically dealt with by a *fit-and-forget* approach, making larger capital intensive network reinforcements that could satisfy the needs for decades to come.

This paradigm was first challenged by the rapid advent of Distributed Generation (DG), mostly through photovoltaic (PV) systems, leading to new risks of over-voltage, reverse flow, and thermal overloads, especially at times of high DG and low consumption [195]. The subsequent adoption of electric vehicles (EVs), battery storage and other Distributed Energy Resources (DERs), increasingly interfaced by controllable electronic power inverters, has shaped a new reality with more challenging quickly diversifying power flow and voltage dynamics requiring a new control paradigm.

Many Distribution System Operators (DSOs) are testing ways to prevent expensive network updates by exploiting DERs and their sensing and actuation capabilities to provide *distributed energy services* (DES) to enable *active distribution networks*. In DES, DERs are embraced (1) to compensate for the negative effects of DG and EVs thereby allowing higher levels of penetration, and (2) to distribute and diversify capital investments on the grid *both in space and time*, to align closer with actual changes that occur organically. The latter benefit can dramatically decrease the capital costs, as larger updates tend to be more conservative due to longer and more uncertain planning horizons [127].

The mounting interest in this problem has led to a plethora of methods that can broadly be categorized in two categories: decentralized control based on local information versus optimization-based control based on system-wide information. The decentralized methods have shown promise in their ability to reduce voltage variability, but require extensive tuning which is impractical and costly for larger networks with many inverters. In addition, these methods yield suboptimal control signals and cannot guarantee the satisfaction of critical system constraints. The optimization-based approaches yield globally optimal power injections and incorporate constraints on voltage and reactive power capacity. A critical assumption is the availability of a communication network to collect measurements from throughout the network and send resulting inverter control signals, which is far from practical for most present day distribution networks. Although distributed implementations try to address this issue, these still require extensive communication between most buses in a network.

The contribution of this work lies in combining the “best of both worlds” by integrating *optimal power flow* and *machine learning* as a means to *learn* a controller for each agent that *locally mimics or reconstructs* the actions from a system-wide OPF controller determined in simulation. As such, the method provides a data-driven approach to tuning controllers that is fully automatic, relying on historical load and generation data and a model of the network. This learning-based optimization approach was first proposed for reactive power control in [193], further formalized with information theory in [54], and recently applied to

design voltage-reactive power droop-curves in [20].

Relevant Advances in Active Distribution Networks

Earlier efforts to address voltage and power flow problems related to higher penetrations of DERs aimed at developing *decentralized control based on local information* [216, 37, 120, 191, 207, 117]. These use heuristics to adjust reactive power output at each inverter based on the local voltage. These methods have shown promise in their ability to reduce voltage variability, but suffer from *extensive tuning* which is impractical for larger networks with many inverters. In addition, these methods yield suboptimal control signals and cannot guarantee the satisfaction of critical system constraints. A step further, we find methods that consider a control theoretic formulation emulating proportional [131] or integral control [227] (the latter to address the steady-state errors inherent to proportional methods). Another body of work considers the use of sensitivities between controllable variables and relevant quantities in the network such as voltage levels or branch flows [16, 174, 172, 32]. Lastly, [227] proposes a simple localized optimization framework aiming to track a local voltage reference, and provides necessary conditions that show that localized control does not converge and may not be stable for longer and more heavily loaded networks.

The emerging need for active distribution grids also motivated the use of Optimal Power Flow (OPF) for Distribution Operation to enable *optimization-based control based on system-wide information* [76]. OPF refers to solving an optimization problem that minimizes some economic or operational objective subject to power flow and other relevant system constraints. Traditionally, OPF is used *offline* as a design tool for network upgrades to size and place equipment, as proposed for capacitor planning in [17, 15], or as a planning tool to schedule the dispatch of generators and control equipment. As an extension of economic dispatch, [80] meant to ensure that the dispatch of generators minimizes economic cost and results into a power flow scenario that is physically feasible and satisfying relevant operating and safety constraints. With the advent of real-time sensing, the implementation of OPF in an *online* setting has been a popular research area in recent years. Historically, the non-convex nature of the OPF problem has limited its online integration due to the lack of convergence guarantees [148], but theoretical advances in optimization led to a breakthrough in solving the OPF problem through convex relaxations [12, 128]. The inherent lack of communication infrastructure in distribution networks subsequently motivated various efforts to implement OPF in a *distributed* fashion, relying on agent-to-agent communication, such as via consensus algorithms [228] or dual decomposition [23, 46, 197, 48]. Recent work in extremum-seeking showed that model-free optimization is possible [7] with provable guarantees for convergence and convexity for a variety of distribution feeder objective functions over a broad range of power flows [8].

Despite the elegance of distributed solutions, the necessary communication infrastructure is still a steep investment at the scale of transforming distribution networks. It remains a challenge to understand what can be done with minimal investments in assets and com-

munication infrastructure by fully exploiting existing infrastructure and the capabilities of naturally arising DERs.

Approach and Contributions

We propose a machine learning-based method to decentralize the implementation of optimal power flow (OPF) to regulate voltage and power flow in distribution grids without the need for communication. The approach consists of four steps. First, for a specific network with a set of controllable DERs, we retrieve data points for loads and generators for K different scenarios, collected over an extended period of time, typically collected by advanced metering infrastructure (AMI). Second, for all scenarios, we run a centralized optimal power flow computation to understand how a group of inverters *could have* best minimized a collective objective by adjusting their real and/or reactive power injections, given a certain available capacity. Third, for each individual DER, we use regression to determine a function that relates its *local* measurements to its optimal power injections, as determined by OPF. Last, we validate these functions as controllers on simulations to determine the power injections based on a new local measurement.

Given the fact that an optimal control action as determined by an OPF solver is a function of states *throughout* the network, our problem is to perform a *reconstruction*, and its efficacy reduces to the following question; to what extent can local state measurements *inform* the reconstruction of a control action that is optimal for the entire network? We treat the reconstruction of each inverter’s control actions as a *compression problem* formulated with *rate distortion theory*, as first proposed in [54]. This framework provides a theoretical lower bound on the *minimum distortion*; the best reconstruction possible given local information. Our formulation naturally extends to consider which nodes in the network a DER should communicate with in order to improve its reconstruction (by further minimizing its distortion).

We show that our method yields close-to-optimal results for various OPF objectives, for both single-phase and unbalanced three-phase networks. As such, the method respects constraints on voltage, equipment specifications and reactive power capacity. The implementation of decentralized OPF is data-driven and needs little or no real-time communication, avoiding expensive investments in infrastructure, controller tuning and maintenance. Our method has the ability to operate in congruence with the existing *legacy control equipment*, such as load tap changers and capacitor banks, making the method relevant in contexts where operators or other autonomous systems make real-time control decisions. Analyzing the structure of the controllers in terms of what local measurements are selected by the machine learning algorithm and which nodes are optimal to communicate with in order to improve the reconstruction provides new insights in the complexity of operating distribution networks with optimal power flow methods.

5.2 A Control-Theoretic Lens for OPF

The dynamics considered in OPF for electric distribution systems with alternating current (AC) are typically not described with differential equations, as these are assumed to be in a *quasi steady-state*. This assumption rests on the condition that the distribution network is connected to the transmission grid through a substation, which can be modeled as an *infinite* bus, meaning that the transmission grid *admits* or *compensates* any changes in power across the distribution network to enable a *fixed* voltage phasor. With this assumption, voltage and power flow across a network can be determined through solving a set of algebraic equations describing fundamental first principles in electric power.

However, when we consider decentralized control solutions, the spatial coupling and the updating nature of local controllers requires a dynamical analysis. We adopt and extend a control-theoretic perspective to optimal power flow (OPF), as initially proposed in [219] for general optimization algorithms and interpreted for power systems in [101]. This lens will help us interpret the challenge of decentralizing OPF with machine learning and it naturally triggers relevant questions around the analysis, development and integration of our solution.

In control theory, the state-space representation a *dynamical system* can be thought of as described by a set of variables including input, state and output variables [81], related by a set of differential and/or algebraic equations. The input variables are also called *exogenous* as these typically enter from outside the system boundary, and can be split up in *controllable* input variables $u \in \mathcal{U}$ (in short, “inputs”) and *uncontrollable* input variables $d \in \mathcal{D}$ (in short, “disturbances”). The state variables $x \in \mathcal{X}$ (or “states”) are all the variables necessary to mathematically describe the system’s full dynamics. The states may contain certain inputs or disturbances. The remaining states, those that are not inputs or disturbances, but that are needed to describe the system dynamics are *endogenous* variables. Endogenous variables have values that are determined by other exogenous variables in the system, and are often called *dependent* variables, such as in econometrics [94]. The output variables $z \in \mathcal{Z}$ (or “outputs”) are all the variables that are measured in some *real-time* capacity, potentially at various sampling-rates.

Consider a network modeled by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with \mathcal{N} a set of nodes representing all buses with cardinality $\nu := |\mathcal{N}|$ and \mathcal{E} a set of edges representing all branches with cardinality $\eta := |\mathcal{E}|$. At each bus, we have 4 variables; voltage magnitude V_n , voltage angle δ_n , and nodal real and reactive power p_n and q_n , for all $n \in \mathcal{N}$. Vectorized for all buses, we have $V, \delta, p, q \in \mathbb{R}^\nu$ and we define the state of the system as

$$x := \begin{bmatrix} V \\ \delta \\ p \\ q \end{bmatrix} \in \mathbb{R}^{4\nu}. \quad (5.1)$$

As suggested above, we partition the state x into controllable inputs u , uncontrollable inputs or disturbances d and endogenous variables x^{end} . This partitioning is done per bus, based on the bus type, as suggested in [101] and summarized in Table 5.1.

	Exogenous		Endogenous
	Controllable - u	Uncontrollable - d	$x^{\text{end}} := x_{\setminus\{u,d\}}$
PQ generation	p_n, q_n		V_n, δ_n
PQ load		p_n, q_n	V_n, δ_n
PV generation	p_n, V_n		q_n, δ_n
slack bus	V_0	δ_0	p_0, q_0

Table 5.1: Partitioning of nodal variables for all bus types.

Note that the exogenous inputs are also treated as state variables: $\{u, d\} \subset x$. In our setting, the state x is the solution to either a set of algebraic equations or an optimization problem. In the former case, we have to solve the power flow equations $F(x) = 0$ to find x ,

$$x \in \{\xi \in \mathbb{R}^{4\nu} : F(\xi) = 0\} =: \mathcal{M}. \tag{5.2}$$

Here, \mathcal{M} is the *power flow manifold*, as defined in [25], representing power flow physics and operating constraints as discussed in Chapter 3.

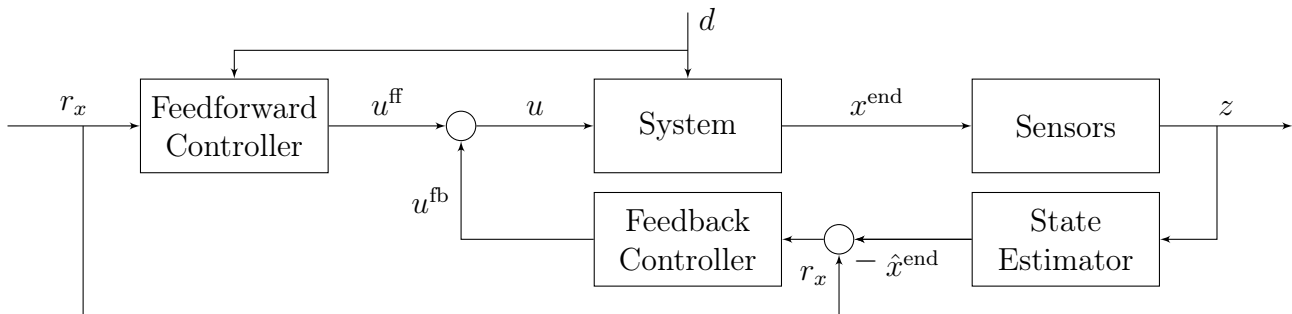


Figure 5.1: Dynamical system representation with feedback and feedforward control.

Figure 5.1 shows a block diagram of a general dynamical system, depicting with both feedback and feedforward control. In feedback control, the output z is used in real-time to inform the input u of the system, typically by comparing z to some defined reference signal or set point r . An example is the use of voltage measurements across a distribution network to automatically adjust the feeder voltage through controlling the load tap changer. In this context, the state vector x may not be fully measured in the output z , and a *state estimator* or *observer* is needed to estimate the full state \hat{x} from the output. Feedforward control is traditionally known as a simple and powerful technique to complement feedback control. It can be used both to improve the response to reference signals and to reduce the effect of measurable disturbances [10]. As such, there are broadly two forms of feedforward control. Firstly, a reference (either on output r_z or state r_x) can be used to determine an input u_r^{ff} that *anticipates* the system’s dynamical response, typically by feeding r_z through an *inverse*

model of the system. Secondly, information about the disturbances d , either through real-time measurement or a prediction, may be used to design an input u_d^{ff} that tries to compensate for its harmful effects by *attenuating* it directly or, again, *anticipating* the system's dynamical response.

By combining feedforward and feedback control in the right way, one has two degrees of freedoms to split the control design problem into two parts. The feedback controller can be designed to give robustness and effective disturbance attenuation, and the feedforward part can be designed independently to give the desired response to command signals or objectives [10].

In this work, we direct our focus on using machine learning to design a *decentralized feedforward controller* $\hat{u}_i := \hat{\pi}_i(d_i) \forall i \in \mathcal{C}$ that locally reconstructs the optimal actions of an OPF problem u_i^* from measurable disturbances d_i , as given in Table 5.1 for different bus types. For a general dynamical system, the method can be interpreted as *learning-based feedforward optimization*. For this paper, we do not consider using machine learning to learn a feedback controller; a policy $\hat{u}_i := \hat{\pi}_i(x_i^{\text{end}})$ that is a function of endogenous dynamic state variables, as proposed in [20]. The main challenge with using these variables is that the historical data of x_i^{end} will be different than the measured state $x_i^{\text{end}} = f(u_i, d_i)$, which is a function of the input. This requires understanding existence and stability of the fixed point solution $(x_i^{+, \text{end}}, \hat{u}_i^+)$ such that $x_i^{+, \text{end}} = f(u_i^+, d_i)$ and $\hat{u}_i^+ = \hat{\pi}_i(x_i^{+, \text{end}})$. If the historical data was affected by control actions, careful analysis is necessary to deal with potential issues related to closed-loop identification [209].

5.3 Learning Decentralized Feedforward Controllers

The goal of the machine learning procedure is to find a model for each DER $i \in \mathcal{C}$ that closely reconstructs the optimal power injection u_i^* based *solely* on local measurements x_i available at node i . We first construct a *central training set* consisting of power flow scenarios that are representative of future behavior on the grid. This data set can contain historical measurements taken from advanced metering infrastructure (AMI) in the grid, augmented with proxy data for variables that were not measured or for scenarios that did not occur in the past but are anticipated to occur in the future. For example, a load not metered by AMI, may be approximated with an average load profile. Examples of anticipated future behaviors not captured in historical measurements are connections of new electric vehicles or solar installations. These can be simulated and added to the training data set used for machine learning. Together, the central training set contains T power flow scenarios augmented with the optimal control set points $u_i^*[t] \forall i \in \mathcal{C}, t = 1, \dots, T$, which were computed in a centralized OPF problem as defined in Sections 3.4-3.7.

Regression is now performed for each individual inverter, by selecting a *local data set* from the central data set, with only those variables x_i that are measured at node i . The variables contained in x_i may differ from node to node based on the local sensing infrastructure. These can be complemented by other predictive variables that are available locally such as time or

weather information. For illustration, we use a running example of a node that can measure only the local load p_i^c , the local PV generation p_i^g , and the present capacity \bar{s}_i . This yields the local training set with base variables X_i and labels Y_i ,

$$[d_i[1] \cdots d_i[T]] , [u_i^*[1] \cdots u_i^*[T]] . \quad (5.3)$$

The next step is to select a machine learning model that is expressive enough to reconstruct valid patterns in the optimal control actions that can be found in the local data, and that is simple enough to generalize well to new scenarios. We compared various model structures in our earlier work [193] and selected a multiple stepwise linear regression algorithm that selects a subset of available nonlinear features [86]. This model is linear in the parameters, but allows for nonlinear transformations of the base variables in $\varphi(d_i)$, concretely quadratic and interaction terms. For instance, if d_i contains 3 base variables as in our example, $\varphi(d_i)$ would yield 3 linear features; $\varphi_1^{(i)} := p_i^c$, $\varphi_2^{(i)} := p_i^g$ and $\varphi_3^{(i)} := \bar{s}_i$, and 6 more nonlinear features, of which 3 interaction terms ($\varphi_4^{(i)} := \varphi_1^{(i)}\varphi_2^{(i)}$ etc.) and 3 quadratic action terms ($\varphi_7^{(i)} := (\varphi_1^{(i)})^2$ etc.). In general, for B_i base variables, we construct $F_i = 2B_i + \binom{B_i}{2}$ features as inputs to the regression problem. The F_i input variables of the t^{th} sample are denoted as $\boldsymbol{\phi}^{(i)}[t] \in \mathbb{R}^{F_i}$. All F_i features across all T data points can be organized as

$$\Phi^{(i)} = \begin{bmatrix} \boldsymbol{\phi}^{(i)\top}[1] \\ \vdots \\ \boldsymbol{\phi}^{(i)\top}[T] \end{bmatrix} = \begin{bmatrix} \varphi_1^{(i)}[1] & \cdots & \varphi_{F_i}^{(i)}[1] \\ \vdots & \ddots & \vdots \\ \varphi_1^{(i)}[T] & \cdots & \varphi_{F_i}^{(i)}[T] \end{bmatrix} \quad (5.4)$$

We use a multiple linear model to relate output Y_i to input matrix $\Phi^{(i)}$,

$$\hat{u}_i(d_i) := \hat{\pi}_i(\beta^{(i)}, \boldsymbol{\phi}^{(i)}[t]) = \beta_0^{(i)} + \beta_1^{(i)}\varphi_1^{(i)}[t] + \dots + \beta_{F_i}^{(i)}\varphi_{F_i}^{(i)}[t]. \quad (5.5)$$

A least squares approach determines the coefficients $\beta^{(i)} = [\beta_0^{(i)}, \dots, \beta_{F_i}^{(i)}]$ that minimize the residuals sum of squares (5.6) given T samples,

$$\text{RSS}(\beta^{(i)}) = \sum_{t=1}^T (u_i^*[t] - \hat{\pi}(\beta^{(i)}, \boldsymbol{\phi}^{(i)}[t]))^2 \quad (5.6)$$

The algorithm is initialized with a multiple linear model of the basic variables in X_i only. At each iteration, the new feature that improves the Bayesian Information Criterion (BIC) [180] the most and sufficiently is added to the model. Subsequently, the variable with the lowest, and sufficiently small contribution is removed. These two steps are iterated until no variables meet the entrance or exit threshold of the algorithm. The goal of the stepwise selection algorithm is to select the subset of features that most accurately predicts the optimal power injection of a DER.

5.4 Reconstruction and Communication

Rate Distortion Framework

As proposed in [54], we approach the problem of how well the decentralized policies $\hat{\pi}_i$ can perform in theory from the perspective of *rate distortion*. Rate distortion theory is a sub-field of information theory, which provides a framework for understanding and computing the minimal *distortion* incurred by any given *compression* scheme. For a detailed overview, see [43, Chapter 10]. In a rate distortion context, we can interpret the fact that the output of each individual policy $\hat{\pi}_i$ depends only on the local state x_i as a compression of the full state x . Note that, in contrast to (5.1), here we consider u_i separately from state x_i , which includes disturbances d_i . We formulate the following variant of the classical rate distortion problem

$$\begin{aligned}
 D^* &= \min_{p(\hat{u}|u^*)} \mathbb{E}[\mathbf{d}(\hat{u}, u^*)] , & (5.7) \\
 \text{s.t.} \quad & I(\hat{u}_i; u_j^*) \leq I(x_i; u_j^*) \triangleq \gamma_{ij} , \\
 & I(\hat{u}_i; \hat{u}_j) \leq I(x_i; x_j) \triangleq \delta_{ij}, \forall i, j \in \mathcal{C} ,
 \end{aligned}$$

where $I(\cdot, \cdot)$ denotes mutual information and $\mathbf{d}(\cdot, \cdot)$ an arbitrary non-negative distortion measure. As usual, the minimum distortion between random variable u^* and its reconstruction \hat{u} may be found by minimizing over conditional distributions $p(\hat{u}|u^*)$.

The novelty in (6.2) lies in the structure of the constraints. Typically, D^* is written as a function $D(R)$, where R is the maximum *rate* or mutual information $I(\hat{u}_i; u_i^*)$. From Figure 6.1b however, we know that pairs of reconstructed and optimal actions cannot share more information than is contained in the intermediate nodes in the graphical model; \hat{u}_{c_1} and $u_{c_1}^*$ cannot share more information than x_{c_1} and $u_{c_1}^*$. This is a simple consequence of the data processing inequality [43, Thm. 2.8.1]. Similarly, the reconstructed optimal actions at two different nodes cannot be more closely related than the measurements x_i 's from which they are computed. The resulting constraints are fixed by the joint distribution of the state x and the optimal actions u^* . That is, they are fully determined by the structure of the optimization problem that we wish to solve.

We emphasize that we have made virtually no assumptions about the distortion function. For the remainder of this paper, we will measure distortion as the mean square error (MSE) deviation between \hat{u}_i and u_i^* , which is a common *loss function* in supervised learning. However, we could also define it to be the suboptimality gap $f_o(x, \hat{u}) - f_o(x, u^*)$, which may be much more complicated to compute. This definition could also allow us to reason explicitly about the cost of decentralization, answering how much distortion is minimally incurred. Alternatively, we could assign more weight to certain scenarios to measure the cost of distortion in specific parts of the operating regime, for instance for scenarios where safety criteria are more stringent such as for contingencies or constraint violations.

Allowing Restricted Communication

Suppose that a decentralized policy $\hat{\pi}_i$ suffers from insufficient mutual information between its local measurement x_i and the optimal action u_i^* . In this case, we would like to quantify the potential benefits of communicating with other nodes $j \neq i$ in order to reduce the distortion limit D^* from (6.2) and improve its ability to reconstruct u_i^* . Here, we build the information-theoretic solution to the problem of how to choose optimally which other data to observe, as proposed in [54]. We assume that in addition to observing its own local state x_i , each $\hat{\pi}_i$ is allowed to depend on at most k other $x_{j \neq i}$.

Lemma 5.4.1. (*Restricted Communication [54]*)

If \mathcal{S}_i is the set of k nodes $j \neq i \in \mathcal{N}$ which \hat{u}_i is allowed to observe in addition to x_i , then setting

$$\mathcal{S}_i = \arg \max_{\mathcal{S}} I(u_i^*; x_i, \{x_j : j \in \mathcal{S}\}) : |\mathcal{S}| = k, \quad (5.8)$$

minimizes the best-case expectation of any distortion measure. That is, this choice of \mathcal{S}_i yields the smallest lower bound D^ from (6.2) of any possible choice of \mathcal{S} .*

Lemma 6.4.1 provides a means of choosing a subset of the state $\{x_j : j \neq i\}$ to communicate to each decentralized policy $\hat{\pi}_i$ that minimizes the corresponding best expected distortion D^* . Practically speaking, this result may be interpreted as formalizing the following intuition: “the best thing to do is to transmit the most information.” In this case, “transmitting the most information” corresponds to allowing $\hat{\pi}_i$ to observe the set \mathcal{S} of nodes $\{x_j : j \neq i\}$ which contains the most information about u_i^* . Likewise, by “best” we mean that \mathcal{S}_i minimizes the best-case expected distortion D^* , for any distortion metric d . Without making some assumption about the structure of the distribution of x and u^* , we cannot guarantee that any particular regressor $\hat{\pi}_i$ will attain D^* . Nevertheless, in a practical situation where sufficient data $\{x[t], u^*[t]\}_{t=1}^T$ is available, we can solve (6.8) by estimating mutual information [111].

5.5 Results for Single-Phase Optimal Power Flow

We evaluate the proposed method on a realistic testbed that is constructed from two independent sources: we construct a 129 node feeder model based on a real distribution feeder from Arizona, Figure 5.2, and populate this with demand measurements [50]. Pecan Street power consumption and PV generation data with a resolution of 15 minutes is obtained from individual residences in Austin, Texas for a period of 330 days. Individual household load and PV time series are selected randomly from the Pecan Street data set [50] and aggregated to match the spot load for each bus as specified in the experiment. We demonstrate Decentralized OPF through two experiments, both adopting the OPF problem defined in Section 3.4. The first experiment assumes controllable DER and loads spread randomly across a network. In this experiment, only reactive power is controlled with the objective to minimize voltage variations throughout the network. The second experiment considers a group of controllable

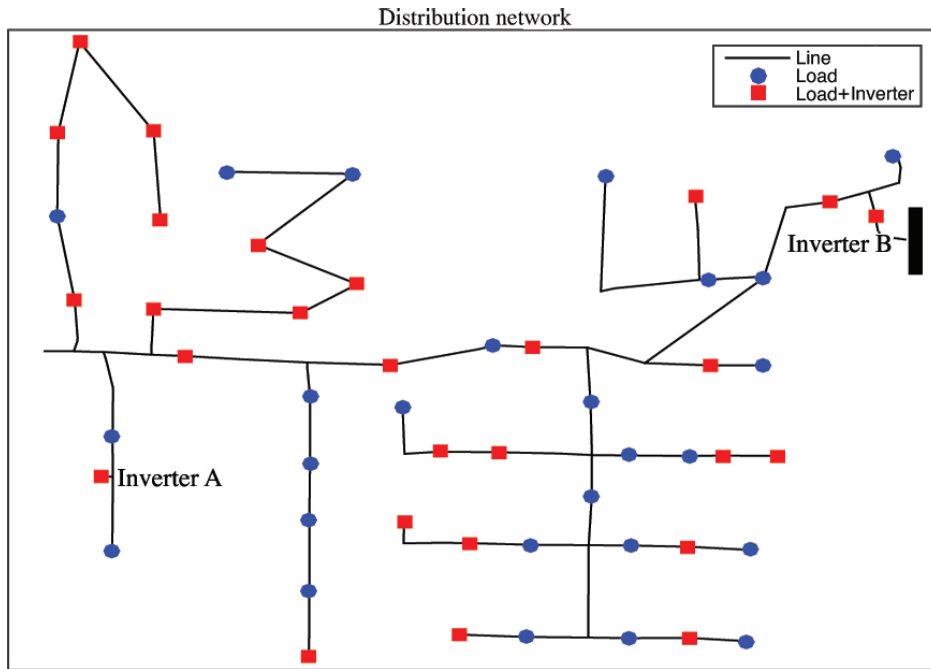


Figure 5.2: Distribution feeder for scenario 1. Substation is located on the far left, locations of loads with PV inverters (squares), and without PV inverters (circles) are included.

DER concentrated in one area, and a group of large loads concentrated in another area. In this experiment, both real and reactive power are controlled and the objective is to produce power locally as much as possible.

