

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Energy Management in Microgrids for Electricity Access

Permalink

<https://escholarship.org/uc/item/5n82g7dv>

Author

Lee, Jonathan Temple

Publication Date

2022

Peer reviewed|Thesis/dissertation

Energy Management in Microgrids for Electricity Access

by

Jonathan Temple Lee

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Energy and Resources

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Duncan Callaway, Chair

Professor Alexandra von Meier

Professor Claire Tomlin

Spring 2022

Energy Management in Microgrids for Electricity Access

Copyright 2022
by
Jonathan Temple Lee

Abstract

Energy Management in Microgrids for Electricity Access

by

Jonathan Temple Lee

Doctor of Philosophy in Energy and Resources

University of California, Berkeley

Professor Duncan Callaway, Chair

This thesis responds to the related imperatives of transitioning to low-carbon electricity systems and increasing global access to energy. It shows that microgrids and decentralized electricity systems are economically and technically capable of providing high levels of electricity access, and argues that incorporating active participation of electricity “prosumers” into energy management systems enables more efficient electricity resource management.

Chapter 1 quantifies the tradeoffs between costs and electricity for autonomous solar and battery systems across sub-Saharan Africa, finding that on average these autonomous systems can achieve high levels of reliability at a cost of on the order of 10 USD cents per ‘9’ of reliability. Moreover, it shows that these costs could drop to as low as 3 cents per 9 as battery costs decline, and that decentralized systems are cost-competitive with legacy grids across much of the continent.

Chapter 2 proposes a load management system to manage electricity consumption in community microgrids with solar photovoltaics and battery storage while accounting for forecast uncertainty. It uses stochastic, model-predictive control techniques to set consumption limits during periods of low solar availability and high-demand. Simulation experiments show the management technique improves system reliability and consumer benefits from electricity through fewer interruptions and better electricity availability to high value uses.

Chapter 3 studies optimal pricing and peer-to-peer energy trading systems in microgrids with 100% renewable energy sources. It promotes a utility-maximization framework from which prices arise from exchanging electricity under scarcity, in contrast to standard marginal-cost based pricing that breaks down in 100% renewable systems. It further proposes a negotiation algorithm for peer-to-peer energy transactions and proves its convergence to optimal exchanges.

Chapter 4 extends the algorithm from Chapter 3 to a more broadly applicable system