Case 1: Minimize Voltage Variations

For scenario 1, 50% of the 53 nodes with loads are randomly selected and equipped with PV installations with peak generation 80% of the peak real power load. For inverter B, Figure 5.3 shows the real power consumption and PV generation profiles for July 4th 2014. 2500 instances of OPF are solved with sampled load and PV generation data to retrieve the optimal reactive power output of all inverters $\{u^*[t]\}_{t=1}^T$, with the objective set with $\alpha = 1, \beta = 2 \cdot 10^{-4}, \gamma = 0$. The data is separated into training and validation data, and the obtained decentralized OPF controllers are simulated on test data that the learning step did not see. For illustration, here we show the results for July 4th 2014. The proposed control approach is simulated and compared to two other approaches: a situation where inverter reactive power capacities are not utilized, and the approach described in [191] where inverters are operated at a constant, non-unity, power factor. In our context, we chose to tune the inverters to operate at lagging (generating) power factor of 0.9 to reduce losses. In addition to the comparison, we extend our method to show it is capable of collaboration with a load tap changer (LTC). We design a scheme in which the inverters operate with controllers to flatten the voltage throughout

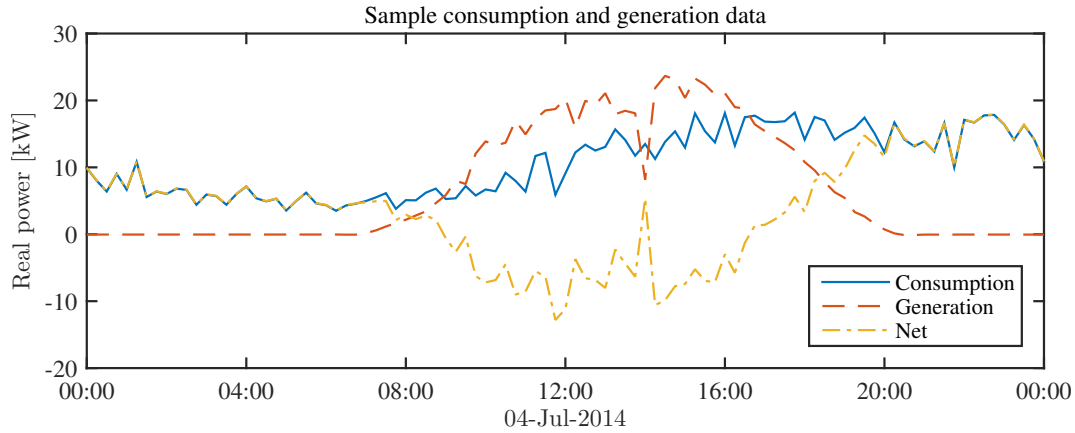


Figure 5.3: Sample real power consumption and PV data at the node of Inverter B on July 4th 2014 from Pecan Street data. The profile is obtained by aggregating data from six individual residences.

the feeder, and then adjust the turn ratio of the LTC at the substation, to safely lower the voltages throughout the feeder. Hence, we consider four approaches:

- a:** no reactive power support
- b:** constant power factor inverter operation
- c:** decentralized OPF (reactive power)
- d:** collaboration decentralized OPF and substation LTC

Figure 5.4 presents voltages at all nodes in the network for all four approaches, indicated with colored surfaces. For approach **a**, the voltage drop in the system is smallest between 10:00–16:00, when most real power demand is supplied by PV systems, see Figure 5.3. During these hours, the combination of real power injection of PV systems and reactive power generation of approach **b** prompts a voltage rise in the distribution feeder. Approach **c** achieves system voltages that are close to the nominal value of 1 p.u., and simultaneously reduces losses as implied by a lower objective function value, depicted Figure 5.5. The transition from peak PV generation to peak consumption between 12.00 and 20.00 causes the system voltage to change significantly without control. For approaches **a** and **b**, the lower bound of the ANSI standard is violated in the evening if traditional voltage regulators are not operated. The effect of approach **c** is obvious at these times: reactive power generated by inverters reduce voltage drop in the system and reduces losses, as depicted in Figure 5.4 and 5.5. Approach **d** exploits the reduced voltage variability achieved with approach **c**, and allows the substation to lower the overall voltage, for instance through adjusting a load tap changer. Finally, the voltages in approach **a** and **b** show a significant change at 14:00. This is caused by a temporary and sudden reduction of PV generation of approximately 60% at

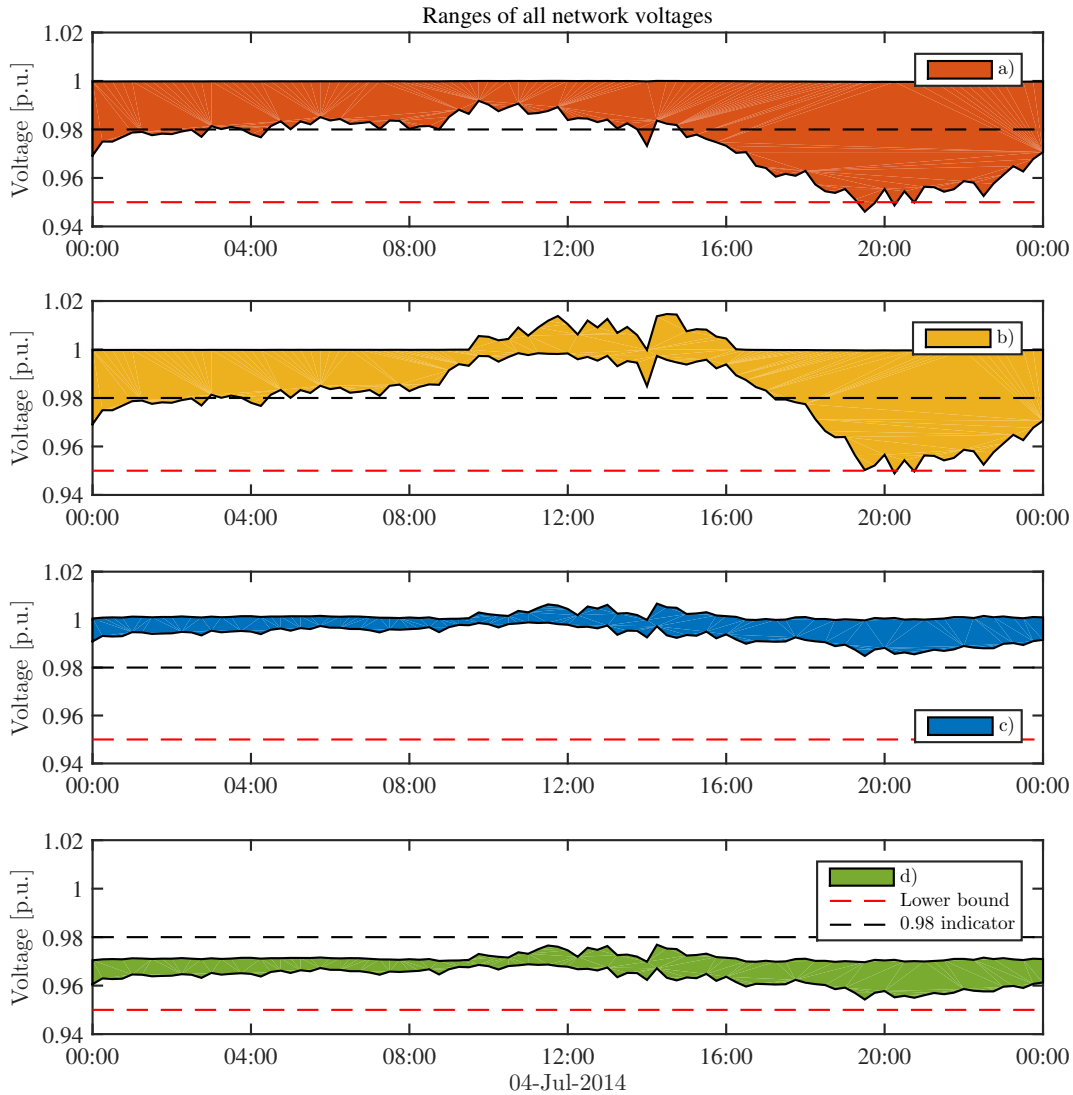


Figure 5.4: All network voltages for approaches **a**, **b**, **c**, and **d**. Colored planes represent the range between the maximum and minimum voltages at any node in the network. Lower voltage bound is indicated with red dashed line, an additional indicator is included as a black dashed line at 0.98 p.u..

14:00, that is best seen in the generation of Figure 5.3. The decentralized OPF controllers **c** act appropriately to this sudden change, damping it significantly.

Figure 5.5 compares the objective function values (3.21) for approaches **a**, **b**, and **c**. Compared to the situation of approach **a**, both approach **b** and **c** have beneficial effect on the objective function. However, approach **c** achieves the best performance at all times. Approach **b** generates reactive power proportional to the real power output of a PV system. This is reflected in a lack of control during hours when PV power is not available. Between

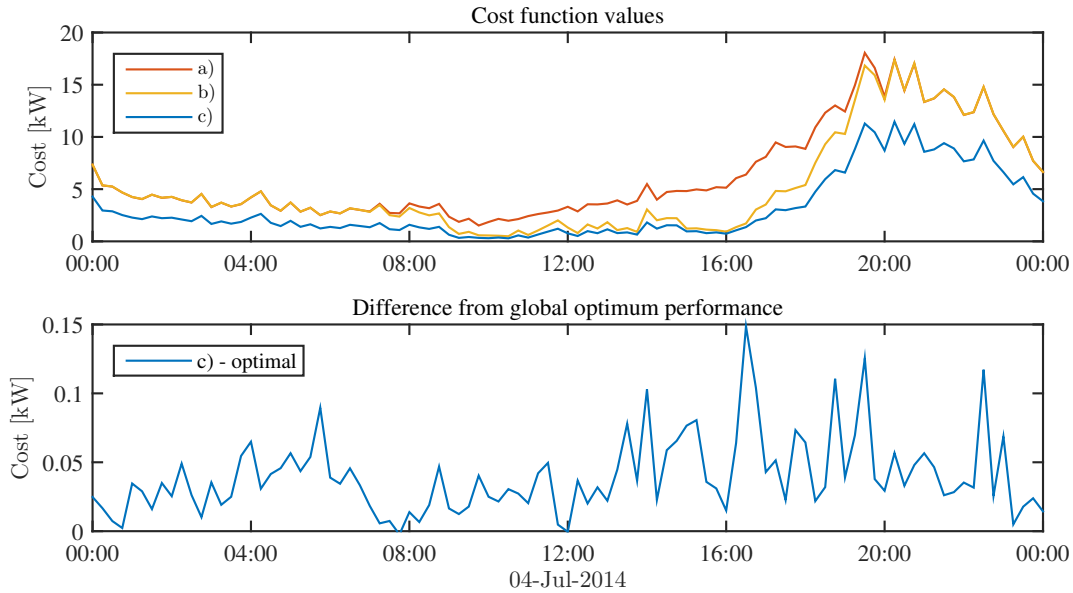


Figure 5.5: Upper figure - Objective function values of **a**, **b**. Lower figure - Difference in objective value for test date between Decentralized OPF control **c** and centralized OPF. The maximum difference is two orders of magnitude smaller than the optimal objective function, amounting to a suboptimality gap of 1.6%.

12.00 and 20.00 the objective function value of approach **a** increases rapidly, which is caused by the transition from peak real power generation to peak real consumption, as can be seen in Figure 5.3. The objective function value of approach **c** also increases, but significantly less.

The lower figure of Figure 5.5 shows the difference in cost between the centralized OPF problem and the Decentralized OPF problem implemented on test data. The maximum difference is 1.6% of the optimal function value, whereas the average difference is 0.15% of the objective function value. Hence, approach **c** achieves near-optimal performance.

Case 2: Localize Power Generation

Scenario 2 aims to show that the method also works for different power dynamics and objectives. We now assign 20 nodes in the lower subradial with controllable DERs, see Figure 5.6. In addition, the spot loads of 12 nodes in the left upper subradial are increased by a factor 3. By setting parameter γ in the objective function sufficiently high, this OPF problem tries minimize power procurement from the transmission grid by matching supply and demand locally, letting power flow from the area with controllable DER to the area with concentrated loads. The total real power capacity of the controllable DERs is assumed to be constant; future work will introduce time-varying capacity profiles based on charging, generation and consumption patterns. We assign 4 kWatt capacity to 20 inverters, totalling to 80 kWatt generation/consumption capacity across the feeder, as indicated by the dotted




Figure 5.6: Total net load delivered by the substation, in the case of: no control, OPF and Decentralized OPF.

red line in Figure 5.6. Figure 5.6 shows how the decentralized OPF controllers elegantly mimic the centralized OPF solutions across the test data set. Notice how the power delivered by the substation is 0 for all times where the network's net load is smaller than the combined power capacity of 80 kWatt.

5.6 Results for Three-Phase Optimal Power Flow

We also show the efficacy of Decentralized OPF for three-phase unbalanced networks, by learning and reconstructing the solutions $u_n^{p,\phi}, u_n^{q,\phi}$, with $\phi \in \{a, b, c\}$, of the centralized OPF problem for balancing voltage magnitudes across 3 phases, as formulated in (3.65), Section 3.7. The feeder model used here is the IEEE 13 node test feeder [107]. Figure 5.7 shows the results of applying Decentralized OPF to balance voltage magnitudes across phases for a scenario with a significant gap between phase b and phase c of ≈ 0.1 p.u.. We see that the voltage magnitudes are tightly balanced across the network. Notice that in order to balance the large gap in the lower half of the feeder (nodes 671, 692, 675, 680, 684), a small increase in gap is incurred in the left upper radial (nodes 645, 646).

5.7 Discussion

Apart from the general performance of Decentralized OPF, we use the scenarios for answering specific questions:

- For scenario 1, how does the structure of different DER $i, j \in \mathcal{C}$ differ across the network, as reflected by parameters $\beta^{(i)}, \beta^{(j)}$ in (5.5), and what does this tell us about the network and OPF problem?

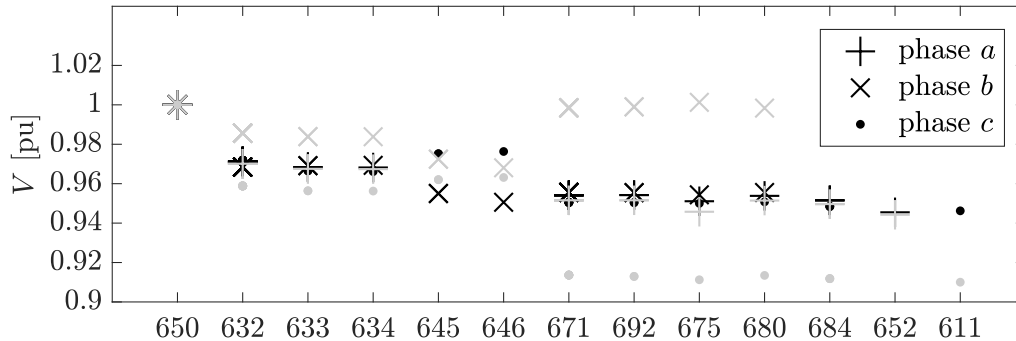


Figure 5.7: Voltage magnitude balancing. In grey, voltage magnitudes across three phases without control. In black, voltage magnitudes across three phases applying decentralized OPF.

- Between scenario 1 and 2, how does the optimal communication strategy differ, as determined via Lemma 6.4.1 and (6.8)?

Interpreting the Machine Learning Models

In our case study, we determined regression models for 27 different inverters. Table 5.2 shows regression results for inverters A and B (both indicated in Figure 5.2). The first two columns present the features selected by the stepwise regression and the values for the β -coefficients in (5.5). The third column shows the standard error of the estimate, and the fourth lists p-values, which here means the probability that coefficient is zero. A p-value of 0.1 implies a 10% chance that the corresponding coefficient is zero. Note how the stepwise regression approach results in two clearly different models. The reactive power output of inverter A depends predominantly on the local reactive power consumption $\varphi_2^{(A)}$, while the output of inverter B is strongly related to the reactive power capacity $\varphi_3^{(B)}$ and $\varphi_2^{(B)}$ is less relevant. Inverter A’s structure can be explained by the fact it is located at the end of the feeder. As a result, the aggregate power flow at A is relatively low and correlated with A’s own consumption. As the objective is to minimize losses and voltage deviations, A tries to produce a signal that looks like providing the neighboring demand for reactive power, thereby minimizing current locally *and* upstream. Inverter B’s structure can be explained by its location close to the feeder head. In this area, the aggregate power flow is much larger, and the output of inverter B is needed more as a “bulk” product to lower the flow on branches elsewhere in the network. This causes inverter B to operate at its maximum capacity most of the time, causing correlation with the output. As the maximum capacity varies with the local real power generation at B (a constraint formulated in (3.24)), the learning algorithm tries to “listen closely” to the present maximum capacity.

This example illustrates that optimal reactive power output of two inverters can have a completely different structure. Therefore, effective design of controllers based on local

Table 5.2: Normalized regression coefficients for inverter A and B, as depicted in Figure 5.2.

A	Est.	SE	p-value	B	Est.	SE	p-value
offset	0.02	0.01	0.01	offset	0.02	0.00	$8.2e^{-4}$
φ_1	0.37	0.01	0	φ_1	0.04	0.00	$3.2e^{-22}$
φ_2	0.77	0.01	0	φ_3	0.96	0.01	0
φ_3	-0.21	0.02	$8.0e^{-24}$	$\varphi_1\varphi_3$	0.06	0.01	$1.2e^{-12}$
$\varphi_1\varphi_2$	0.17	0.01	0	φ_1^2	-0.03	0.00	$3.9e^{-14}$
φ_1^2	-0.06	0.01	$1.4e^{-12}$	φ_3^2	-0.03	0.01	$6.1e^{-9}$

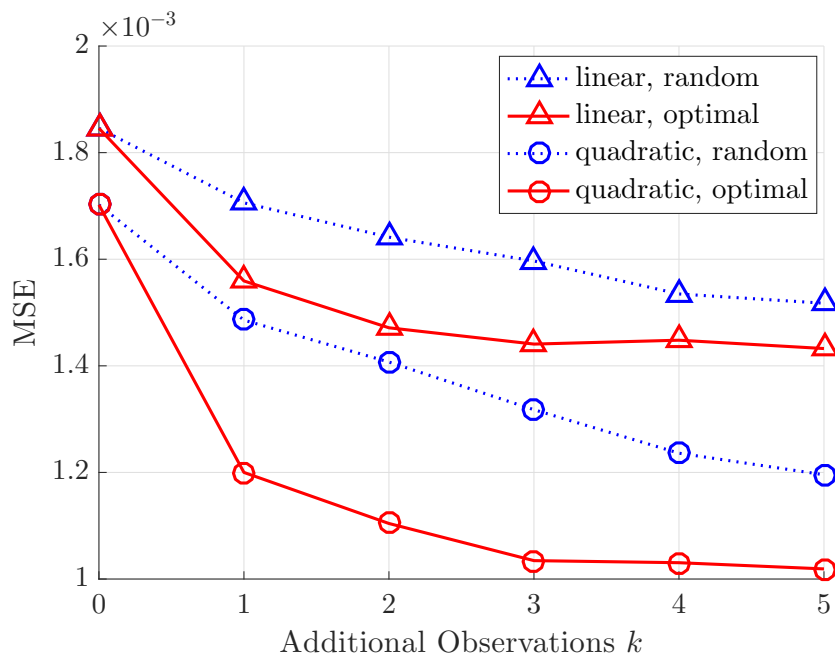


Figure 5.8: Comparison of OPF communication strategies [54].

measurements is challenging, and can clearly benefit from a data driven approach.

The Effect of Communication

In [54], we applied Lemma 6.4.1 to the OPF problem for a smaller network [107], in order to determine the optimal communication strategy to minimize a squared error distortion measure. Fig. 6.4b shows the mean squared error (MSE) distortion measure for an increasing number of observed nodes k and shows how the optimal strategy outperforms an average over random strategies. For the stepwise regression model with interaction and quadratic terms, as formulated in (5.5), the MSE of reconstructing the optimal action u_i for one agent decreases asymptotically to $\sim 40\%$, with a 29% decrease for adding one additional observation. We compare the quadratic model with a linear model in which only the base variables (in

our setting d_i) are used as features. We see that the linear model class has an asymptotic distortion that is $\sim 40\%$ higher than the quadratic model class. A next step could be to further enrich the machine learning model, for instance by using time and weather features and using a more expressive function approximator (such as deep neural networks).

In this paper, we also studied the emerging communication topology that arises from maximizing the mutual information between the optimal actions and the combined available measured variables. For case 1, where we minimize voltage deviations and losses across the network with DERs spread randomly across all buses, we see no particular patterns in the optimal communication infrastructure. Some DERs closer to the end of the network tend to communicate with loads nearby, which is explained by the emerging effect in OPF supply tends to match demand nearby to minimize losses higher upstream in the radial network. For case 2, where we are localizing the production of power by minimizing the power delivered from the substation, and interesting pattern arises from the fact that the loads and DER are concentrated in different areas. As a result, if we allow DERs to observe one extra node anywhere in the network, all select a node in the area with concentrated loads. This makes intuitive sense as the optimal control actions are largely responding to the high demand in the other area. In addition, across all 20 DERs the selected nodes are all in a set of only 3 nodes, suggesting an efficient implementation for a sensor network.

Designing Across Planning and Operation

An important practical consideration of implementing control schemes for active distribution networks is the impact on the planning and operation processes. Recent work has considered the co-optimization of planning and operation by considering both traditional expansion measures (such upgrading transformers or cables) as well as real-time control through DES, utilizing a decision-making process that builds on an iterative AC optimal power flow method [116, 117].

When using supervised learning, some immediate concerns arise around the quality of the training set and the historical data used to construct it. Firstly, a designer wants to have certainty that the historical data reflect the “normal” system behavior expected in future operation. If certain scenarios are not measured historically, the controllers may not *learn* how to respond in an optimal or desired way. If future scenarios are not measured, but can be anticipated, it may be possible to instead *simulate* these and augment these to the historical data. Secondly, a designer might want to know what happens when something changes in the system, such as a new street block connecting to a feeder or the installation of new electric vehicles or other DERs. This concern is valid for any controller. In our future work, we are assessing different learning-based controllers for OPF on their performance across a spectrum of system changes. Again, if some of these changes are anticipated by a planning process, it may be possible to simulate these changes. By adding these to the training set, controllers may learn to *anticipate* these. Otherwise, for smaller changes, a periodic retraining procedure can be used to update the control parameters incrementally by adding new historical measurements.

5.8 Conclusions

In this chapter, we present an integrative approach for decentralizing optimal power flow based on machine learning. A control-theoretic perspective helps formulate the problem as an instance of feedforward control based on measurable disturbances. A rate distortion framework allows us to interpret the decentralized learning approach as a compression and reconstruction problem, providing a theoretical lower bound on the distortion that can be achieved with local information and a procedure to improve controllers by communicating with a node that maximizes the mutual information between the optimal control and the available variables. Experiments on both single- and three-phase unbalanced systems illustrate the relevancy of Decentralized OPF for a broad range of systems and objectives. We discuss some insightful by-products of the approach to understand complex power networks and the behavior of inverters in OPF implementations, and point at some open problems toward practical integration in the planning and operation of distribution grids.

Chapter 6

Fully Decentralized Policies for Learning in Multi-Agent Systems: An Information-Theoretic Approach

“Thus we may have knowledge of the past but cannot control it; we may control the future but have no knowledge of it.”

- Claude E. Shannon

Presented at the 31st Conference on Neural Information Processing Systems in Long Beach, USA, December 2017 [54].

Learning cooperative policies for multi-agent systems is often challenged by partial observability and a lack of coordination. In some settings, the structure of a problem allows a distributed solution with limited communication. Here, we consider a scenario where no communication is available, and instead we learn local policies for all agents that collectively mimic the solution to a centralized multi-agent static optimization problem. Our main contribution is an information theoretic framework based on rate distortion theory which facilitates analysis of how well the resulting fully decentralized policies are able to reconstruct the optimal solution. Moreover, this framework provides a natural extension that addresses which nodes an agent should communicate with to improve the performance of its individual policy.

6.1 Introduction

Finding optimal decentralized policies for multiple agents is often a hard problem hampered by partial observability and a lack of coordination between agents. The distributed multi-agent

problem has been approached from a variety of angles, including distributed optimization [27], game theory [11] and decentralized or networked partially observable Markov decision processes (POMDPs) [157, 91, 151]. In this chapter, we analyze a different approach consisting of a simple learning scheme to design fully decentralized policies for all agents that collectively mimic the solution to a common optimization problem, while having no access to a global reward signal and either no or restricted access to other agents' local state. This algorithm is a generalization of that proposed in our prior work [193] related to decentralized optimal power flow (OPF). Indeed, the success of regression-based decentralization in the OPF domain motivated us to understand when and how well the method works in a more general decentralized optimal control setting.

The key contribution of this work is to view decentralization as a *compression* problem, and then apply classical results from information theory to analyze performance limits. More specifically, we treat the i^{th} agent's optimal action in the centralized problem as a random variable u_i^* , and model its conditional dependence on the global state variables $x = (x_1, \dots, x_n)$, i.e. $p(u_i^*|x)$, which we assume to be stationary in time. We now restrict each agent i to observe only the i^{th} state variable x_i . Rather than solving this decentralized problem directly, we train each agent to replicate what it would have done with full information in the centralized case. That is, the vector of state variables x is *compressed*, and the i^{th} agent must decompress x_i to compute some estimate $\hat{u}_i \approx u_i^*$. In our approach, each agent learns a parameterized Markov control policy $\hat{u}_i = \hat{\pi}_i(x_i)$ via regression. The $\hat{\pi}_i$ are learned from a data set containing local states x_i taken from historical measurements of system state x and corresponding optimal actions u_i^* computed by solving an offline centralized optimization problem for each x .

In this context, we analyze the fundamental limits of compression. In particular, we are interested in unraveling the relationship between the dependence structure of u_i^* and x and the corresponding ability of an agent with partial information to approximate the optimal solution, i.e. the difference – or *distortion* – between decentralized action $\hat{u}_i = \hat{\pi}_i(x_i)$ and u_i^* . This type of relationship is well studied within the information theory literature as an instance of *rate distortion theory* [43, Chapter 13]. Classical results in this field provide a means of finding a lower bound on the expected distortion as a function of the mutual information – or *rate of communication* – between u_i^* and x_i . This lower bound is valid for each specified distortion metric, and for *any* arbitrary strategy of computing \hat{u}_i from available data x_i . Moreover, we are able to leverage a similar result to provide a conceptually simple algorithm for choosing a communication structure – letting the regressor $\hat{\pi}_i$ depend on some other local states $x_{j \neq i}$ – in such a way that the lower bound on expected distortion is minimized. As such, our method generalizes [193] and provides a novel approach for the design and analysis of regression-based decentralized optimal policies for general multi-agent systems. We demonstrate these results on synthetic examples, and on a real example drawn from solving OPF in electrical distribution grids.

6.2 Related Work

Decentralized control has long been studied within the system theory literature, e.g. [138, 186]. Recently, various decomposition based techniques have been proposed for distributed optimization based on primal or dual decomposition methods, which all require iterative computation and some form of communication with either a central node [27] or neighbor-to-neighbor on a connected graph [163, 165, 198]. Distributed model predictive control (MPC) optimizes a networked system composed of subsystems over a time horizon, which can be decentralized (no communication) if the dynamic interconnections between subsystems are weak in order to achieve closed-loop stability as well as performance [39]. The work of [226] extended this to systems with strong coupling by employing time-varying distributed terminal set constraints, which requires neighbor-to-neighbor communication. Another class of methods model problems in which agents try to cooperate on a common objective without full state information as a decentralized partially observable Markov decision process (Dec-POMDP) [157]. [151] introduce networked distributed POMDPs, a variant of the Dec-POMDP inspired in part by the pairwise interaction paradigm of distributed constraint optimization problems (DCOPs).

Although the specific algorithms in these works differ significantly from the regression-based decentralization scheme we consider here, a larger difference is in problem formulation. As described in Sec. 6.3, we study a static optimization problem repeatedly solved at each time step. Much prior work, especially in optimal control (e.g. MPC) and reinforcement learning (e.g. Dec-POMDPs), poses the problem in a dynamic setting where the goal is to minimize cost over some time horizon. In the context of reinforcement learning (RL), the time horizon can be very long, leading to the well known tradeoff between exploration and exploitation; this does not appear in the static case. Additionally, many existing methods for the dynamic setting require an ongoing communication strategy between agents – though not all, e.g. [161]. Even one-shot static problems such as DCOPs tend to require complex communication strategies, e.g. [145].

Although the mathematical formulation of our approach is rather different from prior work, the policies we compute are similar in spirit to other learning and robotic techniques that have been proposed, such as behavioral cloning [176] and apprenticeship learning [1], which aim to let an agent learn from examples. In addition, we see a parallel with recent work on information-theoretic bounded rationality [158] which seeks to formalize decision-making with limited resources such as the time, energy, memory, and computational effort allocated for arriving at a decision. Our work is also related to swarm robotics [28], as it learns simple rules aimed to design robust, scalable and flexible collective behaviors for coordinating a large number of agents or robots.

6.3 General Problem Formulation

Consider a distributed multi-agent problem defined by a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, with \mathcal{N} denoting the nodes in the network with cardinality $|\mathcal{N}| = N$, and \mathcal{E} representing the set of edges between nodes. Figure 6.1a shows a prototypical graph of this sort. Each node has a real-valued state vector $x_i \in \mathbb{R}^{\alpha_i}, i \in \mathcal{N}$. A subset of nodes $\mathcal{C} \subset \mathcal{N}$, with cardinality $|\mathcal{C}| = C$, are controllable and hence are termed “agents.” Each of these agents has an action variable $u_i \in \mathbb{R}^{\beta_i}, i \in \mathcal{C}$. Let $x = (x_1, \dots, x_N)^\top \in \mathbb{R}^{\sum_{i \in \mathcal{N}} \alpha_i} = \mathcal{X}$ denote the full network state vector and $u \in \mathbb{R}^{\sum_{i \in \mathcal{C}} \beta_i} = \mathcal{U}$ the stacked network optimization variable. Physical constraints such as spatial coupling are captured through equality constraints $g(x, u) = 0$. In addition, the system is subject to inequality constraints $h(x, u) \leq 0$ that incorporate limits due to capacity, safety, robustness, etc. We are interested in minimizing a convex scalar function $f_o(x, u)$ that encodes objectives that are to be pursued cooperatively by all agents in the network, i.e. we want to find

$$\begin{aligned} u^* = \arg \min_u \quad & f_o(x, u), \\ \text{s.t.} \quad & g(x, u) = 0, \quad h(x, u) \leq 0. \end{aligned} \tag{6.1}$$

Note that (6.1) is static in the sense that it does not consider the future evolution of the state x or the corresponding future values of cost f_o . We apply this static problem to sequential control tasks by repeatedly solving (6.1) at each time step. Note that this simplification from an explicitly dynamic problem formulation (i.e. one in which the objective function incorporates future costs) is purely for ease of exposition and for consistency with the OPF literature as in [193]. We could also consider the optimal policy which solves a dynamic optimal control or RL problem and the decentralized learning step in Sec. 6.3 would remain the same.

Since (6.1) is static, applying the learned decentralized policies repeatedly over time may lead to dynamical instability. Identifying when this will and will not occur is a key challenge in verifying the regression-based decentralization method, however it is beyond the scope of this work.

Decentralized Learning

We interpret the process of solving (6.1) as applying a well-defined function or stationary Markov policy $\pi^* : \mathcal{X} \rightarrow \mathcal{U}$ that maps an input collective state x to the optimal collective control or action u^* . We presume that this solution exists and can be computed offline. Our objective is to learn C decentralized policies $\hat{u}_i = \hat{\pi}_i(x_i)$, one for each agent $i \in \mathcal{C}$, based on T historical measurements of the states $\{x[t]\}_{t=1}^T$ and the offline computation of the corresponding optimal actions $\{u^*[t]\}_{t=1}^T$. Although each policy $\hat{\pi}_i$ individually aims to approximate u_i^* based on local state x_i , we are able to reason about how well their collective action can approximate π^* . Figure 6.2 summarizes the decentralized learning setup.

More formally, we describe the dependency structure of the individual policies $\hat{\pi}_i : \mathbb{R}^{\alpha_i} \rightarrow \mathbb{R}^{\beta_i}$ with a Markov Random Field (MRF) graphical model, as shown in Figure 6.1b. The \hat{u}_i

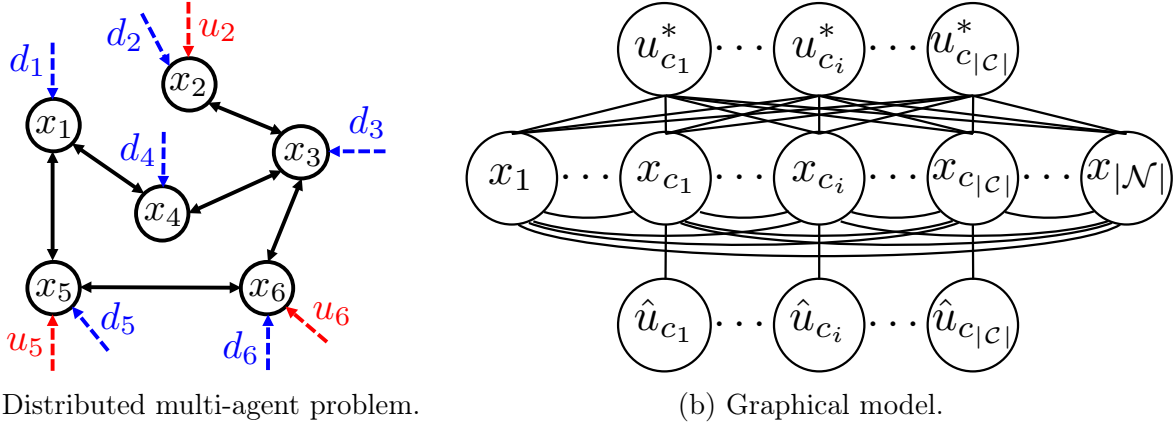


Figure 6.1: (a) Distributed multi-agent problem. The circles denote the local state x_i of an agent (a subset of (5.1)), the dashed arrow denotes its inputs u_i (red) and disturbances d_i (blue), and the double arrows denote physical coupling between the state variables of different agents, (b) Markov Random Field (MRF) graphical model of the dependency structure of all variables in the decentralized learning problem. Note that the state variables x_i and the optimal actions u_i^* form a fully connected undirected network, and the local policy \hat{u}_i only depends on the local state x_i .

are only allowed to depend on local state x_i while the u_i^* may depend on the full state x . With this model, we can determine how information is distributed among different variables and what information-theoretic constraints the policies $\{\hat{\pi}_i\}_{i \in \mathcal{C}}$ are subject to when collectively trying to reconstruct the centralized policy π^* . Note that although we may refer to π^* as globally optimal, this is not actually required for us to reason about how closely the $\hat{\pi}_i$ approximate π^* . That is, our analysis holds even if (6.1) is solved using approximate methods. In a dynamical reformulation of (6.1), for example, π^* could be generated using techniques from deep RL.

A Rate-Distortion Framework

We approach the problem of how well the decentralized policies $\hat{\pi}_i$ can perform in theory from the perspective of *rate distortion*. Rate distortion theory is a sub-field of information theory which provides a framework for understanding and computing the minimal *distortion* incurred by any given *compression* scheme. In a rate distortion context, we can interpret the fact that the output of each individual policy $\hat{\pi}_i$ depends only on the local state x_i as a compression of the full state x . For a detailed overview, see [43, Chapter 10]. We formulate

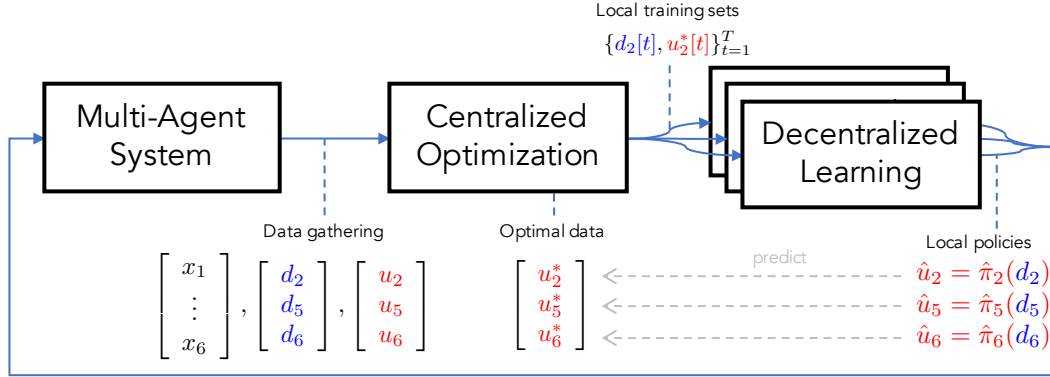


Figure 6.2: A flow diagram explaining the key steps of the decentralized regression method, depicted for the example system in Figure 6.1a. We first collect data from a multi-agent system, and then solve the centralized optimization problem using all the data. The data is then split into smaller training and test sets for all agents to develop individual decentralized policies $\hat{\pi}_i(x_i)$ (or $\hat{\pi}_i(d_i)$ as proposed in the context of OPF) that approximate the optimal solution of the centralized problem. These policies are then implemented in the multi-agent system to collectively achieve a common global behavior.

the following variant of the the classical rate distortion problem

$$\begin{aligned}
 D^* &= \min_{p(\hat{u}|u^*)} \mathbb{E}[\mathbf{d}(\hat{u}, u^*)], & (6.2) \\
 \text{s.t.} \quad & I(\hat{u}_i; u_j^*) \leq I(x_i; u_j^*) \triangleq \gamma_{ij}, \\
 & I(\hat{u}_i; \hat{u}_j) \leq I(x_i; x_j) \triangleq \delta_{ij}, \forall i, j \in \mathcal{C},
 \end{aligned}$$

where $I(\cdot, \cdot)$ denotes mutual information and $\mathbf{d}(\cdot, \cdot)$ an arbitrary non-negative distortion measure. As usual, the minimum distortion between random variable u^* and its reconstruction \hat{u} may be found by minimizing over conditional distributions $p(\hat{u}|u^*)$.

The novelty in (6.2) lies in the structure of the constraints. Typically, D^* is written as a function $D(R)$, where R is the maximum *rate* or mutual information $I(\hat{u}; u^*)$. From Figure 6.1b however, we know that pairs of reconstructed and optimal actions cannot share more information than is contained in the intermediate nodes in the graphical model, e.g. \hat{u}_1 and u_1^* cannot share more information than x_1 and u_1^* . This is a simple consequence of the data processing inequality [43, Thm. 2.8.1]. Similarly, the reconstructed optimal actions at two different nodes cannot be more closely related than the measurements x_i 's from which they are computed. The resulting constraints are fixed by the joint distribution of the state x and the optimal actions u^* . That is, they are fully determined by the structure of the optimization problem (6.1) that we wish to solve.

We emphasize that we have made virtually no assumptions about the distortion function. For the remainder of this chapter, we will measure distortion as the deviation between \hat{u}_i and u_i^* . However, we could also define it to be the suboptimality gap $f_o(x, \hat{u}) - f_o(x, u^*)$,

which may be much more complicated to compute. This definition could allow us to reason explicitly about the cost of decentralization, and it could address the valid concern that the optimal decentralized policy may bear no resemblance to π^* . We leave further investigation for future work.

Example: Squared Error, Jointly Gaussian

To provide more intuition into the rate distortion framework, we consider an idealized example in which the $x_i, u_i \in \mathbb{R}^1$. Let $\mathbf{d}(\hat{u}, u^*) = \|\hat{u} - u^*\|_2^2$ be the squared error distortion measure, and assume the state x and optimal actions u^* to be jointly Gaussian. These assumptions allow us to derive an explicit formula for the optimal distortion D^* and corresponding regression policies $\hat{\pi}_i$. We begin by stating an identity for two jointly Gaussian $X, Y \in \mathbb{R}$ with correlation ρ : $I(X; Y) \leq \gamma \iff \rho^2 \leq 1 - e^{-2\gamma}$, which follows immediately from the definition of mutual information and the formula for the entropy of a Gaussian random variable. Taking $\rho_{\hat{u}_i, u_i^*}$ to be the correlation between \hat{u}_i and u_i^* , $\sigma_{\hat{u}_i}^2$ and $\sigma_{u_i^*}^2$ to be the variances of \hat{u}_i and u_i^* respectively, and assuming that u_i^* and \hat{u}_i are of equal mean (unbiased policies $\hat{\pi}_i$), we can show that the minimum distortion attainable is

$$D^* = \min_{p(\hat{u}|u^*)} \mathbb{E} [\|u^* - \hat{u}\|_2^2] : \rho_{\hat{u}_i, u_i^*}^2 \leq 1 - e^{-2\gamma_i} = \rho_{u_i^*, x_i}^2, \forall i \in \mathcal{C}, \quad (6.3)$$

$$= \min_{\{\rho_{\hat{u}_i, u_i^*}\}, \{\sigma_{\hat{u}_i}\}} \sum_i \left(\sigma_{u_i^*}^2 + \sigma_{\hat{u}_i}^2 - 2\rho_{\hat{u}_i, u_i^*} \sigma_{u_i^*} \sigma_{\hat{u}_i} \right) : \rho_{\hat{u}_i, u_i^*}^2 \leq \rho_{u_i^*, x_i}^2, \quad (6.4)$$

$$= \min_{\{\sigma_{\hat{u}_i}\}} \sum_i \left(\sigma_{u_i^*}^2 + \sigma_{\hat{u}_i}^2 - 2\rho_{u_i^*, x_i} \sigma_{u_i^*} \sigma_{\hat{u}_i} \right), \quad (6.5)$$

$$= \sum_i \sigma_{u_i^*}^2 (1 - \rho_{u_i^*, x_i}^2). \quad (6.6)$$

In (6.4), we have solved for the optimal correlations $\rho_{\hat{u}_i, u_i^*}$. Unsurprisingly, the optimal value turns out to be the maximum allowed by the mutual information constraint, i.e. \hat{u}_i should be as correlated to u_i^* as possible, and in particular as much as u_i^* is correlated to x_i . Similarly, in (6.5) we solve for the optimal $\sigma_{\hat{u}_i}$, with the result that at optimum, $\sigma_{\hat{u}_i} = \rho_{u_i^*, x_i} \sigma_{u_i^*}$. This means that as the correlation between the local state x_i and the optimal action u_i^* decreases, the variance of the estimated action \hat{u}_i decreases as well. As a result, the learned policy will increasingly “bet on the mean” or “listen less” to its local measurement to approximate the optimal action.

Moreover, we may also provide a closed form expression for the regressor which achieves the minimum distortion D^* . Since we have assumed that each u_i^* and the state x are jointly Gaussian, we may write any u_i^* as an affine function of x_i plus independent Gaussian noise. Thus, the minimum mean squared estimator is given by the conditional expectation

$$\hat{u}_i = \hat{\pi}_i(x_i) = \mathbb{E}[u_i^* | x_i] = \mathbb{E}[u_i^*] + \frac{\rho_{u_i^*, x_i} \sigma_{u_i^*}}{\sigma_{x_i}} (x_i - \mathbb{E}[x_i]). \quad (6.7)$$

Thus, we have found a closed form expression for the best regressor $\hat{\pi}_i$ to predict u_i^* from only x_i in the joint Gaussian case with squared error distortion. This result comes as a direct consequence of knowing the true parameterization of the joint distribution $p(u^*, x)$ (in this case, as a Gaussian).

Determining Minimum Distortion in Practice

Often in practice, we do not know the parameterization $p(u^*|x)$ and hence it may be intractable to determine D^* and the corresponding decentralized policies $\hat{\pi}_i$. However, if one can assume that $p(u^*|x)$ belongs to a family of parameterized functions (for instance universal function approximators such as deep neural networks), then it is theoretically possible to attain or at least approach minimum distortion for arbitrary non-negative distortion measures.

Practically, one can compute the mutual information constraint $I(u_i^*, x_i)$ from (6.2) to understand how much information a regressor $\hat{\pi}_i(x_i)$ has available to reconstruct u_i^* . In the Gaussian case, we were able to compute this mutual information in closed form. For data from general distributions however, there is often no way to compute mutual information analytically. Instead, we rely on access to sufficient data $\{x[t], u^*[t]\}_{t=1}^T$, in order to estimate mutual informations numerically. In such situations (e.g. Section 6.5), we discretize the data and then compute mutual information with a minimax risk estimator, as proposed by [111].

6.4 Allowing Restricted Communication

Suppose that a decentralized policy $\hat{\pi}_i$ suffers from insufficient mutual information between its local measurement x_i and the optimal action u_i^* . In this case, we would like to quantify the potential benefits of communicating with other nodes $j \neq i$ in order to reduce the distortion limit D^* from (6.2) and improve its ability to reconstruct u_i^* . In this section, we present an information-theoretic solution to the problem of how to choose optimally which other data to observe, and we provide a lower bound-achieving solution for the idealized Gaussian case introduced in Section 6.3. We assume that in addition to observing its own local state x_i , each $\hat{\pi}_i$ is allowed to depend on at most k other $x_{j \neq i}$.

Theorem 6.4.1. (*Restricted Communication*)

If \mathcal{S}_i is the set of k nodes $j \neq i \in \mathcal{N}$ which $\hat{\pi}_i$ is allowed to observe in addition to x_i , then setting

$$\mathcal{S}_i = \arg \max_{\mathcal{S}} I(u_i^*; x_i, \{x_j : j \in \mathcal{S}\}) : |\mathcal{S}| = k, \quad (6.8)$$

minimizes the best-case expectation of any distortion measure. That is, this choice of \mathcal{S}_i yields the smallest lower bound D^ from (6.2) of any possible choice of \mathcal{S} .*

Proof. By assumption, \mathcal{S}_i maximizes the mutual information between the observed local states $\{x_i, x_j : j \in \mathcal{S}_i\}$ and the optimal action u_i^* . This mutual information is equivalent

to the notion of *rate* R in the classical rate distortion theorem [43]. It is well-known that the distortion rate function $D(R)$ is convex and monotone decreasing in R . Thus, by maximizing mutual information R we are guaranteed to minimize distortion $D(R)$, and hence D^* . \square

Theorem 6.4.1 provides a means of choosing a subset of the state $\{x_j : j \neq i\}$ to communicate to each decentralized policy $\hat{\pi}_i$ that minimizes the corresponding best expected distortion D^* . Practically speaking, this result may be interpreted as formalizing the following intuition: “the best thing to do is to transmit the most information.” In this case, “transmitting the most information” corresponds to allowing $\hat{\pi}_i$ to observe the set \mathcal{S} of nodes $\{x_j : j \neq i\}$ which contains the most information about u_i^* . Likewise, by “best” we mean that \mathcal{S}_i minimizes the best-case expected distortion D^* , for any distortion metric d . As in Section 6.3, without making some assumption about the structure of the distribution of x and u^* , we cannot guarantee that any particular regressor $\hat{\pi}_i$ will attain D^* . Nevertheless, in a practical situation where sufficient data $\{x[t], u^*[t]\}_{t=1}^T$ is available, we can solve (6.8) by estimating mutual information [111].

Example: Joint Gaussian, Squared Error with Communication

Here, we reexamine the joint Gaussian-distributed, mean squared error distortion case from Section 6.3, and apply Thm. 6.4.1. We will take $u^* \in \mathbb{R}^1, x \in \mathbb{R}^{10}$ and u^*, x jointly Gaussian with zero mean and arbitrary covariance. The specific covariance matrix Σ of the joint distribution $p(u^*, x)$ is visualized in Figure 6.3a. For simplicity, we show the squared correlation coefficients of Σ which lie in $[0, 1]$. The boxed cells in Σ in Figure 6.3a indicate that x_9 solves (6.8), i.e. $j = 9$ maximizes $I(u^*; x_1, x_j)$ the mutual information between the observed data and regression target u^* . Intuitively, this choice of j is best because x_9 is highly correlated to u^* and weakly correlated to x_1 , which is already observed by \hat{u} ; that is, it conveys a significant amount of information about u^* that is not already conveyed by x_1 .

Figure 6.3b shows empirical results. Along the horizontal axis we increase the value of k , the number of additional variables x_j which regressor $\hat{\pi}_i$ observes. The vertical axis shows the resulting average distortion. We show results for a linear regressor of the form of (6.7) where we have chosen \mathcal{S}_i optimally according to (6.8), as well as uniformly at random from all possible sets of unique indices. Note that the optimal choice of \mathcal{S}_i yields the lowest average distortion D^* for all choices of k . Moreover, the linear regressor of (6.7) achieves D^* for all k , since we have assumed a Gaussian joint distribution.

6.5 Application to Optimal Power Flow

In this case study, we aim to minimize the voltage variability in an electric grid caused by intermittent renewable energy sources and the increasing load caused by electric vehicle charging. We do so by controlling the reactive power output of distributed energy resources (DERs), while adhering to the physics of power flow and constraints due to energy capacity

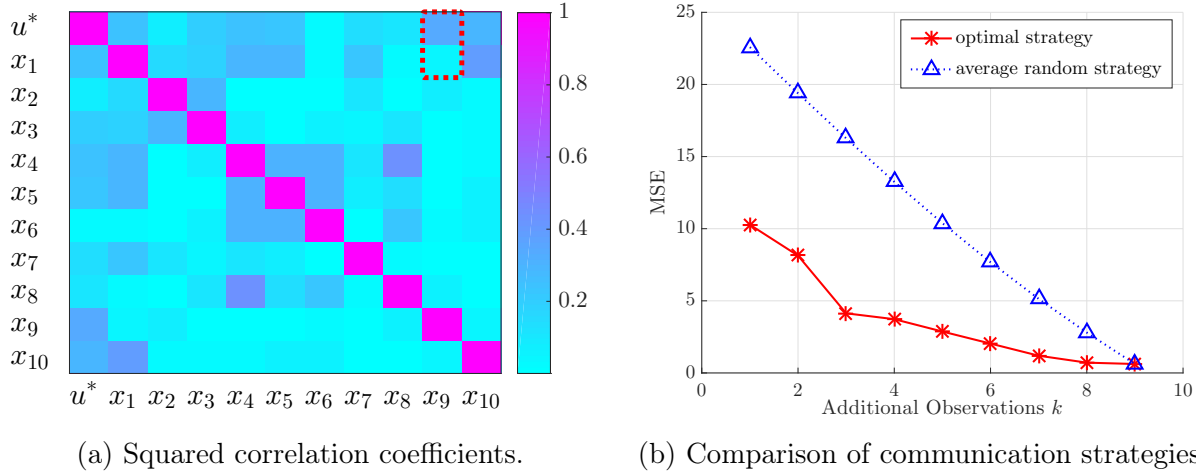


Figure 6.3: Results for optimal communication strategies on a synthetic Gaussian example. (a) shows squared correlation coefficients between of u^* and all x_i 's. The boxed entries correspond to x_9 , which was found to be optimal for $k = 1$. (b) shows that the optimal communication strategy of Thm. 6.4.1 achieves the lowest average distortion and outperforms the average over random strategies.

and safety. Recently, various approaches have been proposed, such as [74] or [228]. In these methods, DERs tend to rely on an extensive communication infrastructure, either with a central master node [223] or between agents leveraging local computation [47]. We study regression-based decentralization as outlined in Section 6.3 and Figure 6.2 to the optimal power flow (OPF) problem [136], as initially proposed by [193]. We apply Thm. 6.4.1 to determine the communication strategy that minimizes optimal distortion to further improve the reconstruction of the optimal actions u_i^* .

Solving OPF requires a model of the electricity grid describing both topology and impedances; this is represented as a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. For clarity of exposition and without loss of generality, we introduce the linearized power flow equations over radial networks, also known as the *LinDistFlow* equations [17] (described in detail in Chapter 3):

$$P_{ij} = \sum_{(j,k) \in \mathcal{E}, k \neq i} P_{jk} + p_j^c - p_j^g, \quad (6.9a)$$

$$Q_{ij} = \sum_{(j,k) \in \mathcal{E}, k \neq i} Q_{jk} + q_j^c - q_j^g, \quad (6.9b)$$

$$y_j = y_i - 2(r_{ij}P_{ij} + \xi_{ij}Q_{ij}) \quad (6.9c)$$

In this model, capitals P_{ij} and Q_{ij} represent real and reactive power flow on a branch from node i to node j for all branches $(i, j) \in \mathcal{E}$, lower case p_i^c and q_i^c are the real and reactive power consumption at node i , and p_i^g and q_i^g are its real and reactive power generation. Complex line impedances $r_{ij} + \sqrt{-1}\xi_{ij}$ have the same indexing as the power flows. The *LinDistFlow*

equations use the squared voltage magnitude y_i , defined and indexed at all nodes $i \in \mathcal{N}$. These equations are included as constraints in the optimization problem to enforce that the solution adheres to laws of physics.

To formulate our decentralized learning problem, we will treat $x_i \triangleq (p_i^c, q_i^c, p_i^g)$ to be the local state variable, and, for all controllable nodes, i.e. agents $i \in \mathcal{C}$, we have $u_i \triangleq q_i^g$, i.e. the reactive power generation can be controlled (y_i, P_{ij}, Q_{ij} are treated as dummy variables). We assume that for all nodes $i \in \mathcal{N}$, consumption p_i^c , q_i^c and real power generation p_i^g are predetermined respectively by the demand and the power generated by a potential photovoltaic (PV) system. The action space is constrained by the reactive power capacity $|u_i| = |q_i^g| \leq \bar{q}_i$. In addition, voltages are maintained within $\pm 5\%$ of 120V, which is expressed as the constraint $\underline{y} \leq y_i \leq \bar{y}$. The OPF problem now reads

$$\begin{aligned} u^* = \arg \min_{q_i^g, \forall i \in \mathcal{C}} \quad & \sum_{i \in \mathcal{N}} |y_i - y_{\text{ref}}|, \\ \text{s.t.} \quad & (6.9), |q_i^g| \leq \bar{q}_i, \underline{y} \leq y_i \leq \bar{y}. \end{aligned} \tag{6.10}$$

Following Figure 6.2, we employ models of real electrical distribution grids (including the IEEE Test Feeders [107]), which we equip with with T historical readings $\{x[t]\}_{t=1}^T$ of load and PV data, which is composed with real smart meter measurements sourced from [50]. We solve (6.10) for all data, yielding a set of minimizers $\{u^*[t]\}_{t=1}^T$. We then separate the overall data set into C smaller data sets $\{x_i[t], u_i^*[t]\}_{t=1}^T$, $\forall i \in \mathcal{C}$ and train linear policies with feature kernels $\varphi_i(\cdot)$ and parameters θ_i of the form $\hat{\pi}_i(x_i) = \theta_i^\top \varphi_i(x_i)$. Practically, the challenge is to select the best feature kernel $\varphi_i(\cdot)$. We extend earlier work which showed that decentralized learning for OPF can be done satisfactorily via a hybrid forward- and backward-stepwise selection algorithm [86, Chapter 3] that uses a quadratic feature kernels.

Figure 6.4a shows the result for an electric distribution grid model based on a real network from Arizona. This network has 129 nodes and, in simulation, 53 nodes were equipped with a controllable DER (i.e. $N = 129, C = 53$). In Figure 6.4a we show the voltage deviation from a normalized setpoint on a simulated network with data not used during training. The improvement over the no-control baseline is striking, and performance is nearly identical to the optimum achieved by the centralized solution. Concretely, we observed: (i) no constraint violations, and (ii) a suboptimality deviation of 0.15% on average, with a maximum deviation of 1.6%, as compared to the optimal policy π^* .

In addition, we applied Thm. 6.4.1 to the OPF problem for a smaller network [107], in order to determine the optimal communication strategy to minimize a squared error distortion measure. Figure 6.4b shows the mean squared error distortion measure for an increasing number of observed nodes k and shows how the optimal strategy outperforms an average over random strategies.

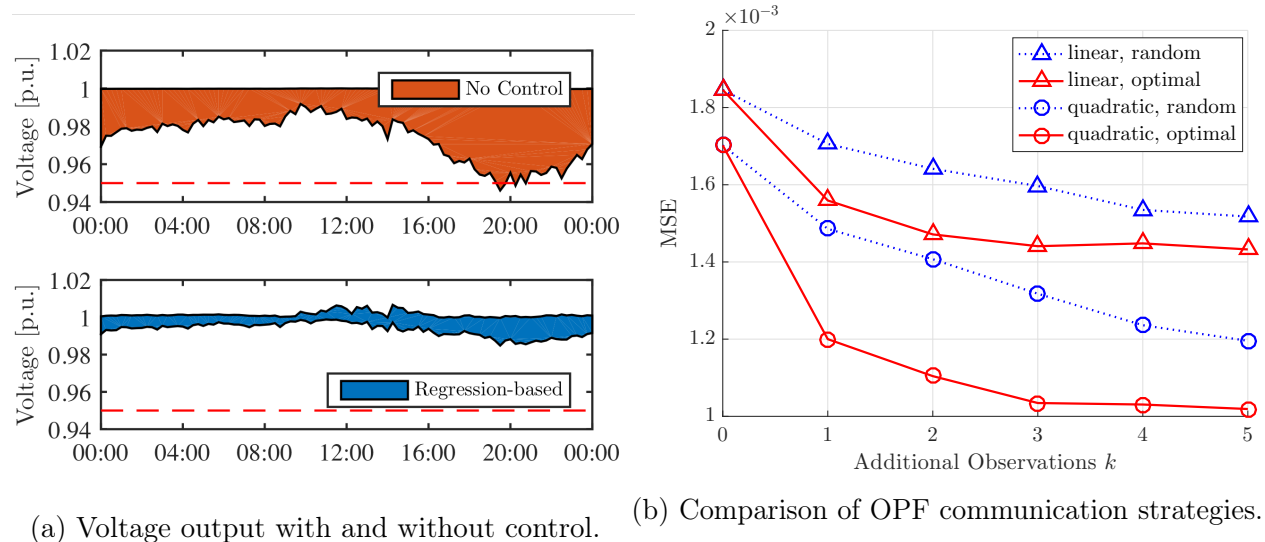


Figure 6.4: Results for decentralized learning on an OPF problem. (a) shows an example result of decentralized learning - the shaded region represents the range of all voltages in a network over a full day. As compared to no control, the fully decentralized regression-based control reduces voltage variation and prevents constraint violation (dashed line). (b) shows that the optimal communication strategy \mathcal{S}_i outperforms the average for random strategies on the mean squared error distortion metric. The regressors used are stepwise linear policies $\hat{\pi}_i$ with linear or quadratic features.

6.6 Conclusions and Future Work

This chapter generalizes the approach developed in Chapter 5 to solve multi-agent static optimal control problems with decentralized policies that are learned offline from historical data. Our rate distortion framework facilitates a principled analysis of the performance of such decentralized policies and the design of optimal communication strategies to improve individual policies. These techniques work well on a model of a sophisticated real-world OPF example.

There are still many open questions about regression-based decentralization. It is well known that strong interactions between different subsystems may lead to instability and suboptimality in decentralized control problems [51]. There are natural extensions of our work to address dynamic control problems more explicitly, and stability analysis is a topic of ongoing work. Also, analysis of the suboptimality of regression-based decentralization should be possible within our rate distortion framework. Finally, it is worth investigating the use of deep neural networks to parameterize both the distribution $p(u^*|x)$ and local policies $\hat{\pi}_i$ in more complicated decentralized control problems with arbitrary distortion measures.

Chapter 7

Customized Differential Privacy for Distributed Optimization in Multi-Agent Systems

“I never said, ‘I want to be alone.’ I only said ‘I want to be let alone!’ There is all the difference.”

- Greta Garbo

Submitted to the IEEE Transactions on Control of Network Systems [56].

Real-time data-driven optimization and control problems over networks may require sensitive information of participating users to calculate solutions and decision variables, such as in traffic or energy systems. Adversaries with access to coordination signals may potentially decode information on individual users and put user privacy at risk. Work in differential privacy for distributed optimization so far has considered focusing on specific instances, protecting either a cost function or constraint parameters, and imposing a uniform privacy level for all agents involved in the mechanism. We develop local differential privacy, which is a strong notion that guarantees user privacy regardless of any auxiliary information an adversary may have, for a large family of convex distributed optimization problems. The nature of the privacy result allows each agent to customize its own level of *local differential privacy* based on local needs and parameter sensitivities. We propose a general sampling based approach for determining sensitivity and derive analytical bounds for specific quadratic problems. We derive a general sub-optimality bound as a function of the cumulative variance of the noise injected by all agents. We analyze inherent trade-offs between privacy and sub-optimality as a function of noise variance and the number of iterations in optimizing. Our algorithm is implemented for distributed optimization in electric grids, and distributed

model-predictive control in building energy management. We propose various allocation schemes to divide the maximum allowable noise, a *privacy budget*, among all participating agents, either through proportional sharing or well studied pricing mechanisms.

7.1 Introduction

Motivation

The electric grid is rapidly changing due to the advent of renewable generation and the electrification of transportation, heating and other vital systems. These changes lead to an unprecedented level of variability, uncertainty and increased risk of overload and power imbalance, and motivate the use of *flexibility* to mitigate these risks in real-time.

In distribution systems, which traditionally had very limited diagnostic and control capabilities, the slow response time of legacy control infrastructure, such as capacitor banks and voltage transformers, prohibits distribution system operators (DSOs) to mitigate these new dynamics, giving way to accelerated wear and failure of network components that can cause cascading blackouts, as is being increasingly experienced in areas with high solar adoption rates [195]. As a result, *distributed energy services* are considered to help regulate power flow and voltages by using inverter-interfaced distributed energy resources (DERs), such as batteries, electric vehicles, solar installations or smart capacitors. Inverters can respond fast and help stabilize voltage and power flow by adjusting their intake or outtake of power, thereby mitigating effects of variability and overload locally, faster and closer to the source. Orchestrating such mitigation in a network with multiple inverters is often cast as an optimal power flow problem (OPF) with power flow and capacity constraints and an objective that minimizes network losses, voltage deviations or real power consumption [80]. In transmission systems, *demand response* schemes are rapidly adopted to mitigate power imbalances resulting from inadequate forecasting and increased variability, by coordinating energy consumers to quickly adjust the consumption of specific devices or DERs.

In both settings, an optimization and control problem is formulated to determine the optimal actions to be computed, communicated and implemented in real-time throughout a network. Often the scale of such a problem desires a distributed implementation that can be solved quickly enough to allow for high frequency control actions. To enable this, a network may be split up into sub-networks governed by different agents, who exchange their local optimization variables with neighbors and/or a central operator to iteratively solve the optimization problem. Exchanging optimization variables between agents and the changes therein may reveal private information, such as whether someone is home and what kind of appliances someone is using [100]. In addition, there is growing understanding that secondary information may be inferred from the communicated variables, including the parameters used in the local objective and constraints, which may reveal sensitive information such as prices and capacity [97].

To make matters more challenging, different DER owners or energy customers may be

competing with each another to serve the utility company with their flexibility. Knowing the capacity and prices negotiated by other players can help in negotiating with the DSO and can lead to strategic behavior that includes untruthful communication in the bidding process for flexibility, thereby harming the quality and adequacy of the distributed optimization or control problem. As such, both personal privacy and commercial privacy needs require the development of agent-to-agent distributed optimization algorithms that can mask sensitive information in objectives and constraints.

In recent years, various privacy-preserving algorithms have been proposed for distributed optimization and control problems, using various privacy metrics. The differential privacy framework [67] has gained the most attention, and is particularly lauded for its robustness to auxiliary side information that an adversary might have to complement information gained from a particular algorithm, providing stronger privacy guarantees than other existing metrics (see [61] for an overview). A recent textbook on differential privacy provides detailed information and various applications [66].

The framework assumes a setting in which sensitive information is stored in a database by a trustworthy curator, which can provide answers to external queries. A system is made differentially private by randomizing its answers in such a way that the distribution over published outputs is not too sensitive to changes in the stored data. Perturbation can be designed to make it provably difficult for an adversary to make inferences about individual records from the published outputs. In the setting of d-OPF, each DER owner is its own curator managing its own locally private information and communication of its optimization variables to neighboring DER owners. In order to preserve differential privacy, each curator has to ensure that the output of queries, that is the communicated variables, remain approximately unchanged if local parameters relating to its objective or constraints are modified.

Related Work

This work complements an existing and rapidly growing body of literature on incorporating *differential privacy* into resource allocation and, most relevant here, in distributed optimization, control and networked systems. Earlier work by Hsu et al. [104] develops differential privacy-preserving algorithms for convex optimization problems that are solved in a *central* fashion, considering general linear programs (LPs) with either private objectives or constraint. Dong et al. [62] consider privacy in a game theoretic environment, motivated by traffic routing in which the origins and destinations of drivers are considered private. Jia et al. [110] consider occupancy-based HVAC control and treat the control objective and the location traces of individual occupants as private variables, using an information-theoretic privacy metric. A recent tutorial paper by Cortés et al. [42] covers differential privacy for *distributed* optimization, and distinguishes between objective-perturbing and message-perturbing strategies for distributed optimization. In the first category, the objective function of each agent is perturbed with noise in a differentially private manner, which guarantees differential privacy at the functional level and is preferred for systems with asymptotically stable dynamics [154].

In the second category, coordination messages are perturbed with noise before sent, either to neighbors or a central node, depending on the specific algorithm. Huang et al. [106] propose a technique for disguising private information in the local objective function for distributed optimization problems with strongly convex separable objective functions and convex constraints. Han et al. [97] consider problems where the private information is encoded in the individual constraints, the objective functions need to be convex and Lipschitz continuously differentiable, and the constraints have to be convex and separable. Other related works are Mo and Murray [144] who aim to preserve privacy of agents' initial states in average consensus and Katewa et al. [118] who explore an alternative trade-off between privacy and the value of cooperation (rather than performance) in distributed optimization. In [185], a privacy-aware optimization algorithm is analyzed using the cryptography notion of zero knowledge proofs. More recently, [229] considers differentially private algorithms over time-varying directed networks.

The above works are restrictive in two ways. Firstly, these consider privacy-preserving mechanisms for constraints, objectives or initial states only. An exception is the work Hsu et al. [104] on linear programs, which can handle both private objectives and constraints. Secondly, the existing work considers uniform privacy levels for agents across a network, not facilitating potential variation in actual privacy needs. As such, this chapter provides a more general set of tools by proposing a mechanism that preserves private objectives and constraints for optimization problems with strongly convex objectives and convex constraints, and allowing for customized local differential privacy across a network of agents.

Contributions

Motivated by personal and commercial privacy concerns in distributed OPF, we investigate the problem of *preserving differential privacy of local objectives and constraints in distributed constrained optimization with agent-to-agent communication*. Compared to previous works on privacy-aware distributed optimization (e. g. [97], [106]), we consider the notion of *local differential privacy*, which is a refined (and more stringent) version of differential privacy considered in other related works (see, for example, [65]). It allows each agent in a network to *customize* its own privacy level, based on individual preferences and characteristics. Furthermore, most previous work considers privacy protection for either individual objective function ([106]) or the individual constraint ([97]), our more general formulation enables us to provide privacy guarantees on both local objective function parameters and local constraint parameters. Specifically, the proposed algorithm solves a general class of convex optimization problems where each agent has a local objective function and a local constraint, and agents communicate with neighbors/adjacent agents with no need for a central authority.

We show that the private optimization algorithm can be formulated as an instance of the Inexact Alternating Minimization Algorithm (IAMA) for Distributed Optimization [163]. This algorithm allows provable convergence under computation and communication errors. This property is exploited to provide privacy by injecting noise large enough to hide sensitive information, while small enough to exploit the convergence properties of the IAMA. We derive

the trade-off between the privacy level and sub-optimality of the algorithm. The trade-off between sub-optimality and differential privacy allows us to determine a *privacy budget* that captures the allowable cumulative variance of noise injected throughout the network that achieves a desired level of (sub-)optimality. We propose two pricing schemes to ensure fair and efficient allocation of the privacy budget over all participating/bidding DER owners.,

7.2 Preliminaries and Problem Statement

Distributed optimization problem

In this section, we consider a distributed optimization problem on a network of M sub-systems (nodes). The sub-systems communicate according to a fixed undirected graph $G = (\mathcal{V}, \mathcal{E})$. The vertex set $\mathcal{V} = \{1, 2, \dots, M\}$ represents the sub-systems and the edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ specifies pairs of sub-systems that can communicate. If $(i, j) \in \mathcal{E}$, we say that sub-systems i and j are neighbors, and we denote by $\mathcal{N}_i = \{j | (i, j) \in \mathcal{E}\}$ the set of the neighbors of sub-system i . Note that \mathcal{N}_i includes i . The cardinality of \mathcal{N}_i is denoted by $|\mathcal{N}_i|$. We use a vector $[v]_i$ to denote the local variable of subsystem i and $[v]_i$ can be of different dimensions for different i . The collection of these local variables is denoted as $v = [v_1^T, \dots, v_M^T]^T$. Furthermore, the concatenation of the local variable $[v]_i$ of sub-system i and the variables of its neighbors $[v]_j, j \in \mathcal{N}_i$ is denoted by z_i . With appropriate selection matrices E_i and F_{ji} , the variables have the following relationship: $z_i = E_i v$ and $[v]_i = F_{ji} z_j, j \in \mathcal{N}_i$, which implies the relation between the local variable $[v]_i$ and the global variable v , i.e. $[v]_i = F_{ji} E_j v, j \in \mathcal{N}_i$. We consider the following distributed optimization problem:

Problem 7.2.1 (Distributed Optimization).

$$\min_{z, v} \sum_{i=1}^M f_i(z_i) \tag{7.1}$$

$$s.t. \quad z_i \in \mathbb{C}_i, \quad z_i = E_i v, \quad i = 1, 2, \dots, M, \tag{7.2}$$

where f_i is the local cost function for node i which is assumed to be strongly convex with a convexity modulus $\rho_{f_i} > 0$, and to have a Lipschitz continuous gradient with a Lipschitz constant $L(\nabla f_i) > 0$. The constraint \mathbb{C}_i is assumed to be a convex set which represents a convex local constraint on z_i , i.e. the concatenation of the variables of sub-system i and the variables of its neighbors.

The above problem formulation is fairly general and can represent a large class of problems in practice. In particular it includes the following quadratic programming problem, which we study as a particular instance in our applications.

Problem 7.2.2 (Distributed Quadratic Problem).

$$\begin{aligned} \min_{z,v} \quad & \sum_{i=1}^M z_i^T H_i z_i + h_i^T z_i \\ \text{s.t.} \quad & C_i z_i \leq c_i, \quad z_i = E_i v, \quad i = 1, 2, \dots, M, \end{aligned} \quad (7.3)$$

where $H_i \succ 0$. In particular, we will assume that the smallest eigenvalue of H_i satisfies $\lambda_{\min}(H_i) := \lambda_{\min}^{(i)} > 0$.

Local Differential Privacy

In this section, we will present definitions and properties for differential privacy. Let $\mathcal{P}, \mathcal{P}'$ be two databases with private parameters in some space \mathcal{X} containing information relevant in executing an algorithm. Let $d : \mathcal{X} \times \mathcal{X} \mapsto [0, \infty)$ denote a metric defined on \mathcal{X} . A mechanism or algorithm \mathcal{A} is a mapping from \mathcal{X} to some set denoting its output space.

Definition 7.2.1 (Differential Privacy). *A randomized algorithm \mathcal{A} is ϵ -differential private if for all $\mathcal{S} \subseteq \text{range}(\mathcal{A})$ and for all database $\mathcal{P}, \mathcal{P}'$ satisfying $d(\mathcal{P}, \mathcal{P}') \leq 1$, it holds that*

$$\Pr[\mathcal{A}(\mathcal{P}) \in \mathcal{S}] \leq e^\epsilon \cdot \Pr[\mathcal{A}(\mathcal{P}') \in \mathcal{S}], \quad (7.4)$$

where the probability space is over the mechanism \mathcal{A} .

This definition of differential privacy is suitable for cases where one uniform level of privacy needs to across all elements in the databases. We now consider a distributed algorithm $\mathcal{A}(\mathcal{P}_1, \dots, \mathcal{P}_M)$ in a network with M agents for solving an optimization problem in a collaborative way, where \mathcal{P}_i denotes the private parameters of agent i . The outputs of the mechanism are the message exchanged between nodes in the network over the time horizon of iterations. This mechanism induces M local mechanisms $\mathcal{A}_1(\mathcal{P}_1, \dots, \mathcal{P}_M), \dots, \mathcal{A}_M(\mathcal{P}_1, \dots, \mathcal{P}_M)$, each executed by one agent. The output of one local mechanism \mathcal{A}_i is the message sent out by node i , i.e. $\text{range}(\mathcal{A}_i) \subseteq \text{range}(\mathcal{A})$. It is important to realize that although one local mechanism, say \mathcal{A}_i , does not necessarily have direct access to the input/database $\mathcal{P}_j, j \neq i$ of other nodes, the output of \mathcal{A}_i could still be affected by $\mathcal{P}_j, j \neq i$ because of the interactions among different nodes. For this reason, we explicitly write $\mathcal{P}_1, \dots, \mathcal{P}_M$ as input to all local mechanism.

We now let each agent i specify its own level of privacy ϵ_i . To formalize this specification, we require a definition of *local differential privacy*:

Definition 7.2.2 (Local Differential Privacy). *Consider a (global) mechanism \mathcal{A} for a network with M nodes, and M local mechanisms $\mathcal{A}_i, i = 1, \dots, M$ induced by \mathcal{A} . We say that*

the mechanism \mathcal{A} is ϵ_i -differentially locally private for node i , if for any $\mathcal{S}_i \in \text{range}(\mathcal{A}_i)$ it satisfies that

$$\frac{\mathbb{P}\{\mathcal{A}_i(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M) \in \mathcal{S}_i\}}{\mathbb{P}\{\mathcal{A}_i(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M) \in \mathcal{S}_i\}} \leq e^{\epsilon_i}, \quad (7.5)$$

where $d(\mathcal{P}_i, \mathcal{P}'_i) \leq 1$. Moreover, we say that the mechanism \mathcal{A} is $(\epsilon_1, \dots, \epsilon_M)$ -differentially private, if \mathcal{A} is ϵ_i -differentially locally private for all nodes, where $i = 1, \dots, M$.

Figure 7.1 presents the concept of local differential privacy pictorially, showing the various considerations that can be taken when designing for local/customized privacy. Firstly, one may desire to include a central node 0 that communicates with all subsystems or implement a fully distributed problem between the subsystems that does not rely on any central node. The former will lead to better convergence properties as information spreads more easily throughout the network, the latter will benefit privacy by making it harder to collect information from across the network. Regardless, the method allows for the privacy to be purely local, strengthening the notion developed in [97], which implements differential privacy in distributed optimization through a trusted central node, assuming a star-shaped communication structure with no agent-to-agent communication. Secondly, subsystems may have varying levels of privacy. In Figure 7.1, subsystem 1 and 3 have a local privacy specification, while subsystems 0 and 2 do not. The systems with local differential privacy have outgoing messages perturbed by noise as indicated by the dashed arrows. As such, the method is flexible to various forms of distributed optimization or control problems with heterogeneous privacy and control properties across its nodes/agents, extending the work in [106], which considers a similar fully distributed algorithm but specifies a uniform privacy level.

7.3 Main Results

Differentially Private Distributed Optimization

In this section, we describe a distributed optimization algorithm (Algorithm 1) for solving Problem 7.2.1 with local differential privacy guarantees, based on the results in [163]. To solve the optimization problem in a distributed way, we split the overall problem into small local problems according to the physical couplings of the sub-systems. This approach gives the resulting algorithm the desired feature that each node only needs to communicate with its neighbors and the computations can be performed in parallel for every subsystem. To guarantee local differential privacy, a noise term is added to the message at each time, before it is sent out to other nodes.

We start with defining the private parameters of agent i to be the collection of parameters for its local objective function and constraints in Problem 7.2.1, $\mathcal{P}_i := (f_i, \mathbb{C}_i)$. Algorithm 1 can be seen as a *global mechanism* \mathcal{A} which takes input data $\mathcal{P}_1, \dots, \mathcal{P}_M$ and produces




Figure 7.1: A system with 4 subsystems. Arrows denote message directions in distributed optimization. Here, node 0 denotes a central node that communicates with all other nodes. Note that Problems 7.2.1 and 7.2.2 can be fully distributed, and hence a central node is not necessary. The other 3 nodes represent subsystems. Nodes that have a privacy specification are patterned with dots and send out messaged perturbed by noise, as indicated by dashed arrows.

Algorithm 1 Differentially private distributed algorithm

Require: Initialize $\mu_i^0 = 0 \in \mathbb{R}^{z_i}$, $\tau^0 = \min_{1 \leq i \leq M} \{\rho_{f_i}\}$ and $\tau^k = \frac{1}{\tau^{0k}}$

for $k = 1, 2, \dots$ **do**

- 1: $z_i^k = \operatorname{argmin}_{z_i \in \mathbb{C}_i} \{f_i(z_i) + \langle \mu_i^{k-1}, -z_i \rangle\} + \delta_i^k$
- 2: Send z_i^k to all the neighbors of agent i .
- 3: $[v^k]_i = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} [z_j^k]_i$.
- 4: Send $[v^k]_i$ to all the neighbors of agent i .
- 5: $\mu_i^k = \mu_i^{k-1} + \tau^k (E_i v^k - z_i^k)$

end for

output $z_j^k, [v^k]_j$ for all $j = 1, \dots, M$ and $k = 1, 2, \dots$. In particular, the output of \mathcal{A} up to iteration K is given by

$$\mathcal{A}^K(\mathcal{P}_1, \dots, \mathcal{P}_M) := (\mathbf{z}^1, \dots, \mathbf{z}^K, [\mathbf{v}]^1, \dots, [\mathbf{v}]^K), \quad (7.6)$$

where we use the bold letter \mathbf{z}^k to denote the collection of (z_1^k, \dots, z_M^k) at time k and $[\mathbf{v}]^k$ to denote the collection of $([v]_1^k, \dots, [v]_M^k)$ at time k . Recall that we are mainly concerned with local privacy at different nodes, we use \mathcal{A}_i to denote the *local mechanism* induced by \mathcal{A} at node i , which takes input data $\mathcal{P}_1, \dots, \mathcal{P}_M$ and produces output $z_i^k, [v^k]_i$ for $k = 1, 2, \dots$, namely

$$\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}_M) := (z_i^1, \dots, z_i^K, [v]_i^1, \dots, [v]_i^K). \quad (7.7)$$

Equipped with the definition of local differential privacy in Definition 7.2.2, we can give the formal definition of *locally differentially private* distributed algorithms.

Definition 7.3.1 (Locally Differentially Private Distributed Algorithms). *Consider the (global) mechanism \mathcal{A} given in Algorithm 1 and the M local mechanisms $\mathcal{A}_j, j = 1, \dots, M$ induced by \mathcal{A} . We say that the mechanism \mathcal{A} is locally ϵ -differentially private for node i after K iterations, if the induced local mechanism \mathcal{A}_i is ϵ -differentially private according to Definition 7.2.1, i.e., for any $\mathcal{S} \in \text{range}(\mathcal{A}_i)$ it satisfies that*

$$\frac{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M) \in \mathcal{S}\}}{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M) \in \mathcal{S}\}} \leq e^\epsilon, \quad (7.8)$$

where \mathcal{P}_i and \mathcal{P}'_i are two different local problems at node i with distance $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i) \leq 1$. Moreover, we say that the mechanism \mathcal{A} is $(\epsilon_1, \dots, \epsilon_M)$ -differentially private, if \mathcal{A} is locally ϵ_i -differentially private for all nodes $i = 1, \dots, M$.

The definition of the distance $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i)$ between two problems \mathcal{P}_i and \mathcal{P}'_i depends on specific applications. For the quadratic problem in (7.3), the local objective f_i and local constraint \mathbb{C}_i are parametrized by the matrix and vectors H_i, h_i, C_i, c_i . One possible definition of the distance for this special case is given by weighted sum of matrix/vector norms

$$\begin{aligned} \text{adj}(\mathcal{P}_i, \mathcal{P}'_i) &= a_1 \|H_i - H'_i\| + a_2 \|h_i - h'_i\| \\ &\quad + a_3 \|C_i - C'_i\| + a_4 \|c_i - c'_i\|, \end{aligned} \quad (7.9)$$

for some weights a_1, a_2, a_3, a_4 . The choice of the weights will put more emphasis sensitivity on a specific matrix or vector and its corresponding parameters. These can be used to define the “region of adjacency”, meaning the parameter differences that are considered to be meaningful for protecting with a differential privacy mechanism. The choice of the norm can also change the behavior, for instance by treating all parameters equal (such as for the ℓ_1 -norm or the Frobenius norm) or only worrying about the maximum distance (such as for the ℓ_∞ -norm). For some applications, only certain entries of the matrices and vectors represent private information. In this case, the distance should be defined only with respect to these “private entries”.

In the original setup of differential privacy where the database representation is considered, there is a natural definition of adjacency: namely two databases are adjacent if they differ by only one element. When extending the adjacency definition to the space of functions, there does not exist a natural candidate for adjacency. We point out that this is a common situation encountered in similar problem setups such as [97, 106]. However, resolving this issue would require an explicit connection between the privacy level ϵ and a concrete specification supplied by the application (such as in [42, Section B]), which is out of the scope for this chapter.

Privacy analysis

In this section, we derive the differential privacy level with Algorithm 1. To state the main result, we first define

$$g(\mathcal{P}_j, \mu) := \arg \min_{z \in \mathbb{C}_j} \{f_j(z) + \langle \mu, -z \rangle\}, \quad (7.10)$$

where \mathcal{P}_j encapsulates the local objective function f_j and constraints \mathbb{C}_j . The main result is given in the following theorem.

Theorem 7.3.1 (Local Differential Privacy). *Consider Algorithm 1 for solving Problem 7.2.1. Assume that each element of the noise vector δ_i^k in Algorithm 1 is independently chosen according the Laplace distribution with the density function $p(\delta_i^k) = \frac{1}{2\sigma_i^k} \exp(-\|\delta_i^k\|/\sigma_i^k)$ for $i = 1, \dots, M$ and $k = 1, \dots, K$. Then for all $i = 1, \dots, M$, this algorithm is locally ϵ_i -differentially private for node i where*

$$\epsilon_i = \Theta_i \sum_{k=1}^K \frac{1}{\sigma_i^k}, \quad (7.11)$$

and

$$\Theta_i := \max_{\text{adj}(\mathcal{P}_i, \mathcal{P}'_i) \leq 1, \mu} \|g(\mathcal{P}_i, \mu) - g(\mathcal{P}'_i, \mu)\| \quad (7.12)$$

is called the sensitivity of the optimization problem.

Proof: The proof is given in Appendix B ■

Sensitivity Calculation

In order to evaluate the privacy level ϵ_i provided in Theorem 7.3.1, we need to calculate the sensitivity Θ_i . This calculation is itself an optimization problem which can be written as follows.

$$\begin{aligned} \max \Theta_i &:= \|z^* - z'^*\|, & (7.13) \\ \text{s.t. } z^* &= \arg \min_{z \in \mathbb{C}_i} f_i(z) + \langle \mu, -z \rangle, \\ z'^* &= \arg \min_{z \in \mathbb{C}'_i} f'_i(z) + \langle \mu, -z \rangle, \\ \text{adj}(\mathcal{P}_i(f_i, \mathbb{C}_i), \mathcal{P}'_i(f'_i, \mathbb{C}'_i)) &\leq 1, \\ &\text{with variables } (z^*, z'^*, f_i, f'_i, \mathbb{C}_i, \mathbb{C}'_i, \mu). \end{aligned}$$

This problem belongs to the class of bi-level optimization problems, for which there is in general no efficient algorithm to find the global optimal solution. In the rest of this section, we will specialize our problem to the the class of quadratic problems in (7.3) and provide some refined analysis. In this case, we represent the optimality condition of z^*, z'^* in terms of the KKT condition of the optimization problem $\min_{z \in \mathbb{C}_i} f_i(z) + \langle \mu, -z \rangle$, the optimization

problem (7.13) can be rewritten explicitly

$$\max \Theta_i := \|z^* - z'^*\|, \quad (7.14)$$

$$\text{s.t. } z^* \in \text{KKT}(H_i, h_i, C_i, c_i, w_i, \mu), \quad (7.15)$$

$$z'^* \in \text{KKT}(H'_i, h'_i, C'_i, c'_i, w'_i, \mu), \quad (7.16)$$

$$\begin{aligned} & a_1 \|H_i - H'_i\| + a_2 \|h_i - h'_i\| \\ & + a_3 \|C_i - C'_i\| + a_4 \|c_i - c'_i\| \leq 1, \end{aligned} \quad (7.17)$$

$$\text{with variables } (z^*, z'^*, H_i, H'_i, h_i, h'_i, C_i, C'_i, c_i, c'_i, w_i, w'_i, \mu),$$

where the set $\text{KKT}(H_i, h_i, C_i, c_i, w_i, \mu)$ is defined as

$$\begin{aligned} \text{KKT}(H_i, h_i, C_i, c_i, w_i, \mu) := \{z \mid & H_i z + h_i - \mu + C_i^T w_i = 0, \\ & C_i z_i \leq c_i, \\ & w_i \geq 0, \\ & w_{i,j} (C_i z - c_i)_j = 0 \text{ for all } j\}, \end{aligned}$$

where w_i represents Lagrangian multipliers for the optimization problem (7.13). Notice that we used the distance $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i)$ as defined for the quadratic problem in (7.9).

We state our first observation regarding the sensitivity of quadratic problems as formulated in Problem 7.2.2.

Lemma 7.3.1 (Sensitivity for Problem 7.2.2). *If we restrict Problem 7.2.1 to be of quadratic form, as defined in Problem 7.2.2, with the adjacency relation $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i)$ defined in (7.9), then the sensitivity Θ_i can be simplified as*

$$\Theta_i := \max_{\text{adj}(\mathcal{P}_i, \mathcal{P}'_i) \leq 1} \|g(\mathcal{P}_i, \mu = 0) - g(\mathcal{P}'_i, \mu = 0)\|, \quad (7.18)$$

that is we lose the explicit dependency on lagrange variable μ .

Furthermore, if additional assumptions are made on the problem statement, an upper bound on Θ_i can be given in closed-form expressions.

Lemma 7.3.2 (Special Case - Protecting H_i in Problem 7.2.2). *Assume $f_i(z_i) := \frac{1}{2} z_i^T H_i x + h_i^T z_i$ and the distance between two problems is defined as $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i) = \|H_i - H'_i\|_2$. That is, the privacy requirement only concerns about the matrix H_i . Also assume that the local variable z_i is bounded as $\|z_i^k\| \leq G_i$, then*

$$\Theta_i \leq \frac{G_i}{\lambda_{\min}^{(i)}}, \quad (7.19)$$

where $\lambda_{\min}^{(i)}$ is the lower bound on the eigenvalues defined in Problem 7.2.2.

Lemma 7.3.3 (Special Case - Protecting h_i in Problem 7.2.2). *Assume $f_i(z_i) := \frac{1}{2}z_i^T H_i x + h_i^T z_i$ and the distance between two problems is defined as $\text{adj}(\mathcal{P}_i, \mathcal{P}'_i) = \|h_i - h'_i\|_2$. That is, the privacy requirement only concerns about the vector h_i . Then*

$$\Theta_i \leq \frac{1}{\lambda_{\min}^{(i)}},$$

where $\lambda_{\min}^{(i)}$ is the lower bound on the eigenvalues defined in Problem 7.2.2.

Proofs of the lemmas are given in Appendix B. These lemmas give closed-form expressions of the upper bounds on Θ_i for two special cases, and they will be useful for our applications in Section 7.4. Notice however that the upperbounds do not serve as a straightforward design principle for scaling eigenvalues to lower sensitivity. As a result of scaling the eigenvalues, the distance metric $\text{adj}(\cdot, \cdot)$ will also change; higher eigenvalues generally result in smaller sensitivity but also a narrower and less meaningful “region of adjacency”, as discussed in Section 7.3, Equation (7.9).

In the rest of this section, we discuss how to estimate the sensitivity for generic quadratic problems using a sample-based method. We point out that a more general discussion of this approach can be found in [30]. For the sake of notation, we write Θ_i in (7.14) as $\Theta_i(\mathcal{P})$ where \mathcal{P} denotes the collection of all variables $(z^*, z'^*, H_i, H'_i, \tilde{h}_i, \tilde{h}'_i, C_i, C'_i, c_i, c'_i, w_i, w'_i)$. Furthermore we use \mathbb{C}_Θ to denote the polynomial constraints given by (7.15), (7.16) and (7.17). With these notations, the optimization problem in (7.14) can be rewritten as

$$\begin{aligned} \min \quad & \gamma, \\ \text{s.t.} \quad & \Theta(\mathcal{P}) - \gamma \leq 0, \quad \forall \mathcal{P} \in \mathbb{C}_\Theta. \end{aligned} \tag{7.20}$$

The idea of sample-based approach is to randomly draw many instances of \mathcal{P} from the set \mathbb{C}_Θ , and find the maximum $\Theta(\mathcal{P})$ using these samples. Namely, we solve the following problem

$$\begin{aligned} \gamma^N := \min \quad & \gamma, \\ \text{s.t.} \quad & \Theta(\mathcal{P}_s) - \gamma \leq 0, \quad \forall \mathcal{P}_s, \quad s = 1, \dots, N, \end{aligned} \tag{7.21}$$

where $\mathcal{P}_1, \dots, \mathcal{P}_N$ are N randomly drawn samples from \mathbb{C}_Θ . More specifically for our problem, we randomly sample the parameters $H_i, H'_i, \tilde{h}_i, \tilde{h}'_i, C_i, C'_i, c_i, c'_i, w_i, w'_i$ from their sets. For each sampled set of parameters, we solve the original optimization problem (7.3) to obtain z^* and z'^* , hence obtain one estimate $\Theta(\mathcal{P}_s) := \|z^* - z'^*\|$. After N samples, the maximal $\Theta(\mathcal{P}_s)$ is set to be γ^N , which gives a lower bound on the sensitivity. To quantify the quality of this approximation, the following definition is introduced.

Definition 7.3.2 (Random Sampling, Definition 1, 2 in [31]). *Let γ^* denote the optimal solution to problem (7.20) and γ^N be a candidate solution retrieved through solving (7.21). We say γ^N is an α -level robustly feasible solution, if $V(\gamma^N) \leq \alpha$, where $V(\gamma^N)$ is defined as*

$$V(\gamma^N) := \mathbb{P}\{\mathcal{P} \in \mathbb{C}_\Theta : \gamma^N \leq \gamma^*\}. \tag{7.22}$$

In other words, $V(\gamma^N)$ is the portion of \mathcal{P} in \mathbb{C}_Θ that was not explored after N samples, which, if explored, would yield a higher function value than γ^N and thereby a tighter lower bound on sensitivity $\Theta(\mathcal{P})$. The following result relates the number of samples N to the quality of the approximation.

Lemma 7.3.4 (Sampling Rule, Corollary 1 in [31]). *For a given $\alpha \in [0, 1]$ and $\beta \in [0, 1]$ and let $N \geq \frac{1}{\alpha\beta} - 1$. Then with probability no smaller than $1 - \beta$, the solution γ^N given by (7.21) is a solution with α -level robust feasibility for Problem (7.20).*

This result gives the minimum number of samples, with which the sample-based approach will with high probability (larger than $1 - \beta$) find an approximate optimal solution to our problem.

Convergence properties of the distributed optimization algorithm

Privacy comes with a price. The noise term in Algorithm 1 makes the convergence rate slower than the case without noise. However, it is possible to prove that even with random noise, the algorithm converges in expectation. To show this, we first write out the dual problem of Problem 7.2.1 as follows.

Problem 7.3.2 (Dual Problem of Problem 7.2.1).

$$\min \quad -D(w) = \underbrace{\sum_{i=1}^M f_i^*(w_i)}_{\phi(w)} + \underbrace{I_{\{E^T w=0\}}(w)}_{\psi(w)},$$

with the matrix $E := [E_1^T, E_2^T, \dots, E_M^T]^T$ and dual variable $w := [w_1^T, \dots, w_M^T]^T$. We use f_i^* to denote the conjugate function of $f_i : \mathbb{C}_i \rightarrow \mathbb{R}$, defined as $f_i^*(w) = \sup_{z \in \mathbb{C}_i} (w^T z - f_i(z))$ and $I_{\mathbb{S}}$ denotes the indicator function on a set \mathbb{S}

$$I_{\mathbb{S}}(x) = \begin{cases} 0 & \text{if } x \in \mathbb{S} \\ \infty & \text{if } x \notin \mathbb{S} \end{cases}.$$

The stochastic proximal-gradient method (stochastic PGM), as given in Algorithm 2, is a method to solve the dual problem above. The proximity operator [160] in Algorithm 2 is defined as

$$\text{prox}_{\tau\psi}(v) := \operatorname{argmin}_x (\psi(x) + (1/2\tau\|x - v\|^2))$$

This algorithm has been studied extensively, see for instance [166, 184]. It addresses optimization problems of the form given in Problem 7.3.3 and with assumptions in Assumption 7.3.4.

Problem 7.3.3 (Dual Problem Form).

$$\min_{w \in \mathbb{W}} \quad \Phi(w) = \phi(w) + \psi(w) \quad .$$

Assumption 7.3.4 (Assumptions for Problem 7.3.3).

- ϕ is a strongly convex function with a convexity modulus $\rho_\phi > 0$, and has a Lipschitz continuous gradient with a Lipschitz constant $L(\nabla\phi) > 0$.
- ψ is a lower semi-continuous convex function, not necessarily smooth.
- The norm of the gradient of the function ϕ is bounded, i.e., $\|\nabla\phi(w)\|^2 \leq B^2$ for all $w \in \mathbb{W}$.
- The variance of the noise e^k is equal to σ^2 , i.e., $\mathbb{E}[\|e^k\|^2] \leq \sigma^2$ for all k .

Algorithm 2 Stochastic Proximal-Gradient Method

Require: Require $w^0 \in \mathbb{W}$ and step size $\tau^k < \frac{1}{\rho_\phi k}$
for $k = 1, 2, \dots$ **do**
 1: $w^k = \text{prox}_{\tau^k \psi}(w^{k-1} - \tau^k(\nabla\phi(w^{k-1}) + e^k))$
end for

The key observation in our proof of convergence is the following lemma, showing that executing Algorithm 1 on the original problem 7.2.1 is equivalent to executing Algorithm 2 on the dual problem 7.3.2.

Lemma 7.3.5 (Equivalence of the Primal and Dual Problems). *Consider using Algorithm 1 on the primal problem, Problem 7.2.1, and using Algorithm 2 on the dual problem, Problem 7.3.2. Further assume that Algorithm 1 is initialized with the sequence μ_j^0, z_j^0 , for $j = 1, \dots, M$, and Algorithm 2 is initialized with the sequence w_j^0 where $w_j^0 = \mu_j^0$ for $j = 1, \dots, M$. Then $w_j^k = \mu_j^k$ for all $k = 1, 2, \dots$ and all $j = 1, \dots, M$, and the error terms in Algorithm 1 and Algorithm 2 have the relationship $e^k = \delta^k = [\delta_1^{kT}, \delta_2^{kT}, \dots, \delta_M^{kT}]^T$.*

Proof sketch: This proof is an extension of the proof for Lemma 3.4 in [164]. We first Problem 7.2.1 in the form of the splitting problem in Problem 3.1 in [164], by defining the two objectives as $f(z) = \sum_{i=1}^M f_i(z_i)$ subject to $z_i \in \mathbb{C}_i$ for all $i = 1, \dots, M$ and $g = 0$ with the optimization variables for the two objectives $z = [z_1^T, z_2^T, \dots, z_M^T]^T$ and v , respectively. The coupling matrices are set to $A = I$, $B = -E = -[E_1^T, E_2^T, \dots, E_M^T]^T$ and $c = 0$. Under the assumption that the first objective $f(z)$ consists of a strongly convex function on z and convex constraints. The convex constraints can be considered as indicator functions, which are convex functions. Due to the fact that the sum of a strongly convex and a convex function is strongly convex, the objective $f(z)$ is a strongly convex function. The second objective g is a convex function. Then, by using the results in Lemma 3.4 in [163], we can prove that Algorithm 1 is equivalent to Algorithm 2, executed on Problem 7.3.2, the dual problem of Problem 7.2.1. ■

Based on the equivalence shown in Lemma 7.3.5, we are ready to provide the following theorem showing the convergence properties of Algorithm 1.

Theorem 7.3.5 (Suboptimality). *Consider Algorithm 1. Assume that the local variables z_i^k are bounded as $\|z_i^k\|^2 \leq G_i^2$ for all $k \geq 0$ and for all $i = 1, \dots, M$. We have that for any $k > 1$,*

$$\mathcal{S}_k := \mathbb{E}[|D(w^k) - D(w^*)|] \leq \frac{4 \sum_{i=1}^M (G_i^2 + \sigma_i^2)}{\rho_\phi^2 k} . \quad (7.23)$$

Proof: Consider applying Algorithm 2 on Problem 7.3.3 with Assumption 7.3.4, we can first show that

$$\mathbb{E}[\|w^k - w^*\|] \leq \frac{4(B^2 + \sigma^2)}{\rho_\phi^2 k} . \quad (7.24)$$

The proof of this claim follows the same flow as the proof of Theorem 1 in [166] by noticing the following two facts: if the function f is convex, closed and proper, then

$$\|\text{prox}_f(x) - \text{prox}_f(y)\| \leq \|x - y\| , \quad (7.25)$$

and the variance of the gradient of ϕ is bounded by $\mathbb{E}[\|\nabla\phi(w^{k-1}) + e^k\|^2] \leq B^2 + \sigma^2$ for all $k > 1$.

We then apply this result to our Problem 7.3.2. The gradient of the first objective is equal to:

$$\nabla_w \phi(w^k) = \nabla_w \sum_{i=1}^M f_i^*(w_i^k) = [z_1^{kT}, \dots, z_M^{kT}]^T$$

with $w^k = [w_1^{kT}, \dots, w_M^{kT}]^T$. With our assumption in the theorem, the dual gradient is bounded as:

$$\|\nabla\phi(w^k)\|^2 \leq \sum_{i=1}^M \|z_i^k\|^2 \leq \sum_{i=1}^M G_i^2 .$$

Finally the claim follows because of the equivalence result in Lemma 7.3.5. ■

Theorem 7.23 shows that the sub-optimality gap for Algorithm 1 is bounded above by a function that is linear in the number of agents M and linear in the noise variances $\sigma_i, i = 1, \dots, M$. The convergence rate is of order $\mathcal{O}\left(\frac{1}{k}\right)$, which is a satisfying result, as compared to current results reported in literature (the motivating work by Han et. al achieved a convergence rate of $\mathcal{O}\left(\frac{1}{\sqrt{k}}\right)$ [97]).

Remark 7.3.6. *It is observed in [154] that a general class of differentially private distributed optimization algorithms cannot be asymptotically stable, if the noise is added to the inter-agent messages. This result is also applicable to the algorithm studied in this paper. However, in this work we do not focus on the asymptotic behavior of the algorithm, but characterize the fundamental quantities of interests for finite k as in (7.11) and (7.23). This is the most interesting for practical implementations, where the algorithm is always executed for a finite number of iterations.*

7.4 Application: Simplified Distributed Optimal Power Flow

This section presents a simplified optimal power flow (OPF) problem that inspires the proposed control approach. We consider the setting of a radial distribution feeder, and consider the flow of real power on its branches. We formulate the power flow model and the OPF objectives and develop the distributed OPF problem according to the quadratic problem, as defined in (7.3). We then discuss the parameters that are subject to privacy requirements and interpret the trade-offs developed in Section 7.3.

Simplified Optimal Power Flow

Solving the simplified OPF problem requires a model of the electric grid describing both topology and impedances. This information is represented as a graph $\mathcal{G} = (\mathbb{V}, \mathcal{E})$, with \mathbb{V} denoting the set of all buses (nodes) in the network, and \mathcal{E} the set of all branches (edges). For ease of presentation and without loss of generality, here we introduce part of the linearized power flow equations over *radial* networks, also known as the *LinDistFlow* equations [17]. In such a network topology, each bus j has one upstream parent bus $\{i \mid (i, j) \in \mathcal{E}\}$ and potentially multiple downstream child buses $\{k \mid (j, k) \in \mathcal{E}\}$. By \mathcal{D}_j we denote the set of all buses downstream of branch (i, j) . We assume losses in the network to be negligible and model the power flowing on a branch as the sum of the downstream net load:

$$P_{ij} \approx \sum_{k \in \mathcal{D}_j} \{p_k^c - p_k^g + u_k\} \quad (7.26)$$

In this model, capital P_{ij} represents real power flow on a branch from node i to node j for all branches $(i, j) \in \mathcal{E}$, lower case p_i^c is the real power consumption at node i , and p_i^g is its real power generation. This nodal consumption and generation is assumed to be uncontrollable. In addition, we consider controllable nodal injection u_i , available at a subset of nodes $i \in \mathcal{C} \subset \mathbb{V}$. In this case study, we aim to prevent overload of real power flow over certain critical branches in an electric network. This aim is formulated through constraints

$$\begin{aligned} \sum_{k \in \mathcal{D}_j} \{p_k^c - p_k^g + u_k\} - \overline{P}_{ij} &\leq 0, \\ \underline{P}_{ij} - \sum_{k \in \mathcal{D}_j} \{p_k^c - p_k^g + u_k\} &\leq 0, \forall (i, j) \in \mathcal{E}_{\text{safe}}, \end{aligned} \quad (7.27)$$

$\mathcal{E}_{\text{safe}} \subset \mathcal{E}$ denotes a subset of branches for which power flow limitations are defined, $\overline{P}_{ij}, \underline{P}_{ij}$ denoting the upper and lower power flow bounds on branch $(i, j) \in \mathcal{E}_{\text{safe}}$. Constraints are included to account for inverter capacity. Each node i is ultimately limited by the local capacity on total apparent power capacity \bar{s}_i . We consider a simple controller architectures that can deliver real power, which we assume are both given their own capacity

$$\underline{u}_i \leq u_i \leq \bar{u}_i, \forall i \in \mathcal{C}. \quad (7.28)$$

We consider a scenario in which the utility negotiates different prices for different capacities, potentially at different points in time, with different third party DER owners. Let u_i refer to the real power used for the optimization scheme from agent i , and π_i denotes the price for procuring a kWatt from agent i . The optimal power flow determines the control setpoints that minimizes an economic objective subject to operational constraints.

$$\begin{aligned} \min_{u_i, i \in \mathcal{C}} \quad & \sum_{i \in \mathcal{C}} \pi_i (u_i)^2, \\ \text{s.t.} \quad & (7.27), (7.28). \end{aligned} \tag{7.29}$$

The OPF problem (7.29) can be recast as an instance of the quadratic distributed optimization problem (7.3). First, note that the objective is quadratic in the optimization variables u_i , and separable per node. Second, for all nodes $i \in \mathbb{V}$, the capacity box constraints (7.28) are linear and fully local. The safety constraints (7.27) require communication to and computation by a central trusted node. To ensure strong convexity of the local problems, the economic cost objectives are shared between each agent i and the central trusted node. Hence, $\forall i \in \mathbb{V} \setminus \{0\}$, the objective reads

$$f_i(u_i) = \frac{\pi_i}{2} (u_i)^2, \tag{7.30}$$

with capacity constraint (7.28). The central node 0 has objective function

$$f_0(z_0) = \sum_{i \in \mathcal{C}} \frac{\pi_i}{2} (u_i)^2, \tag{7.31}$$

with safety constraints 7.27. As such, this distributed problem assumes a star-shaped communication structure, in which the a centrally trusted node receives all u_i, p_i^c, p_i^g from the agents. The agents retrieve iterates of u_i^p from the central node and compute a simple problem with only economic cost and a local capacity constraint.

Private Information in Distributed OPF

In the context of d-OPF, we consider assigning privacy requirements to two sets of parameters; the prices π_i that the DSO charges to different agents in the network, and the capacities $\underline{u}_i, \bar{u}_i$ available to all agents $i \in \mathcal{C}$. Together, these parameters provide important strategic insight into the commercial position of each agent. A DSO may charge different prices for different levels of commitment or for the varying value that the DSO gets from the actions of a specific agent at specific time periods or places in the network. In a natural commercial context, the DSO may have an interest to hide the prices to other agents. In addition, in a negotiation setting, a strategic agent may want to find out the capacity available by other agents in the network to adjust its bid to the DSO, so as to be the first or only agent to be considered, which could lead to asymmetric and potentially unfair bidding situations. As such, in order to give all agents with capacity a fair chance to participate in d-OPF, there is value in hiding the capacity (and price) parameters.

To formulate this in the language of local differential privacy, we need to define the distance metrics for all considered parameters. In the case of both prices and capacity, this is achieved by considering the maximum range in which these parameters can lie. The distance metric proposed is the ℓ_1 -norm. Given this metric, we need to define a proper adjacency relation, which determines the maximum change in a single parameter that we aim to hide with the differentially private algorithm.

Definition 7.4.1. (Adjacency Relation for Distributed OPF): For any databases $D = \{f_i(\pi_i), \mathbb{C}_i(\bar{u}_i, \underline{u}_i)\}$ and $D' = \{f'_i(\lambda'_i), \mathbb{C}'_i(\bar{u}'_i, \underline{u}'_i)\}$, we have $\text{adj}(D, D')$ if and only if there exists $i \in [M]$ such that

$$|\pi_i - \pi'_i| \leq \delta\pi, \quad |\bar{u}_i - \bar{u}'_i| \leq \delta\bar{u}, \quad |\underline{u}_i - \underline{u}'_i| \leq \delta\underline{u}, \quad (7.32)$$

and $\pi_j = \pi'_j, \bar{u}_j = \bar{u}'_j, \underline{u}_j = \underline{u}'_j$ for all $j \neq i$.

By setting $\delta\pi, \delta\bar{u}$ and $\delta\underline{u}$ respectively as the maximum price offered per unit of energy (i.e. $\bar{\pi}$ if $\pi_i \in [0, \bar{\pi}]$) and the maximum capacity in the network (i.e. $\arg \max_{i \in \mathcal{C}} \bar{u}_i$), we ensure that all parameters in the network are properly covered by the definition.

Interpreting Trade-offs

We analyze and interpret the theoretical results that illuminate an inherent trade-off between privacy level and suboptimality. Assuming a fixed noise variance across all iterations, Equations 7.11 and 7.23, we have the following trade-off:

$$\epsilon_i = \Theta_i \frac{K}{\sigma_i}, \quad \mathcal{S}_K \leq \frac{4 \sum_{i=1}^M (G_i^2 + \sigma_i^2)}{\rho_\phi^2 K} \quad (7.33)$$

Remember that better privacy relates to a lower privacy level ϵ_i and a more optimal solution relates to lower suboptimality. Unsurprisingly, an increasing number of iterations leads to worse privacy and better suboptimality. Conversely, a higher noise variance leads better privacy and worse suboptimality. Figure 7.2 shows the region of attainable (ϵ, \mathcal{S}) values for the simplified OPF problem. The left figure shows that, for a fixed reasonable level of privacy ($\epsilon \leq 10^0$), the sub-optimality will decrease for a larger number of iterations K . The right figure shows that for a fixed level of privacy, the sub-optimality bound tightens for higher variance levels. For a fixed level of sub-optimality, a higher noise variance σ achieves a lower (and hence better) privacy level. Ideally, parameters $(\{\sigma_i\}_{i \in \mathcal{C}}, K)$ are chosen along the Pareto front of this graph. As a result, a system designer may want to define specifications,

$$\epsilon_i \leq \bar{\epsilon}, \quad \mathcal{S} \leq \bar{\mathcal{S}}. \quad (7.34)$$

Based on the specifications, we then want to determine feasible values for the number of iterations K and noise variance σ_1 . For the sake of analysis, we let all upper bounds be the same in the second equation, $G_i = G, \forall i \in \mathcal{C}$. In addition, we consider the normalized




Figure 7.2: Achievable tradeoffs between privacy level ϵ and suboptimality \mathcal{S} , with Pareto front. (left) indicates increasing number of iterations and (right) increasing noise variance. We assume $\sigma_i = \sigma, G_i = G, \forall i \in \mathcal{C}$.

noise-to-signal ratio $\nu_i := \frac{\sigma_i}{G}$, which is more intuitive as a tunable parameter. With these steps, we can write

$$\frac{K}{\nu_i} \leq \frac{\bar{\epsilon} G_i}{\Theta_i}, \quad \frac{1 + \nu_i^2}{K} \leq \frac{\bar{\mathcal{S}} \rho_\phi^2}{4MG^2} \quad (7.35)$$

We now make these relationships concrete for the case proposed in Lemma 7.3.2. Using (7.19),




Figure 7.3: Feasible parameter sets (ν, K) for varying levels of $\bar{\epsilon}$ (left) and $\bar{\mathcal{S}}$ (right), assuming $\sigma_i = \sigma, G_i = G, \forall i \in \mathcal{C}$.

we can write

$$\Theta_i \frac{K}{\nu_i} \leq \frac{GK}{\lambda_{\min}^{(i)} \nu}. \quad (7.36)$$

By enforcing the upper-bound specifications, we get

$$\frac{K}{\nu_i} \leq \lambda_{\min}^{(i)} \bar{\epsilon}, \quad \frac{1 + \nu_i^2}{K} \leq \frac{\bar{\mathcal{S}} \rho_\phi^2}{4MG^2} \quad (7.37)$$

These equations tell us two things. Firstly, we see that the smallest eigenvalue of H_i may have a tightening or alleviating effect on the privacy tradeoff equation. Second, the suboptimality

is mostly governed by K , as typically $0 \leq \nu_i^2 \ll 1$ (in other words, we need σ_i to be on the order of G or larger to affect suboptimality). These equations provide a specific test to determine feasible parameters (ν, K) that satisfy a set of specifications $(\bar{\epsilon}, \bar{\mathcal{S}}, M, G, \rho_\phi)$. Note that this set may be empty if the specification are too stringent. Figure 7.3 shows the feasible set for varying levels of specifications $(\bar{\epsilon}, \bar{\mathcal{S}})$.

We now further specify the tradeoff relations in (7.37) for the simplified OPF problem, maintaining the assumption that $G_i = G, \forall i \in \mathcal{C}$. Note that $\rho_\phi = \max_i \pi_i := \pi_{\max}$, and $G = \max(|\underline{u}_i|, |\bar{u}_i|) := u_{\max}$. This yields

$$\frac{K}{\nu_i} \leq \frac{\pi_i \bar{\epsilon}_i}{2}, \quad \frac{1 + \nu_i^2}{K} \leq \frac{\bar{\mathcal{S}} \pi_{\max}^2}{4M u_{\max}^2} \quad (7.38)$$

The first equation shows that the ratio of the number of iterations to the normalized noise needs to be sufficiently small, capped by the specified privacy level $\bar{\epsilon}_i$ and the agent's price π_i . The latter equation shows that with increasing number of agents M injecting noise, we need more iterations or less noise per agent to achieve the same suboptimality. Similarly, if the maximum capacity u_{\max} of the agents, and hence the upper bound G of the optimization variables increases or the maximum price π_{\max} decreases, we require more iterations or lower noise variance to maintain the same level of suboptimality.

7.5 Sharing or Pricing the Privacy Budget?

Theorem 7.3.5 provides a relationship between sub-optimality and the cumulative variance of the Laplacian noise inserted by all subsystems. In real scenarios, a system designer or operator may specify a desired level of (sub-)optimality \mathcal{S}_K achieved after K iterations, that is $\mathbb{E}[D(\lambda^K) - D(\lambda^*)] \leq \mathcal{S}_K$. Rewriting Equation (7.23), we can compute a bound on the amount of cumulative Laplacian noise allowed at run-time

$$\sum_{i=1}^M \sigma_i^2 \leq \frac{1}{4} \rho_\phi^2 K \mathcal{S}_K - \sum_{i=1}^M G_i^2 \triangleq \Sigma_{\text{budget}} \quad (7.39)$$

Hence, once \mathcal{S}_K is specified, a *cumulative privacy budget* Σ_{budget} is set. Remember that each individual agent may define a different distance metric, or different weights, which may lead to different sensitivities and hence different noise levels required by the various agents. Since the privacy budget is limited, a fair and transparent allocation procedure is required to divide the allowable noise over all agents. Here, we propose two approaches for going about such allocation.

The first approach, would entail a proportional allocation. This could be done in two ways. The first way involves splitting the noise variance budget by the number of agents: $\sigma_i^2 = \frac{1}{M} \Sigma_{\text{budget}}, \forall i = 1, \dots, M$. In this case, the allowed noise variance σ_i^2 determines the maximum level of local differential privacy ϵ_i that can be maintained, given a set sensitivity Θ_i (or vice versa), as outlined in (7.11)), which will differ from agent to agent. The second way involves

setting all local differential privacy levels ϵ_i equal, and, given set sensitivities Θ_i , splitting the noise among all agents. This is equivalent to equating Equation (7.11) for all agents $i = 1, \dots, M$, leading to M equations for $\sigma_1, \dots, \sigma_M$:

$$\frac{\Theta_1}{\sigma_1} = \frac{\Theta_2}{\sigma_2} = \dots = \frac{\Theta_M}{\sigma_M}, \quad \sum_{i=1}^M \sigma_i^2 = \Sigma_{\text{budget}}, \quad (7.40)$$

where we assume that the noise variances are constant for all time steps $k = 1, \dots, K$.

The second approach we anticipate is a pricing scheme, in which the value of privacy is left to a market or negotiation. In the context of d-OPF, it is natural to assume that different DER owners with varying privacy levels will have varying degrees of willingness to pay or incur a deduction on their revenue for preserving privacy of their local parameters. Here, we propose two scenarios to perform allocation via the so-called *Kelly mechanism* [121]. We assume a one-directional bid w_i done by all agents after seeing a price π_i^σ given by the network operator to the agent. With the bid, the operator constructs a surrogate utility function $w_i \log \sigma_i^2$. The operator then determines the allocation and payment by maximizing the sum of surrogate utility functions as:

$$\sigma_{\text{Kelly}}^2(w) = \arg \max_{\sum_{i=1}^M \sigma_i^2 \leq \Sigma_{\text{budget}}} \sum_{i=1}^M w_i \log \sigma_i^2. \quad (7.41)$$

The authors of [225] show that this mechanism works neatly if one assumes that buyers are *price-takers*. In the case buyers are strategic, realizing that the eventual price is influenced by all bids, i.e. $\pi_i^\sigma(w)$, the Kelly mechanism may not yield an efficient Nash equilibrium. To account for such behavior, the same authors combined the Kelly mechanism with the celebrated Vickrey-Clark-Groves (VCG) mechanism [214, 40, 95], proposing the VCG-Kelly mechanism. The VCG mechanism is lauded because it incentivizes all participating buyers to report truthfully about the value they assign to the goods sold. It has an intuitive payment scheme, which says that each player i should pay the difference in utility of the other players $j \neq i$ between the scenarios that player i does and does not participate, which leads to each player caring about both its own utility and that of others, leading to truthful reporting. The VCG-Kelly mechanism extends the VCG mechanism to problems with divisible goods [225]. It considers a more general class of surrogate utility functions $V_i(w_i, \sigma_i^2) = w_i f_i(\sigma_i^2)$, with f_i 's strictly increasing, strictly concave, and twice differentiable. The allocation rule is then similar in form:

$$\sigma_{\text{VCGK}}^2(w) = \arg \max_{\sum_{i=1}^M \sigma_i^2 \leq \Sigma_{\text{budget}}} \sum_{i=1}^n V_i(w_i, \sigma_i^2), \quad (7.42)$$

and the payment scheme reads:

$$m_{i,\text{VCGK}}(w) = \left(\max_{\sum_{i=1}^M \sigma_i^2 \leq \Sigma_{\text{budget}}, \sigma_i^2 = 0} \sum_{j=1, j \neq i}^M V_j(w_j, \sigma_j^2) \right) - \sum_{j=1, j \neq i}^M V_j(w_j, \sigma_{i,\text{VCGK}}^2). \quad (7.43)$$

The actual form of the utility functions in the context of our application is left as an open problem, which requires further investigation. Readers interested in further details and examples of the VCG-Kelly mechanism are directed to [225].

We propose these allocation schemes to trigger a discussion about what a fair and workable division of the allowable noise variance in our privacy-preserving looks like. This may vary based on the application. Given that privacy in engineered systems is a value of increasing importance, it may be wise to consider the impact of pricing privacy for participants with varying socioeconomic background. In situations where certain agents have significantly less resources but privacy is equally important, a proportional scheme may be the more ethical approach to take. If pricing is still used and it is anticipated that some participants outbid others, one may think about allowing each participant to have a minimum amount of privacy, which can be translated into an extra linear constraint in the Kelly mechanisms.

7.6 Conclusions and Future Work

In this chapter, we developed local ϵ -differential privacy for distributed optimization, building on recent advances in inexact alternating minimization algorithm (IAMA). By exploiting the IAMA's convergence properties under the existence of errors in communication and computation, we are able to add noise to agent-to-agent communication in a way that preserves privacy in the specifications of user objectives and constraints while still guaranteeing convergence. We achieve a convergence rate of order $\mathcal{O}(\frac{1}{k})$, which compares favorably to state-of-the-art algorithms. In addition, our method allows the protection of private parameters in both constraints and objective functions, and facilitates specifying customized privacy levels. We analyzed the trade-offs between privacy and suboptimality for various levels of noise and number of iterations, and gave a method to determine feasible values of noise variance and number of iterations given specifications on privacy and suboptimality. We propose different alternatives to allocate the allowable noise variance across participating agents, either via proportional sharing or market mechanisms that incentivize for truthful reporting and allow efficient Nash equilibria, which can be used to implement the algorithm in fair and efficient ways.

Chapter 8

The Epistemology and Dynamics of Automated Decision-Making

“Everything that we see is a shadow cast by that which we do not see.”

- Martin Luther King, Jr.

Submitted to the International Conference on Machine Learning [60].

Machine learning (ML) is increasingly deployed in real world contexts, supplying “actionable insights” and forming the basis of automated decision-making systems. These systems can be affected by *pre-existing bias* in training data, but also *technical and emerging biases* which arise as context-specific artifacts of system implementation. For the data-driven methodologies presented in this dissertation, we analyzed and pointed out different ways in which bias and error may affect their performance and lead to violation of relevant values such as privacy breaches or new safety risks. In this chapter, we take a more general view at the role of bias in automated decision-making. We reframe technical bias as an epistemological problem and emerging bias as a dynamical feedback phenomenon, call for changing machine learning practice, and encourage reflection on how *our positionality shapes our epistemology*.

8.1 Introduction

Data-driven decision-making has emerged in high-stakes social domains such as medical clinics, criminal justice, and public transportation. Motivated by the promise of new capabilities, quality improvements, or efficiency gains, learning-based technologies have captured the attention and imagination of many domain experts. At the same time, dilemmas have emerged surrounding the bias that creeps into machine learning systems, leading to erroneous decision-making, disparate treatment or outcomes [18], and representational harm [44].

While many technical tools are being proposed to mitigate these issues [99, 109], there is insufficient understanding of *how the machine learning design and deployment process, and its practitioners* can safeguard critical human values such as safety, fairness, and freedom of expression. The recently launched *AI Now Institute* stresses that no single technical tool will cleanly solve these problems and identifies “a deep need for interdisciplinary, socially aware work that integrates the long history of bias research from the social sciences and humanities into the field of AI research” [35].

But how can ML practitioners, who lack a consistent language beyond technical descriptions and solution spaces for well-defined problems, start to grapple with these challenges? And how can we do so in a way that embraces constructive dialogue rather than dismissive critique? Until the machine learning field embraces a dynamic perspective of ML-based decision-making and reflects on the discipline’s epistemology, we will not be able to contend effectively with issues such as fairness, accountability, and transparency.

8.2 A Broader View On Bias

While often championed as a way to overcome human biases in decision-making, there are many ways that machine learning systems can fail to achieve this goal, potentially causing systematic and unfair discrimination against certain individuals or groups of individuals. Most literature addressing issues of fairness has focused on the ways in which ML models can inherit *pre-existing biases* in its training data. Limiting ourselves to these biases is problematic in two ways.

Firstly, it narrows us to look at how these biases lead to *allocative harm*; a primarily economic view of how systems allocate or withhold an opportunity or resource, such as being granted a loan or sent back to prison. In her NIPS 2017 keynote, Kate Crawford made the case that at the root of all forms of allocative harm are biases that cause *representational harm*. This perspective requires us to move beyond biases in the data set and “think about the role of ML in harmful representations of human identity,” and how these biases “reinforce the subordination of groups along the lines of identity” and “affect how groups or individuals are understood socially,” thereby also contributing to harmful attitudes and cultural beliefs in the longer term [44]. It is fair to say that representation issues have been largely neglected by the ML community, mainly because they are hard to formalize and track. *Responsible representation* requires analyses beyond scrutinizing a training set, for instance by questioning how sensitive attributes might be represented by different features and classes of models.

Secondly, while ML systems are increasingly implemented to provide “actionable insights” and guide decisions in the real world, the core methods still fail to effectively address the inherent *dynamic nature* of interactions between the automated decision making process and the environment or individuals that are acted upon, especially in contexts where observations or human responses (such as clicks and likes) are *fed back* along the way to update the algorithm’s parameters, allowing biases to be further reinforced and amplified.

The tendency of ML-based decision-making systems to formalize and reinforce socially sensitive phenomena necessitates a broader taxonomy of biases that includes risks beyond those pre-existing in the data. As argued in 1996 by Friedman and Nissenbaum, two other sources of bias are equally likely to occur when designing and employing any computer systems that are used to mediate decisions, namely *technical bias* and *emerging bias* [82, 84].

While understanding pre-existing bias has lent itself reasonably well to statistical approaches for understanding a given data set, technical and emerging bias require engaging with the domain of application and the ways in which the algorithm is used and integrated. For automated decision-making tools to be responsibly integrated in any context, it is critical that designers (1) assess technical bias by reflecting on their *epistemology* and that of users and stakeholders to consider how a new tool may affect the ways critical responsibilities are carried out, and (2) assess emerging bias by studying the *feedback mechanisms* that create intimate coupling between evolving algorithms and the environment they act upon.

In what follows, we further reflect on these challenges. We aim to spark a constructive and informed debate about how ethics may enter the process of machine learning design and work towards what constitutes a *value-sensitive design* approach for data-driven decision-making.

8.3 Technical Bias Is About Epistemology

Friedman describes a source of technical bias as “the attempt to make human constructs amenable to computers - when, for example, we quantify the qualitative, make discrete the continuous, or formalize the nonformal” [82]. This is where the capacity of machine learning to make predictions (often better than human beings can) based on patterns in high-dimensional data clashes with the complex reality of many sociotechnical systems. On the one hand, machine learning provides a window to phenomena that we would otherwise not be able to perceive, let alone experiment with. On the other hand, the black-box nature of complex machine learning models, with its reliance on reduction and its tendency to encode correlation rather than cause-effect relationships, may violate common knowledge or produce unethical or catastrophic errors if interpreted incorrectly. While technical bias is domain-specific, to facilitate a broad discussion, we discuss four general categories as they relate to the more general process of modeling.

Firstly, existing or collected data X is at some point measured and transformed into a computer readable scale. Depending on the objects measured, each variable may have a different scale, such as nominal, ordinal, interval or ratio. Consider for example Netflix’s decision to let viewers rate movies with “likes” instead of a 1-5 star rating. As such, movie rating moved from an ordinal scale (a number score in which order matters, but the interval between scores does not) to a nominal scale (mutually exclusive labels: you like a movie or you don’t). Arguably, the nominal scale makes it easier for viewers to rate and provide data to find preference patterns across the customer base. However, note that it also affects how viewers are *represented* and what content gets recommended by the ML system. As

such, measurement matters and can produce *measurement bias*, so careful consideration is necessary to understand how this may affect a system's outcomes [98].

Secondly, based on the data gathered X and available domain knowledge, we may *engineer features* and *select model classes*. Features $\varphi(X)$ can be the available data attributes, transformations thereof based on knowledge and hypotheses, or generated/discovered in an automated fashion. A feature may also be a *proxy* representing a relevant phenomenon that cannot be measured directly. Since each feature can be regarded as a model of the system or population under study, it is relevant to ask how representative it is and why it may be predictive of the outcome. A model class $f(\cdot; \theta)$, with parameters θ , should be selected mostly based on how complex the phenomenon that we are representing or trying to predict is, and on the amount and quality of the available data. Is the individual or object that is subject to the decision easily reduced to numbers or equations to begin with? What information in the data x is inherently lost by virtue of the mapping $f(\varphi(x); \theta)$ having a limited complexity, regardless of the best fit θ^* ? In other words, ML can be seen as a *compression* problem in which a complex phenomenon is stored as a pattern in a finite-dimensional parameter space, meaning part of the information about the modeled phenomenon is inherently lost. The process of representation, abstraction and compression can be collectively described as inducing *modeling bias*. The challenges described here can be seen as a more complete cross-disciplinary version of the technical bias-variance trade-off [86]. From an information theoretic perspective, modeling bias influences the extent to which distortion can be minimized when *reconstructing* a phenomenon from a compressed or sampled version of the original [43, 54].

Thirdly, label data Y is used to represent the output of the model. Training labels may be the actual outcome for historical cases or some discretized or proxy version in cases where the actual outcomes cannot be measured or exactly quantified. Consider for example the use of records of arrest to predict crime rather than the facts of whether the crime was actually committed, which is often not possible to know with full certainty. How representative are such records of real crime across all subpopulations? What core information do they miss for representing the intended classes? And what bias lies hidden in them? We propose to refer to such issues as *label or output bias*.

Lastly, given a certain parameterization $(\varphi(\cdot), f(\cdot, \theta))$ and training data (X, Y) , a model is trained and *tuned to optimize certain objectives*. At this stage, various metrics may inform the model builder on where to tweak the model. Do we minimize the number of false positives or false negatives? In recidivism prediction, a false positive may be someone who gets sent to jail with no need, a false negative a dangerous person let free. It was shown recently that there are inherent trade-offs between prioritizing for equal prediction accuracy across groups versus for an equal likelihood of false positives and negatives across groups [123]. In addressing fairness, the many different metrics have motivated many definitions [152], illuminating the inherent ambiguity and context-dependence of such issues. This yields the question of what the right balance is for a given context, and who gets to decide. We coin the effects of striking such trade-offs *optimization bias*.

The many questions posed above are meant to illuminate the range of places in the

machine learning design process where issues of *epistemology* arise. These questions seek *justification* and often require *value judgment*. How do we represent phenomena in ways that are deemed correct? What evidence is needed in order to justify an action or decision? What are legitimate outcomes of an intervention? And how do we deal with inherent trade-offs of fairness? These challenges are deeply context-specific, often ethical, and challenge us to understand our epistemology and that of the domain we are working in.

The detrimental effects of overlooking these questions in practice are obvious in high-stakes domains, such as predictive policing and sentencing, where the decision to treat crime as a prediction problem reduces the perceived autonomy of individuals, fated to either commit or not commit a crime. This problem was convincingly laid out by Barabas et al., who argue that rather than prediction, “machine learning models should be used to surface covariates that are fed into a *causal model* for understanding the social, structural and psychological drivers of crime” [13]. It is a strong message with many challenges, but it points in the right direction: in these contexts, the domain expert should really be in the driver seat, with the machine learning work forming a constitutive rather than replacing role in hypothesizing and testing data-driven models that reflect rigorous causal relationships. It forces ML practitioners to be humble and reflect on how our skills and tools may benefit *or* hurt an existing decision-making process.

8.4 Emerging Bias Is About Dynamics

Recently, convincing examples of emerging bias have surfaced in technologies that we use in our daily lives. Consider, for instance, the way in which divisive *filter bubbles* or *echo chambers* on social media emerge from algorithms trained on users’ clicking behaviors, or how recommended content on Youtube has a tendency to turn towards the extreme and radical [206]. Rightfully so, in situations where a “machine learning system is unleashed in feedback with humans, that system is a reinforcement learning system, not a machine learning system” [169]. And reinforcement learning to determine decisions or actions is a *feedback control* problem, for which the effect of feedback can affect the stability of the *overall system* as bias accrues over time. In the context of predictive policing, where discovered crime data (e.g., arrest records) are used to predict the location of new crimes and determine police deployment, runaway feedback loops may cause increasing surveillance of particular neighborhoods regardless of the true crime rate [70], and has led to over-policing of “high-risk” individuals [196].

To acknowledge emerging bias, we adopt a conceptual dynamical model describing the feedback between a machine learning based decision system and the environment it is acting on (Figure 1). From the environment, the decision maker will consider observations, historical data, measurements, and responses. The machine learning system acts on the environment through decisions, control actions, or interventions. For example, in the case of predictive policing, “the environment” describes a city, and “the decision maker” is the police department, which determines where to send police patrols or invest in social interventions.

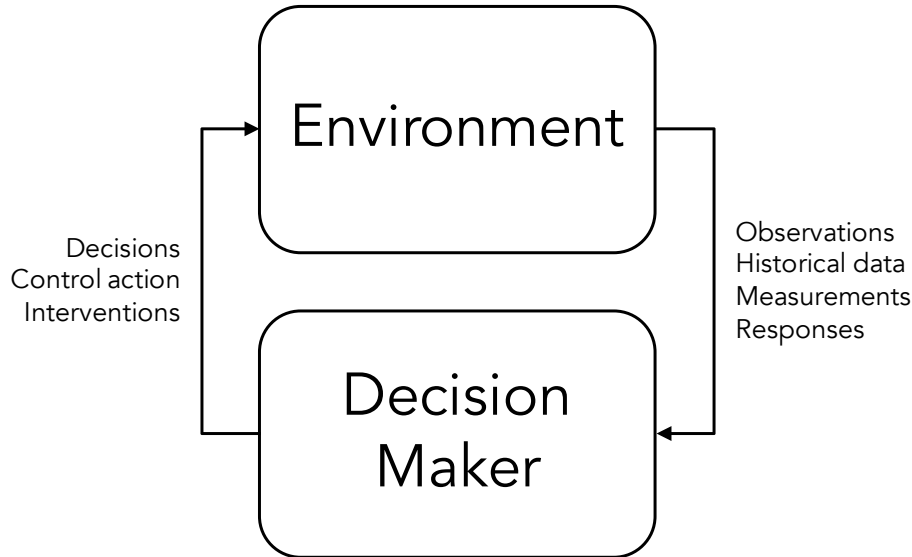


Figure 8.1: A Simple Feedback Model

The dynamical perspective offered by the conception of a feedback model allows for a focus on interactions, which can add clarity to debates over key issues like fairness and interpretability. Without a dynamic view of biases, it can be difficult to judge the fairness of the resulting decisions. To provide a convincing example, Hardt et al. showed how *static* observational metrics (properties of the joint distribution of input, model and output) are not able to distinguish between two scenarios with identical joint distributions but completely different interpretations for fairness [99]. This finding further motivates the case for causal reasoning, and also proves the necessity of analyzing the *dynamics of bias and fairness* in their particular context. For instance, a one-step feedback model, incorporating temporal indicators of well-being for individuals affected by decisions, offers a way of comparing competing definitions of fairness [133]. Similarly, calls for “interpretability” and proposed solutions often omit key operative words – Interpretability to whom? And for what purpose? [124]. The dynamic viewpoint adds clarity to these questions by focusing on causes and effects of decision making systems, and situating interpretability in the appropriate context.

Beyond providing a more realistic and workable frame of thinking about bias and related issues, the feedback system perspective may also allow inspiration to be drawn from areas of *Systems Theory* that have traditionally studied feedback and dynamics. For instance, from the field of *System Identification*, which uses statistical methods to build mathematical models of dynamical systems from measured data [134], often to be employed to control *safety-critical dynamical systems* such as airplanes or electric power generators with strict safety requirements [10]. Inspiration may be drawn from the rich literature on *closed-loop identification*, which considers the identification of models with data gathered from experiments in which the same model is continuously used to safeguard a system [209].

8.5 Our Positionality Shapes Our Epistemology

As ML experts, we are wired to analyze challenges in ways that *abstract, formalize and reduce complexity*. In our daily practice, we are concerned with model selection, over- and under-fitting, exploration/excitation trade-offs, and parameter estimation. The proliferation of biases in the systems we design understandably triggers us to think rigorously about its technical roots and what techno-fixes may be able to prevent negative impact. However, it is of crucial importance to acknowledge that reduction, formalization, and sources of feedback are *themselves* inherent sources of bias. These epistemological challenges differ tremendously from application to application and ultimately shape the way a user or decision-maker justifies decisions and affects individuals. Our technology *intimately touches and embodies values* deemed critical in employing the intended decision-making system. As such, we need to go beyond our formal tools and analyses to engage with others and reflect on our own epistemology in order to determine *responsible ways* in which technology can help put values into practice, and what its fundamental limits are.

With a plethora of issues surfacing, it is easy to either criticize ML all together, or otherwise dismiss requests to fundamentally revisit its role in enabling data-driven decision-making in sensitive environments [13]. Instead of either extreme, we propose three principles to nourish debate on the middle ground:

1. Do fairness forensics [44], by keeping track of biases in an open and transparent way and engaging in constructive dialogue with domain experts, to understand proven ways of formalizing complex phenomena and to breed awareness about how bias works and when/where users should be cautious.
2. Determine what values are relevant in building a decision-making system and how they might be embodied or challenged in the design of a system by engaging with users and beneficiaries [210, 85].
3. Acknowledge that your positionality shapes your epistemology [201]. Our personal backgrounds, the training we received, the people we represent or interact with all have an impact on how we look at problems. As ML practitioners, we should set aside some time and energy for critical self-reflection and harbor communication with the groups affected by the systems we design.

As Takacs describes it, the benefits of self-reflection go well beyond arriving at the “best solution” to a complex problem [201, 202]. “This means learning to listen with open minds and hearts, learning to respect different ways of knowing the world borne of different identities and experiences, and learning to examine and re-examine one’s own worldviews. [...] When we constantly engage to understand how our positionality biases our epistemology, we greet the world with respect, interact with others to explore and cherish their differences, and live life with a fuller sense of self as part of a web of community.” Martin Heidegger argued in his essay “The Question Concerning Technology” that all tools simultaneously *reveal* new insights

about ourselves and the world while also *concealing* old traditions and potential futures [102]. As machine learning systems rapidly change our own worldviews in ways we cannot fully anticipate, self-reflection and awareness of our epistemology may become ever-more important for machine learning practitioners.

Chapter 9

Conclusions

“Imagination is more important than knowledge. For knowledge is limited, whereas imagination embraces the entire world, stimulating progress, giving birth to evolution.”

- Albert Einstein

We are working towards a world in which we have an increasingly better handle on how data-driven technologies can be integrated in existing and new critical infrastructures while safeguarding and representing important values such as safety, privacy and social justice. We have the technical rigor to understand and mitigate new technical vulnerabilities, and the vocabulary to situate our technological development in its societal and human context, to reason about the ways in which values are affected and can be embodied. In the context of electricity distribution, we have figured out how to leverage the availability of new sources of sensing, data and actuation to enable high levels of penetration of renewable energy sources in distribution networks, while ensuring the safety, security, stability and reliability of the network, and affordable access and privacy for the citizens relying on it.

This dissertation has been a first step towards addressing the general ambition for doing science and engineering as well as helping the energy transition in electricity distribution. Taking distribution system operation as our motivating domain of problems and research questions, this dissertation primarily tried to utilize new forms of sensing, data collection and actuation to provide higher resolution observability and controllability, both in space and time, for distribution system operators (DSOs), needed to enable the safe and reliable integration of distributed generation and electrification. In Chapter 4, we challenged the inherent limitations of traditional state estimation by developing an estimation technique that relies on a limited set of real-time sensors by leveraging the predictive power of historical load information. This technique can be seen as a critical tool in the transition towards higher resolution state estimation schemes which are economically unfeasible for the foreseeable

future. Our estimation method tries to do the best with the information available, and can help DSOs to better understand what is going on in their networks, and form a critical step towards more active control strategies for power flow and voltage regulation. Its implementation on a real network in The Netherlands suggests its immediate relevance to DSOs. Chapters 5 and 6 presented a novel methodology for enabling sophisticated safety-aware decentralized control of distributed energy resources (DERs) for voltage and power flow regulation. We studied the ability of machine learning algorithms to mimic the execution of centralized optimal power flow (OPF) problems based on solely local information. This method will enable DSOs to do automated data-driven construction of DER controllers for such complex control tasks, avoiding intensive and expensive manual tuning and maintenance. It also circumvents the need for the capital intensive communication infrastructure needed to implement more sophisticated control schemes such as OPF, thereby providing a tool that can be implemented by many DSOs today. The theory developed to design fully decentralized policies that mimic network optimization problems is novel and may inform other engineering problems where decisions are decentralized using machine learning, or where compression is otherwise occurring naturally in the machine learning design process, providing fundamental limits on reconstruction and a strategy for communication or feature selection to improve reconstruction. Chapter 7 completes the work on electricity distribution by considering the impact of distributed optimization schemes for OPF on the privacy of participating DER owners. Though much more work needs to be done to properly define workable notions of privacy in this context, the concept of local differential privacy for distributed optimization provides a fairly general toolkit for studying and designing technical privacy mechanisms in contexts where distributed optimization is implemented.

With the dramatic rise of automated decision-making applications leveraging the commodification of data storage, computing, sensing, actuation, and algorithms, many opportunities and aspirations for making lives better around the world are accompanied by relevant concerns around the safe, beneficial and just integration of such technology. Though most contributions in this dissertation are of a technical nature, the author has spend time understanding the social implications of automated decision-making [55], and organizing changes on the Berkeley campus to stimulate more cross-disciplinary debate and scholarship around this topic [53]. Chapter 8 attempts to draw connections between the technical ways in which biases occur in the design and deployment of automated decision-making systems, and the inherent social and political realities that these intervene in. Arguments were made for more reflection on our practices, tools and ways of knowing and engagement with the domain experts, beneficiaries and other relevant stakeholders, as a way to complement the technical rigor of our work, together working towards a world in which automated decision-making is integrated and adopted to genuinely improve lives and prevent harmful outcomes.

9.1 Future Research Directions

The main research areas of improving observability and controllability in distribution networks through data-driven monitoring and control still have relevant follow up questions to be pursued, both theoretically and in terms of practical implementation. As the grid is transitioning, new issues arise requiring further attention from system and control theorists. In addition, the multi-disciplinary nature of this dissertation has helped to start addressing the question of beneficial integration of data-driven monitoring and control methods in a more comprehensive way. However, the focus of this dissertation has been on a limited set of questions, leaving space for future work. Lastly, some of the methodologies developed for distribution system operation may be generalized and find application in other domains. Here, we list a set of research directions that is by no means exhaustive, but reflects the main interests of the author at this point in time.

Operationalizing Learning-based Monitoring and Control Systems

The implementation of data-driven methods, such as decentralized optimal power flow (Chapter 5) or forecast-based state estimation (Chapter 4), rely on either supervised or reinforcement learning techniques. In the former case, a critical challenge is constructing a data set that sufficiently represents the phenomena that will be encountered during implementation, and to train a machine learning model that generalizes well. In Chapter 5, we point at various challenges that may show up in operationalizing a data-driven system, including a limited ability to collect data before system design or encountering system changes. The context in which a system is implemented will dictate the extent to which inherent errors and biases can be accepted or mitigated. It is relevant to study how our work may be complemented by a safety controllers that kick in when the learning-based controller behaves in unexpected ways. Recent advances in safe learning for dynamical systems may be applicable to understand whether the learning process, which we have considered as offline, can be moved online by employing safety controllers, thereby allowing the learning-based controllers to be determined and updated as new data becomes available [3, 78].

Studying Further Decentralization and Grid-Forming Control in Power Systems

The inevitable transition to power systems fueled by solely renewable energy sources is only getting started. If higher levels of penetration of distributed generation are pursued, distribution systems run into inherent limitations [204]. When most of the energy is generated locally, the source of stability from large-scale centralized generation will deplete, and local networks will have to organize the control of a stable voltage signal by themselves, necessitating *grid-forming* rather than *grid-following* inverter and DER control architectures [159, 112, 114]. Motivated by the exciting idea that local networks will have to become more self-sufficient, this transition may contribute to more equitable access and control over electric energy.

However, much work remains to develop well-understood, scalable and affordable grid-forming control techniques; a challenge that necessitates an integral approach, designing the hardware and control strategies in unison. It is relevant to understand how data-driven methodologies, such as those developed in this dissertation, can facilitate this design challenge. Since the dynamics in situations of low voltage and low inertia with inverter-interfaced generation are inherently different from the quasi-steady state power flow dynamics considered in this dissertation, characterizing these dynamics is a first necessity before data-driven monitoring and control methods can be considered.

Develop a More General Information-Theoretic Framework for Data-Driven Controllers

The rate distortion framework developed in Chapter 6 provides a more general perspective for analyzing machine learning models with access to limited information. The framing of such problems as compression and reconstruction is useful and complements the standard machine learning practice of training, testing and validation. It provides fundamental limits on how well a model can ever do given limited information, which may help as a guideline for model class selection, feature selection and communication/sensing. It is relevant to further explore the information theory literature, to understand how this analysis may contribute to other critical controller specifications, such as (closed-loop) stability, robustness and algorithmic convergence and optimality.

Studying Dynamics of Automated Decision-Making

As discussed in Chapter 8, many automated decision-making systems applied in societal contexts such as credit lending, medical practices or traffic allocation are often designed from a static point of view, not taking into account the dynamic effects that arise as the system is implemented. Dynamic coupling between the decision-making system and its environment may indicate bias in a training set that was gathered without decision-making occurring; an issue requiring an online learning, reinforcement learning or adaptive control approach. But even if a system learns online, the dynamics of the directions in which the system learns may be unstable, leading to new emerging biases that may be harmful or otherwise undesirable. The importance of such dynamics seems to be largely overlooked among most practitioners and users of automated decision-making and more literature search needs to be done to understand gaps in the literature. An effective analysis of dynamics seems to be context-specific, as often the environment acted upon are human beings or other complex processes that require cross-disciplinary analyses to model dynamics and inform the iterative design of automated decision-making systems.

Developing a Value-Sensitive Design Approach for Automated Decision-Making

The rapid adoption of computer systems in the 1980s and 1990s forms an important source of inspiration for understanding how we may deal with the rapid integration of algorithms to automate decision-making. In those days, the societal implications of computers and the needs of users were a motivation to develop *value-sensitive design* (VSD) approaches. VSD is a theoretically grounded approach to the design of technology that accounts for human values in a principled and comprehensive manner by actively engaging with the ethical values of direct and indirect stakeholders of the technology [82]. This theory and methodology has analyzed the sources of bias in computer systems [84], it has been adopted in the context of information systems [210, 85] and it has found its way into other domains of engineering and science [83], also influencing relevant European science and research policy narratives around *responsible innovation* [200, 211]. In Chapter 8, we are making connections between the VSD perspective and the implications of automated decision-making. These connections form the first steps towards a comprehensive approach to help practitioners and stakeholders reason about how values may be affected by automated decision-making systems, and provide principles to integrate values and track the expression of values throughout the design and deployment process.

Bibliography

- [1] P. Abbeel and A. Y. Ng. “Apprenticeship Learning via Inverse Reinforcement Learning”. In: *International Conference on Machine Learning*. New York, NY, USA: ACM, 2004.
- [2] A. Abur and A. Gomez Exposito. *Power System State Estimation: Theory and Implementation*. Vol. 24. Power Engineering (Willis). CRC Press, Mar. 24, 2004.
- [3] A. K. Akametalu et al. “Reachability-based safe learning with Gaussian processes”. In: *53rd IEEE Conference on Decision and Control*. 53rd IEEE Conference on Decision and Control. Dec. 2014, pp. 1424–1431.
- [4] D. Amodei et al. “Concrete Problems in AI Safety”. In: *arXiv:1606.06565 [cs]* (June 21, 2016). arXiv: 1606.06565.
- [5] R. Arghandeh, A. Onen, and R. Broadwater. “Distributed energy storage system control for optimal adoption of electric vehicles”. In: *2012 IEEE Power and Energy Society General Meeting*. 2012 IEEE Power and Energy Society General Meeting. July 2012, pp. 1–8.
- [6] R. Arghandeh et al. “On the definition of cyber-physical resilience in power systems”. In: *Renewable and Sustainable Energy Reviews* 58 (2016), pp. 1060–1069.
- [7] D. B. Arnold et al. “Extremum Seeking control of smart inverters for VAR compensation”. In: *2015 IEEE Power Energy Society General Meeting*. 2015 IEEE Power Energy Society General Meeting. July 2015, pp. 1–5.
- [8] D. B. Arnold et al. “Model-free optimal control of var resources in distribution systems: An extremum seeking approach”. In: *IEEE Transactions on Power Systems* 31.5 (2016), pp. 3583–3593.
- [9] J. Arrillaga, M. H. J. Bollen, and N. Watson. “Power quality following deregulation”. In: *Proceedings of the IEEE* 88.2 (Feb. 2000), pp. 246–261.
- [10] K. J. Åström and R. M. Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- [11] R. J. Aumann and J. H. Dreze. “Cooperative games with coalition structures”. In: *International Journal of Game Theory* 3.4 (Dec. 1, 1974), pp. 217–237.
- [12] X. Bai et al. “Semidefinite programming for optimal power flow problems”. In: *International Journal of Electrical Power & Energy Systems* 30.6 (2008), pp. 383–392.

- [13] C. Barabas et al. “Interventions over Predictions: Reframing the Ethical Debate for Actuarial Risk Assessment”. In: 1st Conference on Fairness, Accountability, and Transparency. New York, NY, USA, Feb. 2018. arXiv: 1712.08238.
- [14] M. Baradar and M. R. Hesamzadeh. “Calculating negative LMPs from SOCP-OPF”. In: *Energy Conference (ENERGYCON), 2014 IEEE International*. IEEE, 2014, pp. 1461–1466.
- [15] M. Baran and F. Wu. “Optimal capacitor placement on radial distribution systems”. In: *IEEE Transactions on Power Delivery* 4.1 (Jan. 1989), pp. 725–734.
- [16] M. E. Baran and I. M. El-Markabi. “A multiagent-based dispatching scheme for distributed generators for voltage support on distribution feeders”. In: *IEEE Transactions on power systems* 22.1 (2007), pp. 52–59.
- [17] M. E. Baran and F. F. Wu. “Optimal sizing of capacitors placed on a radial distribution system”. In: *Power Delivery, IEEE Transactions on* 4.1 (1989), pp. 735–743.
- [18] S. Barocas and A. D. Selbst. “Big data’s disparate impact”. In: *Cal. L. Rev.* 104 (2016), p. 671.
- [19] M. Behl, A. Jain, and R. Mangharam. “Data-Driven Modeling, Control and Tools for Cyber-Physical Energy Systems”. In: 7th International Conference on Cyber-Physical Systems (ICCPs). Vienna, Austria: ACM, Apr. 2016.
- [20] F. Bellizio et al. “Optimized local control schemes for active distribution grids using machine learning techniques”. In: *Submitted to IEEE PES GM*. Nov. 2017.
- [21] B. Biegel. *The Danish Smart Grid Model*. Smartgriddemark.com. May 2014. URL: <http://www.smartgriddemark.com/> (visited on 05/05/2014).
- [22] E. Blood, B. Krogh, and M. Ilic. “Electric power system static state estimation through Kalman filtering and load forecasting”. In: Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century. IEEE, July 2008.
- [23] S. Bolognani and S. Zampieri. “A Distributed Control Strategy for Reactive Power Compensation in Smart Microgrids”. In: *IEEE Transactions on Automatic Control* 58.11 (Nov. 2013), pp. 2818–2833.
- [24] S. Bolognani and S. Zampieri. “On the Existence and Linear Approximation of the Power Flow Solution in Power Distribution Networks”. In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016), pp. 163–172.
- [25] S. Bolognani and F. Dörfler. “Fast power system analysis via implicit linearization of the power flow manifold”. In: *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*. IEEE, 2015, pp. 402–409.
- [26] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [27] S. Boyd et al. “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers”. In: *Foundations and Trends® in Machine Learning* 3.1 (July 26, 2011), pp. 1–122.

- [28] M. Brambilla et al. “Swarm robotics: a review from the swarm engineering perspective”. In: *Swarm Intelligence* 7.1 (Mar. 1, 2013), pp. 1–41.
- [29] C. Cadwalladr. “‘I made Steve Bannon’s psychological warfare tool’: meet the data war whistleblower”. In: *The Guardian* (Mar. 18, 2018).
- [30] G. Calafiore and M. Campi. “The scenario approach to robust control design”. In: *IEEE Transactions on Automatic Control* 51.5 (May 2006), pp. 742–753.
- [31] G. Calafiore and M. C. Campi. “Uncertain convex programs: randomized solutions and confidence levels”. In: *Mathematical Programming* 102.1 (2005), pp. 25–46.
- [32] V. Calderaro et al. “Optimal decentralized voltage control for distribution systems with inverter-based distributed generators”. In: *Power Systems, IEEE Transactions on* 29.1 (2014), pp. 230–241.
- [33] D. Callaway and I. Hiskens. “Achieving Controllability of Electric Loads”. In: *Proceedings of the IEEE* 99.1 (Jan. 2011), pp. 184–199.
- [34] D. Callaway. *CITRIS Research Exchange: Demand-Side Modeling, Estimation and Control in Electric Power Systems*. UC, Nov. 13, 2013.
- [35] A. Campolo et al. *AI Now 2017 Report*. New York, NY, USA: AI Now Institute, 2017.
- [36] D. Cardwell. “Solar Power Battle Puts Hawaii at Forefront of Worldwide Changes”. In: *The New York Times* (Apr. 18, 2015).
- [37] P. M. Carvalho, P. F. Correia, and L. A. Ferreira. “Distributed reactive power generation control for voltage rise mitigation in distribution networks”. In: *IEEE Transactions on Power Systems* 23.2 (2008), pp. 766–772.
- [38] *Center for Human-Compatible AI*. 2018. URL: <http://humancompatible.ai/about> (visited on 04/29/2018).
- [39] P. D. Christofides et al. “Distributed model predictive control: A tutorial review and future research directions”. In: *Computers & Chemical Engineering* 51 (2013), pp. 21–41.
- [40] E. H. Clarke. “Multipart pricing of public goods”. In: *Public Choice* 11.1 (Sept. 1, 1971), pp. 17–33.
- [41] *Conference on Fairness, Accountability, and Transparency in Socio-Technical Systems (FAT*)*. 2018. URL: <https://fatconference.org/> (visited on 04/29/2018).
- [42] J. Cortés et al. “Differential privacy in control and network systems”. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*. 2016 IEEE 55th Conference on Decision and Control (CDC). Dec. 2016, pp. 4252–4272.
- [43] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [44] K. Crawford. “Keynote: The Trouble with Bias”. NIPS 2017. Long Beach, CA, USA, Dec. 2017.

- [45] E. Dall’Anese, G. Giannakis, and B. Wollenberg. “Optimization of unbalanced power distribution networks via semidefinite relaxation”. In: North American Power Symposium. Sept. 2012.
- [46] E. Dall’Anese, H. Zhu, and G. Giannakis. “Distributed Optimal Power Flow for Smart Microgrids”. In: *IEEE Transactions on Smart Grid* 4.3 (Sept. 2013), pp. 1464–1475.
- [47] E. Dall’Anese, S. V. Dhople, and G. Giannakis. “Optimal dispatch of photovoltaic inverters in residential distribution systems”. In: *Sustainable Energy, IEEE Transactions on* 5.2 (2014), pp. 487–497.
- [48] E. Dall’Anese and A. Simonetto. “Optimal power flow pursuit”. In: *IEEE Transactions on Smart Grid* 9.2 (2018), pp. 942–952.
- [49] J. Dalton. “Scotland produces record amount of energy from renewables as green schemes generate two-thirds of electricity”. In: *The Independent* (Mar. 31, 2018).
- [50] *Dataport*. Pecan Street Inc. Dataport. 2017. URL: <http://www.pecanstreet.org/>.
- [51] E. J. Davison and T. N. Chang. “Decentralized stabilization and pole assignment for general proper systems”. In: *IEEE Transactions on Automatic Control* 35.6 (1990), pp. 652–664.
- [52] P. De Martini. “Business & Operational Implications of the Convergence of Bulk Power & Distributed Energy”. UC Berkeley, Apr. 3, 2014.
- [53] R. Dobbe. *Time to redefine what it means to be in tech*. Medium. Dec. 4, 2017. URL: <https://medium.com/@roeldobbe/from-techno-optimism-to-techno-dystopia-time-to-redefine-what-it-means-to-be-in-tech-7d65a83939a> (visited on 05/04/2018).
- [54] R. Dobbe, D. Fridovich-Keil, and C. Tomlin. “Fully Decentralized Policies for Multi-Agent Systems: An Information Theoretic Approach”. In: Conference on Neural Information Processing Systems. Long Beach, CA, USA, Dec. 2017.
- [55] R. Dobbe et al. “Automating Us - The entanglement of people and machines”. In: *The Berkeley Science Review* (Fall 2015 Issue 2015).
- [56] R. Dobbe et al. “Customized Local Differential Privacy for Multi-Agent Distributed Optimization”. In: *IEEE Transactions on Control of Network Systems* (In preparation) (May 2018).
- [57] R. Dobbe et al. “Decentralized Optimal Power Flow”. In: *IEEE Transactions on Smart Grid* (Under review) (Apr. 2018).
- [58] R. Dobbe et al. “Forecasting-Based State Estimation in Single- and Three-Phase Distribution Systems”. In: *IEEE Transactions on Power Systems* (Under Review) (May 2018).
- [59] R. Dobbe et al. “Real-time distribution grid state estimation with limited sensors and load forecasting”. In: *7th International Conference on Cyber-Physical Systems (ICCPs)*. ACM, May 2016, pp. 1–10.

- [60] R. Dobbe et al. “The Epistemology and Dynamics of Automated Decision-Making”. In: *International Conference on Machine Learning (under review)*. Apr. 2018.
- [61] R. Dong. “Dissertation talk: Privacy-by-Design and Data Markets for the Internet of Things”. UC, 2017.
- [62] R. Dong et al. “Differential privacy of populations in routing games”. In: *2015 54th IEEE Conference on Decision and Control (CDC)*. 2015 54th IEEE Conference on Decision and Control (CDC). Dec. 2015, pp. 2798–2803.
- [63] F. Dörfler and F. Bullo. “Synchronization and Transient Stability in Power Networks and Nonuniform Kuramoto Oscillators”. In: *SIAM Journal on Control and Optimization* 50.3 (Jan. 1, 2012), pp. 1616–1642.
- [64] F. Dörfler, J. W. Simpson-Porco, and F. Bullo. “Breaking the Hierarchy: Distributed Control and Economic Optimality in Microgrids”. In: *IEEE Transactions on Control of Network Systems* 3.3 (Sept. 2016), pp. 241–253.
- [65] J. C. Duchi, M. I. Jordan, and M. J. Wainwright. “Local privacy and statistical minimax rates”. In: *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013, pp. 429–438.
- [66] C. Dwork and A. Roth. “The Algorithmic Foundations of Differential Privacy”. In: *Foundations and Trends® in Theoretical Computer Science* 9.3 (Aug. 11, 2014), pp. 211–407.
- [67] C. Dwork et al. “Calibrating noise to sensitivity in private data analysis”. In: *TCC*. Vol. 3876. Springer, 2006, pp. 265–284.
- [68] K. J. Dyke, N. Schofield, and M. Barnes. “The impact of transport electrification on electrical networks”. In: *IEEE Transactions on Industrial Electronics* 57.12 (2010), pp. 3917–3926.
- [69] *Economic Benefits of Increasing Electric Grid Resilience to Weather Outages*. Executive Office of the President, Aug. 2013.
- [70] D. Ensign et al. “Runaway feedback loops in predictive policing”. In: *Conference on Fairness, Accountability and Transparency*. New York, NY, USA, Feb. 2018.
- [71] *EU GDPR Information Portal*. EU GDPR Portal. 2018. URL: <http://eugdpr.org/eugdpr.org.html> (visited on 04/29/2018).
- [72] Eugene Shlatz, Nathan Buch, and Melissa Chan. *Distributed Generation Integration Cost Study - Analytical Framework*. California Energy Commission, Nov. 2013.
- [73] *Fairness, Accountability, and Transparency in Machine Learning*. 2018. URL: <https://www.fatml.org/> (visited on 04/29/2018).
- [74] M. Farivar, L. Chen, and S. Low. “Equilibrium and dynamics of local voltage control in distribution systems”. In: *2013 IEEE 52nd Annual Conference on Decision and Control (CDC)*. 2013 IEEE 52nd Annual Conference on Decision and Control (CDC). Dec. 2013, pp. 4329–4334.

- [75] M. Farivar and S. H. Low. “Branch flow model: Relaxations and convexification—Part I”. In: *IEEE Transactions on Power Systems* 28.3 (2013), pp. 2554–2564.
- [76] M. Farivar et al. “Optimal inverter VAR control in distribution systems with high PV penetration”. In: *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–7.
- [77] J. Fialka. *As Hawaii Aims for 100% Renewable Energy, Other States Watching Closely*. Scientific American. Apr. 27, 2018. URL: <https://www.scientificamerican.com/article/as-hawaii-aims-for-100-renewable-energy-other-states-watching-closely/> (visited on 04/29/2018).
- [78] J. F. Fisac et al. “A general safety framework for learning-based control in uncertain robotic systems”. In: *arXiv preprint arXiv:1705.01292* (2017).
- [79] M. Flores et al. “P4 medicine: how systems medicine will transform the healthcare sector and society”. In: *Personalized medicine* 10.6 (2013), pp. 565–576.
- [80] S. Frank and S. Rebennack. “An introduction to optimal power flow: Theory, formulation, and examples”. In: *IIE Transactions* 48.12 (Dec. 1, 2016), pp. 1172–1197.
- [81] B. Friedland. *Control system design: an introduction to state-space methods*. Courier Corporation, 2012.
- [82] B. Friedman. “Value-sensitive Design”. In: *interactions* 3.6 (Dec. 1996), pp. 16–23.
- [83] B. Friedman, D. G. Hendry, and A. Borning. “A Survey of Value Sensitive Design Methods”. In: *Foundations and Trends® in Human–Computer Interaction* 11.2 (Nov. 22, 2017), pp. 63–125.
- [84] B. Friedman and H. Nissenbaum. “Bias in Computer Systems”. In: *ACM Trans. Inf. Syst.* 14.3 (July 1996), pp. 330–347.
- [85] B. Friedman et al. “Value sensitive design and information systems”. In: *Early engagement and new technologies: Opening up the laboratory*. Springer, 2013, pp. 55–95.
- [86] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*. Vol. 1. Springer series in statistics Springer, Berlin, 2001.
- [87] L. Gan and S. H. Low. “Convex relaxations and linear approximation for optimal power flow in multiphase radial networks”. In: *Power Systems Computation Conference*. Aug. 2014.
- [88] L. Gan et al. “Exact convex relaxation of optimal power flow in radial networks”. In: *IEEE Transactions on Automatic Control* 60.1 (2015), pp. 72–87.
- [89] G. Giannakis et al. “Monitoring and Optimization for Power Grids: A Signal Processing Perspective”. In: *IEEE Signal Processing Magazine* 30.5 (Sept. 2013), pp. 107–128.
- [90] M. Gol and A. Abur. “A Hybrid State Estimator For Systems With Limited Number of PMUs”. In: *IEEE Transactions on Power Systems* 30.3 (May 2015), pp. 1511–1517.

- [91] C. V. Goldman and S. Zilberstein. “Decentralized control of cooperative systems: Categorization and complexity analysis”. In: *Journal of Artificial Intelligence Research (JAIR)* 22 (2004), pp. 143–174.
- [92] A. Gomez Exposito and A. Abur. “Generalized observability analysis and measurement classification”. In: *20th International Conference on Power Industry Computer Applications., 1997.* 20th International Conference on Power Industry Computer Applications., 1997. May 1997, pp. 97–103.
- [93] R. C. Green, L. Wang, and M. Alam. “The impact of plug-in hybrid electric vehicles on distribution networks: A review and outlook”. In: *Renewable and Sustainable Energy Reviews* 15.1 (2011), pp. 544–553.
- [94] W. H. Greene. *Econometric analysis (International edition)*. 7th ed. Pearson, 2000.
- [95] T. Groves. “Incentives in Teams”. In: *Econometrica* 41.4 (1973), pp. 617–631.
- [96] J. Guerrero et al. “Advanced Control Architectures for Intelligent Microgrids - Part II: Power Quality, Energy Storage, and AC/DC MicroGrids”. In: *IEEE Transactions Industrial Electronics* 99 (2013).
- [97] S. Han, U. Topcu, and G. J. Pappas. “Differentially Private Distributed Constrained Optimization”. In: *IEEE Transactions on Automatic Control* 62.1 (Jan. 2017), pp. 50–64.
- [98] M. Hardt and S. Barocas. “NIPS 2017 - Tutorial on Fairness in Machine Learning”. Tutorial. Long Beach, CA, USA, Dec. 2017.
- [99] M. Hardt, E. Price, and N. Srebro. “Equality of opportunity in supervised learning”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 3315–3323.
- [100] G. W. Hart. “Nonintrusive appliance load monitoring”. In: *Proceedings of the IEEE* 80.12 (Dec. 1992), pp. 1870–1891.
- [101] A. Hauswirth et al. “Online optimization in closed loop on the power flow manifold”. In: *PowerTech, 2017 IEEE Manchester*. IEEE, 2017.
- [102] M. Heidegger and W. Lovitt. *The question concerning technology, and other essays*. Harper & Row New York, 1977.
- [103] J. van den Hoven. “Value Sensitive Design and Responsible Innovation”. In: *Responsible Innovation*. Ed. by R. Owen, J. Bessant, and g. Heintz. John Wiley & Sons, Ltd, 2013, pp. 75–83.
- [104] J. Hsu et al. “Privately Solving Linear Programs”. In: *arXiv:1402.3631 [cs]* (Feb. 14, 2014). arXiv: 1402.3631.
- [105] S.-J. Huang and K.-R. Shih. “Short-term load forecasting via ARMA model identification including non-Gaussian process considerations”. In: *IEEE Transactions on Power Systems* 18.2 (May 2003), pp. 673–679.

- [106] Z. Huang, S. Mitra, and N. Vaidya. “Differentially Private Distributed Optimization”. In: *Proceedings of the 2015 International Conference on Distributed Computing and Networking*. ICDCN '15. New York, NY, USA: ACM, 2015, 4:1–4:10.
- [107] *IEEE Distribution Test Feeders*. Distribution Test Feeders. 2017. URL: <http://ewh.ieee.org/soc/pes/dsacom/testfeeders/>.
- [108] *International Workshop on Artificial Intelligence Safety Engineering - WAISE 2018*. Home | WAISE 2018. 2018. URL: <https://www.waise2018.com/scope-topics> (visited on 04/29/2018).
- [109] S. Jabbari et al. “Fair learning in Markovian environments”. In: *arXiv preprint arXiv:1611.03071* (2016).
- [110] R. Jia et al. “Privacy-enhanced Architecture for Occupancy-based HVAC Control”. In: *Proceedings of the 8th International Conference on Cyber-Physical Systems*. ICCPS '17. New York, NY, USA: ACM, 2017, pp. 177–186.
- [111] J. Jiao et al. “Minimax estimation of functionals of discrete distributions”. In: *IEEE Transactions on Information Theory* 61.5 (2015), pp. 2835–2885.
- [112] B. B. Johnson et al. “Synthesizing Virtual Oscillators to Control Islanded Inverters”. In: *IEEE Transactions on Power Electronics* 31.8 (Aug. 2016), pp. 6002–6015.
- [113] M. Jordan. *Artificial Intelligence — The Revolution Hasn't Happened Yet*. Medium. Apr. 19, 2018. URL: <https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7> (visited on 04/29/2018).
- [114] T. Jouini, C. Arghir, and F. Dorfler. “Grid-forming control for power converters based on matching of synchronous machines”. In: *arXiv preprint arXiv:1706.09495* (2017).
- [115] R. E. Kalman. “A new approach to linear filtering and prediction problems”. In: *Journal of basic Engineering* 82.1 (1960), pp. 35–45.
- [116] S. Karagiannopoulos, P. Aristidou, and G. Hug. “Co-optimisation of planning and operation for active distribution grids”. In: *PowerTech*. Manchester, UK: IEEE, 2017.
- [117] S. Karagiannopoulos, P. Aristidou, and G. Hug. “Hybrid approach for planning and operating active distribution grids”. In: *IET Generation, Transmission & Distribution* 11.3 (2017), pp. 685–695.
- [118] V. Katewa, F. Pasqualetti, and V. Gupta. “On privacy vs. cooperation in multi-agent systems”. In: *International Journal of Control* (May 8, 2017).
- [119] G. Kaur and M. Vaziri. “Effects of distributed generation (DG) interconnections on protection of distribution feeders”. In: *IEEE Power Engineering Society General Meeting, 2006*. IEEE Power Engineering Society General Meeting, 2006. 2006.
- [120] A. Keane et al. “Enhanced utilization of voltage control resources with distributed generation”. In: *IEEE Transactions on Power Systems* 26.1 (2011), pp. 252–260.

- [121] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. “Rate control for communication networks: shadow prices, proportional fairness and stability”. In: *Journal of the Operational Research Society* 49.3 (Mar. 1, 1998), pp. 237–252.
- [122] W. H. Kersting. *Distribution System Modeling and Analysis, Third Edition*. CRC Press, Jan. 24, 2012. 459 pp.
- [123] J. Kleinberg, S. Mullainathan, and M. Raghavan. “Inherent Trade-Offs in the Fair Determination of Risk Scores”. In: 8th Innovations in Theoretical Computer Science Conference (ITCS 2017). Vol. 67. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, 43:1–43:23. arXiv: 1609.05807.
- [124] N. Kohli, R. Barreto, and J. A. Kroll. “Translation Tutorial: A Shared Lexicon for Research and Practice in Human-Centered Software Systems”. In: 1st Conference on Fairness, Accountability, and Transparency. New York, NY, USA, Feb. 2018, p. 7.
- [125] B. Kroposki et al. “Benefits of power electronic interfaces for distributed energy systems”. In: IEEE Power Engineering Society General Meeting. 2006.
- [126] P. Kundur et al. “Definition and classification of power system stability IEEE/CIGRE joint task force on stability terms and definitions”. In: *IEEE Transactions on Power Systems* 19.3 (Aug. 2004), pp. 1387–1401.
- [127] S. Lacey and R. Hanley. *The Interchange Podcast with Ryan Hanley: Building a Blueprint for the Transactive Grid at PG&E, SolarCity, Tesla and AMS*. Green Tech Media. Mar. 6, 2018. URL: <https://www.greentechmedia.com/articles/read/the-blueprint-for-the-distributed-grid> (visited on 03/31/2018).
- [128] J. Lavaei and S. Low. “Zero Duality Gap in Optimal Power Flow Problem”. In: *IEEE Transactions on Power Systems* 27.1 (Feb. 2012), pp. 92–107.
- [129] B. C. Lesieutre et al. “Examining the limits of the application of semidefinite programming to power flow problems”. In: *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. IEEE, 2011, pp. 1492–1499.
- [130] P. Lewis. “‘Our minds can be hijacked’: the tech insiders who fear a smartphone dystopia”. In: *The Guardian* (Oct. 6, 2017).
- [131] N. Li, G. Qu, and M. Dahleh. “Real-time decentralized voltage control in distribution networks”. In: *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*. IEEE, 2014, pp. 582–588.
- [132] R. Lim and M. Lewis. *The Retail Experience Economy: The Behavioural Revolution*. UK: Retail Economics and Squire Patton Boggs, 2017.
- [133] L. T. Liu et al. “Delayed Impact of Fair Machine Learning”. In: *arXiv:1803.04383 [cs, stat]* (Mar. 12, 2018). arXiv: 1803.04383.
- [134] L. Ljung. *System identification*. Springer, 1998.

- [135] R. Louca, P. Seiler, and E. Bitar. “Nondegeneracy and inexactness of semidefinite relaxations of optimal power flow”. In: *arXiv preprint arXiv:1411.4663* (2014).
- [136] S. Low. “Convex Relaxation of Optimal Power Flow; Part I: Formulations and Equivalence”. In: *IEEE Transactions on Control of Network Systems* 1.1 (Mar. 2014), pp. 15–27.
- [137] S. H. Low. “Convex relaxation of optimal power flow—Part II: Exactness”. In: *IEEE Transactions on Control of Network Systems* 1.2 (2014), pp. 177–189.
- [138] J. Lunze. *Feedback control of large-scale systems*. Prentice Hall New York, 1992.
- [139] M. Maasoumy et al. “Model Predictive Control Approach to Online Computation of Demand-Side Flexibility of Commercial Buildings HVAC Systems for Supply Following”. In: *IEEE American Control Conference, ACC 2014*. 2014.
- [140] R. Madani, S. Sojoudi, and J. Lavaei. “Convex Relaxation for Optimal Power Flow Problem: Mesh Networks”. In: *IEEE Transactions on Power Systems* 30.1 (Jan. 2015), pp. 199–211.
- [141] E. Manitsas et al. “Distribution System State Estimation Using an Artificial Neural Network Approach for Pseudo Measurement Modeling”. In: *IEEE Transactions on Power Systems* 27.4 (Nov. 2012), pp. 1888–1896.
- [142] A. von Meier et al. “Micro-synchrophasors for distribution systems”. In: *IEEE 5th Innovative Smart Grid Technologies Conference, Washington, DC*. 2014.
- [143] P. Mirowski et al. “Demand forecasting in smart grids”. In: *Bell Labs technical journal* 18.4 (2014), pp. 135–158.
- [144] Y. Mo and R. M. Murray. “Privacy Preserving Average Consensus”. In: *IEEE Transactions on Automatic Control* 62.2 (Feb. 2017), pp. 753–765.
- [145] P. J. Modi et al. “ADOPT: Asynchronous distributed constraint optimization with quality guarantees”. In: *Artificial Intelligence* 161.1 (2005), pp. 149–180.
- [146] D. K. Molzahn, B. C. Lesieutre, and C. L. DeMarco. “A sufficient condition for global optimality of solutions to the optimal power flow problem”. In: *IEEE Transactions on Power Systems* 29.2 (2014), pp. 978–979.
- [147] D. K. Molzahn, B. C. Lesieutre, and C. L. DeMarco. “Investigation of non-zero duality gap solutions to a semidefinite relaxation of the optimal power flow problem”. In: *System Sciences (HICSS), 2014 47th Hawaii International Conference on*. IEEE, 2014, pp. 2325–2334.
- [148] J. A. Momoh et al. “Challenges to optimal power flow”. In: *IEEE Transactions on Power Systems* 12.1 (Feb. 1997), pp. 444–455.
- [149] A. Monticelli and F. F. Wu. “Network Observability: Theory”. In: *IEEE Transactions on Power Apparatus and Systems* PAS-104.5 (May 1985), pp. 1042–1048.

- [150] H. Mori and M. Ohmi. “Probabilistic short-term load forecasting with Gaussian processes”. In: 13th International Conference on Intelligent Systems Application to Power Systems. Nov. 2005.
- [151] R. Nair et al. “Networked Distributed POMDPs: A synthesis of distributed constraint optimization and POMDPs”. In: *AAAI*. Vol. 5. 2005, pp. 133–139.
- [152] A. Narayanan. “FAT* tutorial: 21 fairness definitions and their politics”. New York, NY, USA, Feb. 2018.
- [153] H. Nissenbaum. “Privacy as contextual integrity”. In: *Wash. L. Rev.* 79 (2004), p. 119.
- [154] E. Nozari, P. Tallapragada, and J. Cortes. “Differentially Private Distributed Convex Optimization via Functional Perturbation”. In: *IEEE Transactions on Control of Network Systems* PP.99 (2016).
- [155] F. Oldewurtel et al. “A framework for and assessment of demand response and energy storage in power systems”. In: Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid, 2013 IREP Symposium. Aug. 2013.
- [156] F. Oldewurtel et al. “Energy efficient building climate control using Stochastic Model Predictive Control and weather predictions”. In: *American Control Conference (ACC), 2010*. American Control Conference (ACC), 2010. June 2010, pp. 5100–5105.
- [157] F. A. Oliehoek and C. Amato. *A concise introduction to decentralized POMDPs*. Vol. 1. Springer, 2016.
- [158] P. A. Ortega et al. “Information-Theoretic Bounded Rationality”. In: *arXiv preprint* (2015). arXiv: 1512.06789.
- [159] E. Ortjohann et al. “Grid-Forming Three-Phase Inverters for Unbalanced Loads in Hybrid Power Systems”. In: 4th World Conference on Photovoltaic Energy Conference. Vol. 2. IEEE, May 2006, pp. 2396–2399.
- [160] N. Parikh and S. Boyd. “Proximal Algorithms”. In: *Foundations and Trends® in Optimization* 1.3 (Jan. 13, 2014), pp. 127–239.
- [161] L. Peshkin et al. “Learning to cooperate via policy search”. In: *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2000, pp. 489–496.
- [162] A. Primadianto and C. N. Lu. “A Review on Distribution System State Estimation”. In: *IEEE Transactions on Power Systems* 32.5 (Sept. 2017), pp. 3875–3883.
- [163] Y. Pu, M. N. Zeilinger, and C. N. Jones. “Inexact fast alternating minimization algorithm for distributed model predictive control”. In: 53rd IEEE Conference on Decision and Control. Los Angeles, CA, USA, Dec. 2014, pp. 5915–5921.
- [164] Y. Pu, C. N. Jones, and M. N. Zeilinger. “Inexact Alternating Minimization Algorithm for Distributed Optimization with an Application to Distributed MPC”. In: *arXiv:1608.00413 [math]* (Aug. 1, 2016). arXiv: 1608.00413.

- [165] R. L. Raffard, C. J. Tomlin, and S. P. Boyd. “Distributed optimization for cooperative agents: Application to formation flight”. In: *Conference on Decision and Control*. Nassau, The Bahamas: IEEE, 2004.
- [166] A. Rakhlin, O. Shamir, and K. Sridharan. “Making gradient descent optimal for strongly convex stochastic optimization”. In: *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*. 2012, pp. 449–456.
- [167] C. E. Rasmussen. “Gaussian processes in machine learning”. In: *Advanced lectures on machine learning*. Springer, 2004, pp. 63–71.
- [168] L. J. Ratliff et al. “Social game for building energy efficiency: Incentive design”. In: *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*. IEEE, 2014, pp. 1011–1018.
- [169] B. Recht. *The Ethics of Reward Shaping*. arg min blog. Apr. 2018. URL: <http://benjamin-recht.github.io/2018/04/16/ethical-rewards/> (visited on 04/19/2018).
- [170] S. Reed. “Dutch Utility Bets Its Future on an Unusual Strategy: Selling Less Power”. In: *The New York Times* (Aug. 18, 2017).
- [171] *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance)*. May 4, 2016.
- [172] B. Robbins, C. Hadjicostis, and A. Dominguez-Garcia. “A Two-Stage Distributed Architecture for Voltage Control in Power Distribution Systems”. In: *IEEE Transactions on Power Systems* 28.2 (May 2013), pp. 1470–1482.
- [173] B. A. Robbins and A. D. Dominguez-Garcia. “Optimal reactive power dispatch for voltage regulation in unbalanced distribution systems”. In: *IEEE Transactions on Power Systems* 31.4 (2016), pp. 2903–2913.
- [174] K. M. Rogers et al. “An authenticated control framework for distributed voltage support on the smart grid”. In: *IEEE Transactions on Smart Grid* 1.1 (2010), pp. 40–47.
- [175] P. Rousseaux, T. Van Cutsem, and T. E. Dy Liacco. “Whither dynamic state estimation?” In: *International Journal of Electrical Power & Energy Systems* 12.2 (Apr. 1990), pp. 104–116.
- [176] C. Sammut. “Automatic construction of reactive control systems using symbolic machine learning”. In: *The Knowledge Engineering Review* 11.1 (1996), pp. 27–42.
- [177] M. Sankur et al. “A Linearized Power Flow Model for Optimization in Unbalanced Distribution Systems”. In: *arXiv preprint arXiv:1606.04492* (2016).

- [178] M. Sankur et al. “Optimal Voltage Phasor Regulation for Switching Actions in Distribution Systems”. In: *IEEE Transaction on Smart Grids* (Under Review) (Apr. 2018).
- [179] L. Schenato et al. “Bayesian linear state estimation using smart meters and PMUs measurements in distribution grids”. In: *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 2014 IEEE International Conference on Smart Grid Communications (SmartGridComm). Nov. 2014, pp. 572–577.
- [180] G. Schwarz. “Estimating the dimension of a model”. In: *The annals of statistics* 6.2 (1978), pp. 461–464.
- [181] A. D. Selbst and S. Barocas. *The Intuitive Appeal of Explainable Machines*. SSRN Scholarly Paper ID 3126971. Rochester, NY: Social Science Research Network, Feb. 19, 2018.
- [182] N. Seltenrich. *The New Grid - Plugging into California’s clean-energy future*. Breakthroughs - The magazine of the UC Berkeley College of Natural Resources. 2013. URL: https://nature.berkeley.edu/breakthroughs/fa12/the_new_grid (visited on 03/21/2014).
- [183] *Senate Bill 350 Clean Energy and Pollution Reduction Act of 2015*. 2015. URL: https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201520160SB350 (visited on 04/29/2018).
- [184] O. Shamir and T. Zhang. “Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes”. In: *International Conference on Machine Learning*. 2013, pp. 71–79.
- [185] Y. Shoukry et al. “Privacy-aware quadratic optimization using partially homomorphic encryption”. In: 55th Conference on Decision and Control (CDC). Las Vegas, NV, USA: IEEE, Dec. 2016, pp. 5053–5058.
- [186] D. D. Siljak. *Decentralized control of complex systems*. Courier Corporation, 2011.
- [187] J. W. Simpson-Porco, F. Dörfler, and F. Bullo. “Synchronization and power sharing for droop-controlled inverters in islanded microgrids”. In: *Automatica* 49.9 (Sept. 2013), pp. 2603–2611.
- [188] J. W. Simpson-Porco et al. “Stability, power sharing, & distributed secondary control in droop-controlled microgrids”. In: *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*. IEEE, 2013, pp. 672–677.
- [189] R. Singh, B. C. Pal, and R. A. Jabr. “Distribution system state estimation through Gaussian mixture model of the load as pseudo-measurement”. In: *IET Generation, Transmission & Distribution* 4.1 (Jan. 1, 2010), pp. 50–59.
- [190] *Small-angle approximation*. In: *Wikipedia*. Page Version ID: 806174630. Oct. 20, 2017.

- [191] J. Smith et al. “Smart inverter volt/var control functions for high penetration of PV on distribution systems”. In: *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*. Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES. Mar. 2011, pp. 1–6.
- [192] S. Sojoudi and J. Lavaei. “Physics of power networks makes hard optimization problems easy to solve”. In: Power and Energy Society General Meeting. IEEE, July 2012.
- [193] O. Sondermeijer et al. “Regression-based Inverter Control for Decentralized Optimal Power Flow and Voltage Regulation”. In: Power & Energy Society General Meeting. Boston, MA, USA: IEEE, July 2016.
- [194] D. Sperling. *Three Revolutions: Steering Automated, Shared, and Electric Vehicles to a Better Future*. None edition. Washington, DC: Island Press, Mar. 1, 2018.
- [195] E. Stewart et al. *Analysis of High-Penetration Levels of Photovoltaics into the Distribution Grid on Oahu, Hawaii: Detailed Analysis of HECO Feeder WF1*. National Renewable Energy Laboratory (NREL), Golden, CO., 2013.
- [196] M. Stroud. *Chicago’s predictive policing tool just failed a major test*. The Verge. Aug. 19, 2016. URL: <https://www.theverge.com/2016/8/19/12552384/chicago-heat-list-tool-failed-rand-test> (visited on 04/24/2018).
- [197] P. Šulc, S. Backhaus, and M. Chertkov. “Optimal distributed control of reactive power via the alternating direction method of multipliers”. In: *IEEE Transactions on Energy Conversion* 29.4 (2014), pp. 968–977.
- [198] A. X. Sun, D. T. Phan, and S. Ghosh. “Fully decentralized AC optimal power flow algorithms”. In: Power & Energy Society General Meeting. Vancouver, Canada: IEEE, July 2013.
- [199] J. Sun. “Impedance-Based Stability Criterion for Grid-Connected Inverters”. In: *IEEE Transactions on Power Electronics* 26.11 (Nov. 2011), pp. 3075–3078.
- [200] B. Taebi et al. “Responsible innovation as an endorsement of public values: the need for interdisciplinary research”. In: *Journal of Responsible Innovation* 1.1 (Jan. 2, 2014), pp. 118–124.
- [201] D. Takacs. “How does your positionality bias your epistemology?” In: *Thought & Action* 27 (2003).
- [202] D. Takacs. “Positionality, epistemology, and social justice in the classroom”. In: *Social Justice* 29.4 (2002), pp. 168–181.
- [203] A. Taub and M. Fisher. “Where Countries Are Tinderboxes and Facebook Is a Match”. In: *The New York Times* (Apr. 21, 2018).
- [204] J. A. Taylor, S. V. Dhople, and D. S. Callaway. “Power systems without fuel”. In: *Renewable and Sustainable Energy Reviews* 57 (2016), pp. 1322–1336.

- [205] *The Impact of Localized Energy Resources on Southern California Edison's Transmission and Distribution System*. Southern California Edison Distribution Engineering and Advanced Technology, May 2012.
- [206] Z. Tufekci. "Opinion | YouTube, the Great Radicalizer". In: *The New York Times* (Mar. 10, 2018).
- [207] K. Turitsyn et al. "Options for Control of Reactive Power by Distributed Photovoltaic Generators". In: *Proceedings of the IEEE* 99.6 (June 2011), pp. 1063–1073.
- [208] A. Ulbig, T. S. Borsche, and G. Andersson. "Impact of low rotational inertia on power system stability and operation". In: *IFAC Proceedings Volumes* 47.3 (2014), pp. 7290–7297.
- [209] P. Van den Hof. "Closed-loop issues in system identification". In: *Annual reviews in control* 22 (1998), pp. 173–186.
- [210] J. Van den Hoven. "ICT and value sensitive design". In: *The information society: Innovation, legitimacy, ethics and democracy in honor of Professor Jacques Berleur SJ* (2007), pp. 67–72.
- [211] J. Van den Hoven. "Value sensitive design and responsible innovation". In: *Responsible innovation: Managing the responsible emergence of science and innovation in society* (2013), pp. 75–83.
- [212] E. Veldman and R. Verzijlbergh. "Distribution Grid Impacts of Smart Electric Vehicle Charging From Different Perspectives". In: *IEEE Transactions on Smart Grid* 6.1 (Jan. 2015), pp. 333–342.
- [213] E. Veldman et al. "Impact of electrification of residential heating on loading of distribution networks". In: *2011 IEEE Trondheim PowerTech*. 2011 IEEE Trondheim PowerTech. June 2011.
- [214] W. Vickrey. "Counterspeculation, Auctions, and Competitive Sealed Tenders". In: *The Journal of Finance* 16.1 (Mar. 1, 1961), pp. 8–37.
- [215] A. Von Meier. *Electric power systems: a conceptual introduction*. John Wiley & Sons, 2006.
- [216] P. N. Vovos et al. "Centralized and distributed voltage control: Impact on distributed generation penetration". In: *IEEE Transactions on power systems* 22.1 (2007), pp. 476–483.
- [217] D. Wakabayashi. "Uber's Self-Driving Cars Were Struggling Before Arizona Crash". In: *The New York Times* (Mar. 23, 2018).
- [218] J. Walrand. *Probability in Electrical Engineering and Computer Science: An Application-Driven Course*. 1st. Vol. 1. Quoi?, Mar. 2014.
- [219] J. Wang and N. Elia. "A control perspective for centralized and distributed convex optimization". In: *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 3800–3805.

- [220] Y. Wang et al. “Research on resilience of power systems under natural disasters—A review”. In: *IEEE Transactions on Power Systems* 31.2 (2016), pp. 1604–1613.
- [221] Y. Weng, R. Negi, and M. D. Ilić. “Probabilistic Joint State Estimation for Operational Planning”. In: *IEEE Transactions on Smart Grid* PP.99 (2017).
- [222] B. Xu and A. Abur. “Observability analysis and measurement placement for systems with PMUs”. In: *IEEE PES power systems conference and exposition*. Vol. 2. Citeseer, 2004, pp. 943–946.
- [223] Y. Xu et al. “Multi-Timescale Coordinated Voltage/Var Control of High Renewable-Penetrated Distribution Systems”. In: *IEEE Transactions on Power Systems* (2017).
- [224] R. V. Yampolskiy. *AI Is the Future of Cybersecurity, for Better and for Worse*. Harvard Business Review. May 8, 2017. URL: <https://hbr.org/2017/05/ai-is-the-future-of-cybersecurity-for-better-and-for-worse> (visited on 04/29/2018).
- [225] S. Yang and B. Hajek. “VCG-Kelly Mechanisms for Allocation of Divisible Goods: Adapting VCG Mechanisms to One-Dimensional Signals”. In: *IEEE Journal on Selected Areas in Communications* 25.6 (Aug. 2007), pp. 1237–1243.
- [226] M. N. Zeilinger et al. “Plug and play distributed model predictive control based on distributed invariance and optimization”. In: *Conference on Decision and Control*. Florence, Italy: IEEE, 2013.
- [227] B. Zhang, A. D. Dominguez-Garcia, and D. Tse. “A local control approach to voltage regulation in distribution networks”. In: *North American Power Symposium (NAPS)*. IEEE, 2013.
- [228] B. Zhang et al. “An Optimal and Distributed Method for Voltage Regulation in Power Distribution Systems”. In: *IEEE Transactions on Power Systems* 30.4 (July 2015), pp. 1714–1726.
- [229] J. Zhu et al. “Differentially Private Distributed Online Algorithms Over Time-Varying Directed Networks”. In: *IEEE Transactions on Signal and Information Processing over Networks* 4.1 (2018), pp. 4–17.
- [230] M. Zook et al. “Ten simple rules for responsible big data research”. In: *PLOS Computational Biology* 13.3 (Mar. 30, 2017).

Appendix A

Power System Stability

This chapter formed the basis for motivating much of the research in this dissertation, earlier on in the PhD process. As such, it is still relevant as a reference to provide more insight into the trends affecting electricity distribution, as covered more concisely in Chapter 2. It frames the overall notions used to describe efficacy of power systems in Section A.1. It then provides a classification of power systems stability for distribution networks in Section A.2. It ends with a description of how trends in energy systems are impacting power system stability in Section A.3.

A.1 Classic Notions of Efficacy in Power Systems

Throughout the literature on power systems, different concepts and words with varying definitions are used to describe the efficacy of operation of power systems and power quality. Hence, we define the different criteria considered in the power systems and control communities. Here, we adopt and extend some definitions from [126] and [9]. These studies themselves have tried to formally define the different notions and provide the field with advise for terminology. At the highest level of power systems operations, we make a distinction between the terms reliability, security, stability and resilience, as proposed in [126]:

- *Reliability* - refers to the probability of its satisfactory operation over the long run. It denotes the ability to supply adequate electric service on a nearly continuous basis, with few interruptions over an extended time period.
- *Security* - refers to the degree of risk in its ability to survive imminent disturbances (contingencies) without interruption of customer service. It relates to robustness of the system to imminent disturbances, and hence, depends on the system operating condition as well as the contingent probability of disturbances.
- *Stability* - refers to the continuance of intact operation following a disturbance. It depends on the operating condition and the nature of the physical disturbance. Notice

that this is different from the notion of stability in the dynamical systems and control literature [10]. Whenever we refer to stability, we will provide further clarity if the context does not provide sufficient clarification.

- *Resilience* - refers to the capacity to withstand extreme situations (e.g. extreme weather conditions), to prevent the spread of outages, and to recover quickly when outages occur.

Reliability is the overall objective in power system design and operation. To be reliable, the system must be secure. To be secure, the system must be stable, but in addition, it must also be secure against other contingencies that would not be classified as stability problems, such as damage to equipment due to explosive failure of a cable. Security and stability can be further distinguished in terms of resulting consequences. For example, two systems may both be stable with equal margins, but one may be more secure because consequences of instability are less severe. Stability is hence an integral component of a power system security and reliability assessment. Grid resilience is an objective of increasing importance, as severe weather is the leading cause of power outages in the United States. Between 2003 and 2012, an estimated 679 widespread power outages occurred due to severe weather, with an estimated annual cost average of \$18-33 billion [69]. As a result of climate change and the frequency of severe weather events increasing, the study of resilience has been getting more attention. See [220] for a review of research on resilience in power systems, and [6] for a discussion on definitions.

A.2 Power System Stability in Distribution Networks

Many major blackouts caused by power system instability have illustrated the importance of this phenomenon. Historically, in transmission networks transient instability has been the main focus of the industry's attention concerning system stability. With the evolution of power systems through continuing growth in interconnections, use of new technologies and controls, and the increased operation in highly stressed conditions, different forms of system instability have emerged.

The overarching definition as proposed in [126] is: *Power system stability is the ability of an electric power system, for a given initial operating condition, to regain a state of operating equilibrium after being subjected to a physical disturbance, with most system variables bounded so that practically the entire system remains intact.* This definition applies to an interconnected power system as a whole. Power system stability is similar to the stability of any dynamical system, and therefore has fundamental mathematical underpinnings. Stability of an electric power system is a property of the system motion around an equilibrium set, that is the initial operating condition.

Power systems are subjected to a wide range of disturbances, small and large. Small disturbances in the form of load changes occur continually. They must also be able to ride through numerous severe disturbances, such as the tripping of a generator or short circuit on

a transmission line. A large disturbance may lead to structural changes due to isolation of faulted elements, and may cause damage to many vital pieces of equipment. For example, in a cascading event, a fault on a critical element can lead to protective relaying, which causes variations in power flows, bus voltages, and rotor speeds. Voltage variations can then trigger voltage regulators and switch off other equipment. Loads can either switch off or be severely damaged, depending on their individual characteristics. An unstable system condition could lead to cascading outages and a shutdown of a major portion of the power system.

In contrast to transmission networks, there is not yet a clear all-encompassing definition of stability for the distribution grid. Common practice is to use transmission level definitions. However, structural differences between transmission and distribution, varying characteristics of (aggregated) loads in distribution networks, typically lower inertia of distributed generation, intermittency of distributed renewable energy resources, and presence of power electronic inverters all give rise to new forms of dynamic phenomena and hence new equilibria between supply and demand in distribution networks [188]. Therefore, it is essential to assess stability of distribution network stability in dynamic and transient conditions and clearly classify when stability is met.

Power system stability is essentially a single problem. However, because of high dimensionality and complexity of stability problems, it helps to make simplifying assumptions to analyze specific types of problems using an appropriate degree of detail of system representation and appropriate analytical techniques. We revisit the power system stability as defined for transmission networks in [126], and define its distribution network equivalent. We distinguish between four different notions of stability:

1. *Frequency Stability* - the ability of a power system to maintain steady frequency following a severe system upset resulting in a significant imbalance between generation and load. It depends on the ability to maintain/restore equilibrium between system generation and load, with minimum unintentional loss of load. Instability that may result occurs in the form of sustained frequency swings leading to tripping of generating units and/or loads.
2. *Voltage Stability* - the ability of a power system to maintain steady voltages at all buses in the system after being subjected to a disturbance from a given initial operating condition. It depends on the ability to maintain/restore equilibrium between load demand and load supply from the power system. Instability that may result occurs in the form of a progressive fall or rise of voltages of some buses.
3. *Synchronization* - the ability of synchronous machines of an interconnected power system to remain in synchronism after being subjected to a disturbance. For Distributed Generation and the interfacing inverters, it depends on the ability to track the right operating point to maintain/restore equilibrium between the input energy of the generator and the desired output frequency and angle of the network. Instability that may result occurs in the form of increasing angular swings of some generators leading to their loss of synchronism with other generators. There are five conditions that

must be met before the synchronization process takes place. The source (generator or sub-network) must have equal line voltage, frequency, phase sequence (for three phases), phase angle, and waveform to that of the system to which it is being synchronized. As such, some frame synchronization often to be an overarching problem, as it indirectly addresses frequency and voltage stability [63, 187].

4. *Waveform Stability* - the ability of a power system to operate at a near-perfect waveform at its fundamental frequency and to maintain this waveform after being subjected to a harmonic disturbance. Instability that may result occurs in the form of imperfect sine waves that can cause serious damage and progressive tripping of various electric components in the network [9, 96].

A.3 Impact of Trends on Power System Stability

In this section we discuss critical challenges that arise with the trends mentioned in the Introduction. We explain the impact on the different notions of power system stability.

High penetration of distributed generation

In order for California to reach its 2020 goals to have 33% of its loads served by renewable energy sources, it will heavily rely on the installation of distributed PV and wind turbine installations connected to mid- and low-voltage distribution networks. The CAISO estimates that by 2020 12000 MW, which is roughly 20% of today's average consumption, is generated by distributed PV sources.

- *Intermittency* - The short term stochastic intermittency effects of PV generation due to clouds or of wind due to wind intermittency can create imbalances leading to unstable frequency trajectories, that can have its impact on the larger network. [34] shows that in order to balance the intermittency effects of distributed PV generation on local frequency stability one needs over 80% balancing capacity (with respect to installed PV capacity) versus 8% for state level balancing. California has already put in place regulation to enforce the installation of balancing capacity with the installation of new PV panels. Under- or over-generation of distributed generators can cause voltage restrictions to be violated in congested or underserved areas of the network. Distributed generators need to resynchronize to the network in order to inject power around intermittency.
- *Forecast errors* - Large deviations from the forecasted distributed generation are a big concern for transmission networks, as slower spinning generation is scheduled based on forecasts. Deviations lead to imbalance, which if not compensated for can lead to over/under-generation and voltage instability. For islanded microgrids such phenomena can easily lead to outage because of the inherently lower inertia in such systems.

- *Ramping* - For California estimates are that the high penetration of PV will lead to dramatic ramping phenomena in the net load (consumption minus local generation) of the overall system. For 2020 a projected 13 GW ramp up of net load in just 3 hrs around sunset is expected. As for now, there is a big shortage of fast regulating sources that can supply this ramp (as solar power diminishes) and prevent major potential outages.
- *Lack of inertia* - AC power systems robustness to uncertainty and disturbances relies on sufficient inertia, which for transmission networks is provided by large spinning generators. As distributed generation takes over the role of large fuel-powered spinning generation the grid's inertia is inherently lowered. As a result, the system is more sensitive to perturbations, and deviations from operating equilibria are more likely to lead to cascading events and outages.

These phenomena highlight the tremendous impact that DG will have on the local power system stability. With penetration currently increasing at a fast rate, we will likely face new phenomena that motivate the careful study of and design for stability in distribution networks.

Electrification and diversified dynamic behavior of loads

Increasingly more aspects of our economy are electrifying. Transportation will transition to electrified vehicles that charge on the power grid. Smart applications with embedded computation integrate sophisticated electronics that diversify the impact on the grid. Demand response programs use flexibility of demand and control the dynamics of loads in novel ways (directly or indirectly) to regulate frequency for the transmission network. As more critical processes rely on electronic interfaces the vulnerability of the grid is increased, sometimes leading to stricter requirements on reliability, and thus on stability and power quality.

- *Nonlinear loads* - some loads cause the current to vary disproportionately with the voltage during each cyclic period. These are classified as nonlinear loads, and the current taken by them has a non-sinusoidal waveform. When there is significant impedance in the path from the power source to a nonlinear load, these current distortions will also produce distortions in the voltage waveform at the load.
- *Intermittency* - It is inherently impossible to know when all loads will be on or off. If large loads (as related to the inflowing power of a (sub)network) switch on or off this can create local imbalances and potentially cascade to voltage sag/sweep and destabilize and desynchronize the system.
- *Forecast error* - Large deviations from the forecasted load consumption are a big concern for transmission networks, as slower spinning generation is scheduled based on forecasts. Deviations lead to imbalance, which if not compensated for can lead to

over/under-generation and voltage instability. For islanded microgrids such phenomena can easily lead to outage because of the inherently lower inertia in such grids.

- *Electric Vehicles* - Plug-in electric vehicles create new load peaks by connecting around similar time slots [5]. Having a high penetration of controlled (vs uncontrolled) EVs could have detrimental effects to local frequency (imbalance) and voltage (congestion) [155].
- *Demand Response* - By exploiting the flexibility of demand from buildings, ancillary services are now provided to the grid for regulating frequency. Examples of loads involved are HVAC systems (using fast fans [139] , using slower heating systems [156]), thermostatically controlled loads and batteries with stored energy. These services are now used for transmission frequency regulation, but could as well be used for local balancing. Applicability depends on the available control actuation bandwidth, which varies greatly from seconds to minutes depending on the available actuation and control infrastructure and load dynamics. It is yet to be understood what the effect of demand response schemes is on local stability (especially if many neighboring resources are triggered in parallel).
- *Critical loads* - Critical loads such as hospitals and data centers rely on increasingly stricter power supply standards to run critical processes and electronic equipment. Apart from their new and possibly unreliable effect on the grid, new equipment also brings more vulnerability, and as such potentially stricter stability requirements on distribution networks.

Availability of novel measurement and control actuation devices

There is not yet alarming attention for dealing with stability in distribution networks. But as measurement and actuation devices are installed and model accuracy improves, we form a better idea of the phenomena and impact on stability and components in the distribution system as it occurs in practice. This can help setting thresholds for what phenomena or scenarios are worth preventing or stabilizing, and how actuation can help doing so in the most effective and efficient way. Apart from the potential benefits, exploiting sensors and control actuation can also lead to new stability problems.

- *Advanced Metering Infrastructure* - The rapid deployment of “smart meters” allows utilities to gain more understanding of electricity demand. Such data gathered historically can reveal patterns that can help shape plan the dimensioning and operation of the grid more effectively and less conservatively. It can also serve as a source of information to design data-driven demand response programs to actively adjust demand to avoid network contingencies and optimize network objectives. With the measuring of household demand come new concerns of privacy, demanding engineering tradeoffs between utility of data for network operation and the protection of sensitive information.

- *Synchrophasors* - Phasor Measurement Units are able to measure and compare voltage and current phasors at high temporal resolution, and have made their way into the standard practice in the operation of transmission systems, acting as a means of assessing real-time stability and understanding transient dynamics. In the last decade, synchrophasors have been developed for distribution networks, providing new options for actively managing distribution systems with diverse resources and growing complexity [142]. As this technology is currently experimented with, the applications are still becoming more apparent, including using it for predictive maintenance, fault and event detection, state estimation and power flow and voltage regulation.
- *Power electronic inverters* - To interface increasingly sophisticated loads and distributed generators, power electronic devices are implemented. The potential impact of these devices on local frequency stability, voltage stability and synchronization of microgrids has not studied well until recently [187, 64, 63]. Inverters have shown to have very diverse impacts on the harmonic waveform as introduced in different ways (series, parallel), and can also impact the voltage stability [199]. Apart from their disturbing effects, inverters also provide controllability. They can regulate the injected/absorbed real and reactive power through droop control, and modulate voltage magnitude, frequency and phase angle for synchronization [125]. Moreover, with its circuitry one can adjust the overall impedance to compensate for the nonlinear effect on the current/voltage waveform coming from the load/generator.
- *Modeling errors* - Planning of the distribution grid uses numerous types of models for transient, dynamic and steady state configurations. Most models are parametrized based on equipment data originating from a Geographic Information System (GIS) database. Errors in these databases can lead to large simulation errors, which may not be caught during existing validation schemes. As penetration of renewables and other stochastic sources increase, these small errors may become more pertinent and mask integration and future planning issues. Incorrect parameter estimation is another source of risk that can lead to modeling errors.
- *Control actuation errors* - Fast local compensation by power electronic equipment to respond to voltage variations might lead to overcompensation, causing voltage drops for the noncompensated parts of the network. Widespread use of compensation equipment may even become a voltage stability issue, if not designed and implemented carefully [9]. Unaccounted delays in control architecture or communication might yield dissatisfactory operating results for control strategies. Unforeseen interaction of different timescales or control loops in the dynamics might also impede on the effectiveness of control schemes [64].
- *Measurement errors* - Wear and fraud in measurement devices can be fed back to a control loop and cause tremendous actuation errors leading to various dynamic phenomena, that can affect the voltage stability and harmonic wave form. Misplaced or

mis-integrated sensors could provide a monitoring system or control loop with inadequate information. Packet losses could harm the feasibility of a computation needed for safe control and operation. All such scenarios can be detrimental depending on the safety-criticality of the equipment and control strategies that rely on the information gathered by measurements.

Aging infrastructure

In many countries, the energy market and infrastructure have been heavily deregulated over the last few decades. Depending on the pricing models of a market, utility companies are facing economic challenges to meet the technical investments needed for upgrading the grid for future reliability, given the scenarios for increased distributed generation and electrification of the economy. For transmission networks the safety criteria and shared reliance on the infrastructure have thus far often guaranteed the necessary political and economic buy-in to innovate. For distribution networks, the investment responsibility is often more complicated and depending on the area. An area with relatively little economic income is less likely to find the resources to upgrade its distribution network. Nonetheless, the users of such networks most often have a free hand to connect new loads and generators that affect and likely increase the impact on the network. Decreasing reliability could likely increase the investment of users in their network. Also, the possible effects of one controlled area on others might yield debate between users and stimulate a system wide solution [9]. For a given distribution network, one should carefully consider the practical issues related to economic perspective, load and generation impact on the network, and the necessary political involvement/agreement to realize the necessary upgrade for reliable and stable operation. As such, the aging infrastructure is an overarching trend that influences all the other challenges listed.

Appendix B

Proofs for Theorems and Lemmas Local Differential Privacy

Proof of Theorem 7.3.1

Proof: To show that the proposed mechanism \mathcal{A} in Algorithm 1 has the promised privacy guarantee, we first show that \mathcal{A}_i is locally ϵ_i -differentially private for node i . To this end, we study the quantity of interest

$$\frac{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M) \in \mathcal{Z}\}}{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M) \in \mathcal{Z}\}} \quad (\text{B.1})$$

of node i . In the proof we use random variables Z_i^k, V_i^k , for $k = 1, \dots, K$, to denote the output of the local mechanism \mathcal{A}_i (c.f. Eq. (7.7)) with input $(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M)$, and use $Z_i'^k, V_i'^k$, for $k = 1, \dots, K$, to denote the output of the local mechanism \mathcal{A}_i with input $(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M)$ where the two local problems \mathcal{P}_i and \mathcal{P}'_i have distance $d(\mathcal{P}_i, \mathcal{P}'_i) \leq 1$. With these notations, we can rewrite the above expressions as

$$\frac{\mathbb{P}\{Z_i^1 = z_i^1, V_i^1 = v_i^1, \dots, Z_i^K = z_i^K, V_i^K = v_i^K\}}{\mathbb{P}\{Z_i'^1 = z_i^1, V_i'^1 = v_i^1, \dots, Z_i'^K = z_i^K, V_i'^K = v_i^K\}}$$

for some realizations $z_i^k, v_i^k, k = 1, \dots, K$.

It is straightforward to show that if $Z_i^k = Z_i'^k, V_i^k = V_i'^k$ for all $k = 1, \dots, K$, then we have $Z_j^k = Z_j'^k, V_j^k = V_j'^k$ for all $j \neq i$ and for all $k = 1, \dots, K$. For the simplicity of notation, we use Z^k and V^k to denote the tuple (Z_1^k, \dots, Z_M^k) and (V_1^k, \dots, V_M^k) , respectively. Furthermore notice that V^k is a deterministic function of Z^k , as shown in the step 3 of Algorithm 1. Hence

we have

$$\begin{aligned}
 & \frac{\mathbb{P}\{Z_i^1 = z_i^1, V_i^1 = v_i^1, \dots, Z_i^K = z_i^K, V_i^K = v_i^K\}}{\mathbb{P}\{Z_i'^1 = z_i^1, V_i'^1 = v_i^1, \dots, Z_i'^K = z_i^K, V_i'^K = v_i^K\}} \\
 &= \frac{\mathbb{P}\{Z^1 = z^1, V^1 = v^1, \dots, Z^K = z^K, V^K = v^K\}}{\mathbb{P}\{Z'^1 = z^1, V'^1 = v^1, \dots, Z'^K = z^K, V'^K = v^K\}} \\
 &= \frac{\mathbb{P}\{Z^1 = z^1, \dots, Z^K = z^K\}}{\mathbb{P}\{Z'^1 = z^1, \dots, Z'^K = z^K\}} \\
 &= \prod_{k=1}^K \frac{\mathbb{P}\{Z^k = z^k | Z^t = z^t, t < k\}}{\mathbb{P}\{Z'^k = z^k | Z'^t = z^t, t < k\}}
 \end{aligned}$$

where $z^k, v^k, k = 1, \dots, K$ are realizations of the random variables. Now we analyze the term

$$\frac{\mathbb{P}\{Z^k = z^k | Z^t = z^t, t < k\}}{\mathbb{P}\{Z'^k = z^k | Z'^t = z^t, t < k\}}$$

for each $k = 1, \dots, K$.

According to Step 1 of Algorithm 1, z_i^k is determined by μ_i^{k-1} and δ_i^k , in other words $Z^k - m^{k-1} - (Z_0, \dots, Z^{k-1})$ forms a Markov chain where the random vector m^k denotes the Lagrangian $(\mu_1^k, \dots, \mu_M^k)$ in the algorithm. Consequently, it holds that

$$\frac{\mathbb{P}\{Z^k = z^k | Z^t = z^t, t < k\}}{\mathbb{P}\{Z'^k = z^k | Z'^t = z^t, t < k\}} = \frac{\mathbb{P}\{Z^k = z^k | m^{k-1} = \mu^{k-1}\}}{\mathbb{P}\{Z'^k = z^k | m'^{k-1} = \mu'^{k-1}\}}$$

where m^k and m'^k denote the random Lagrangians with inputs $(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M)$ and $(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M)$, respectively. Furthermore, it can be checked easily that m^{k-1} and m'^{k-1} are completely determined by Z^t, Z'^t for $t < k$. Since we have $Z^t = Z'^t = z^t$ for all $t < k$, we also have $\mu^{k-1} = \mu'^{k-1}$.

With the above observation, we can continue our derivation

$$\begin{aligned}
& \frac{\mathbb{P}\{Z^k = z^k | Z^t = z^t, t < k\}}{\mathbb{P}\{Z'^k = z^k | Z'^t = z^t, t < k\}} = \frac{\mathbb{P}\{Z^k = z^k | m^{k-1} = \mu^{k-1}\}}{\mathbb{P}\{Z'^k = z^k | m'^{k-1} = \mu'^{k-1}\}} \\
&= \frac{\mathbb{P}\{Z^k = z^k | m^{k-1} = \mu^{k-1}\}}{\mathbb{P}\{Z'^k = z^k | m'^{k-1} = \mu'^{k-1}\}} \\
&= \prod_{j=1}^M \frac{\mathbb{P}\{Z_j^k = z_j^k | m_j^{k-1} = \mu_j^{k-1}\}}{\mathbb{P}\{Z_j'^k = z_j^k | m_j'^{k-1} = \mu_j'^{k-1}\}} \\
&= \prod_{j=1}^M \frac{\mathbb{P}\{g(\mathcal{P}_j, \mu_j^{k-1}) + \delta_j^k = z_j^k\}}{\mathbb{P}\{g(\mathcal{P}'_j, \mu_j'^{k-1}) + \delta_j^k = z_j^k\}} \\
&= \prod_{j=1}^M \frac{\exp(-\|z_j^k - g(\mathcal{P}_j, \mu_j^{k-1})\|/\sigma_j^k)}{\exp(-\|z_j^k - g(\mathcal{P}'_j, \mu_j'^{k-1})\|/\sigma_j^k)} \\
&\leq \prod_{j=1}^M \exp(\|g(\mathcal{P}_j, \mu_j^{k-1}) - g(\mathcal{P}'_j, \mu_j'^{k-1})\|/\sigma_j^k) \\
&= \exp(\|g(\mathcal{P}_i, \mu_i^{k-1}) - g(\mathcal{P}'_i, \mu_i'^{k-1})\|/\sigma_i^k) \\
&\leq \exp(\Theta_i/\sigma_i^k)
\end{aligned}$$

with Θ_i defined in (7.12). Finally, we obtain

$$\begin{aligned}
& \frac{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}_i, \dots, \mathcal{P}_M) \in \mathcal{Z}\}}{\mathbb{P}\{\mathcal{A}_i^K(\mathcal{P}_1, \dots, \mathcal{P}'_i, \dots, \mathcal{P}_M) \in \mathcal{Z}\}} \\
&= \prod_{k=1}^K \frac{\mathbb{P}\{Z^k = z^k | Z^t = z^t, t < k\}}{\mathbb{P}\{Z'^k = z^k | Z'^t = z^t, t < k\}} \\
&\leq \prod_{k=1}^K \exp(\Theta_i/\sigma_i^k) \\
&= \exp\left(\sum_{k=1}^K \Theta_i/\sigma_i^k\right)
\end{aligned}$$

This proves the privacy guarantee in (7.11). ■

Proof of Lemmas

Proof of Lemma 7.3.1: We show that for the quadratic problem, the optimization problem in (7.14) is in fact independent of the value of μ . To see this, we define $\tilde{h}_i := h_i - \mu$

and $\tilde{h}'_i := h'_i - \mu$. The optimization problem (7.14) can be rewritten as

$$\begin{aligned}
 & \text{maximize } \Theta_i := \|z^* - z'^*\| & (\text{B.2}) \\
 & \text{s. t. } z^* \in \text{KKT}(H_i, \tilde{h}_i, C_i, c_i, w_i, 0) \\
 & \quad z'^* \in \text{KKT}(H'_i, \tilde{h}'_i, C'_i, c'_i, w'_i, 0) \\
 & \quad a_1 \|H_i - H'_i\| + a_2 \|\tilde{h}_i - \tilde{h}'_i\| \\
 & \quad + a_3 \|C_i - C'_i\| + a_4 \|c_i - c'_i\| \leq 1 \\
 & \quad \text{with variables } z^*, z'^*, H_i, H'_i, \tilde{h}_i, \tilde{h}'_i, \\
 & \quad C_i, C'_i, c_i, c'_i, w_i, w'_i
 \end{aligned}$$

where we used the fact $\tilde{h}_i - \tilde{h}'_i = h_i - h'_i$. It can be seen that the above problem formulation is independent of μ , which proves the claim. \blacksquare

Proof of Lemma 7.3.2: Using Lemma 7.3.1, we consider the case with $g(P_i, \mu = 0) = \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h^T z$ (we omit the index i for simplicity). Define

$$z^* := \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h^T z$$

and

$$z'^* := \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H' z + h^T z$$

where H, H' satisfies $\|H - H'\|_2 \leq 1$.

The optimality of z^* and feasibility of z'^* implies

$$\langle H z^* + h, z'^* - z^* \rangle \geq 0,$$

and the optimality of z'^* and the feasibility of z^* implies that

$$\langle H' z'^* + h, z^* - z'^* \rangle \geq 0.$$

Manipulations of the above two inequalities yield

$$\langle H' z'^* - H z^*, z^* - z'^* \rangle \geq 0$$

which can be rewritten as

$$\langle H' \Delta + (H - H') z^*, \Delta \rangle \leq 0$$

where $\Delta := z^* - z'^*$. This implies that

$$\Delta^T H' \Delta \leq \Delta^T (H' - H) z^*$$

Using the fact that $\Delta^T H' \Delta \geq \lambda_{\min}(H') \|\Delta\|^2$ and $\Delta^T (H' - H) z^* \leq \|\Delta\| \|z^*\| |\lambda_{\max}(H' - H)|$, we conclude that

$$\begin{aligned} \|z^* - z'^*\| &\leq \frac{|\lambda_{\max}(H' - H)|}{\lambda_{\min}(H')} \cdot \|z^*\| \\ &\leq \frac{1}{\lambda_{\min}^{(i)}} G_i \end{aligned}$$

where the last inequality uses the assumption that $|\lambda_{\max}(H' - H)| = \|H - H'\|_2 \leq 1$, $\|z^*\| \geq G_i$, and the minimum eigenvalue of H' is lower bounded by $\lambda_{\min}^{(i)}$ ■

Proof of Lemma 7.3.3: Using Lemma 7.3.1, we consider the case with $g(P_i, \lambda = 0) = \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h^T z$ (we omit the index i for simplicity). Define

$$z^* := \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h^T z$$

and

$$z'^* := \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h'^T z$$

where h, h' satisfies $\|h - h'\|_2 \leq 1$.

Since $H \succ 0$, we define $H = D \cdot D^T$ with D invertible and rewrite the above expression as

$$\begin{aligned} z^* &= \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} z^T H z + h^T z \\ &= \operatorname{argmin}_{z \in \mathbb{C}} \frac{1}{2} \|D^T z + D^{-1} h\|^2 \end{aligned}$$

Let $v = D^T z$. The optimization problem above becomes

$$v^* = \operatorname{argmin}_{v \in \bar{\mathbb{C}}} \frac{1}{2} \|v + D^{-1} h\|^2 ,$$

where we define $\bar{\mathbb{C}} := \{D^T z \mid z \in \mathbb{C}\}$. Similarly we have

$$v'^* = \operatorname{argmin}_{v \in \bar{\mathbb{C}}} \frac{1}{2} \|v + D^{-1} h'\|^2 .$$

Hence v^* (res. v'^*) can be seen as the projection of the point $-D^{-1} h$ (res. $-D^{-1} h'$) onto the set $\bar{\mathbb{C}}$. Since \mathbb{C} is convex, we know that $\bar{\mathbb{C}}$ is also convex, and it holds that

$$\|v^* - v'^*\| \leq \|D^{-1}(h - h')\| .$$

Using $z = (D^T)^{-1} v$, we obtain

$$\begin{aligned} \|z^* - z'^*\|_2 &\leq \|D^{-1}\|_2 \cdot \|D^{-1} \cdot (h - h')\|_2 \\ &\leq \|D^{-1}\|_2^2 \cdot \|h - h'\|_2 \\ &\leq \frac{1}{\lambda_{\min}(H)} \cdot \|h - h'\|_2 \\ &\leq \frac{1}{\lambda_{\min}^{(i)}} \end{aligned}$$

where the last inequality uses the assumption that $\|h - h'\| \leq 1$ and minimum eigenvalue of H is lower bounded by $\lambda_{min}^{(i)}$. ■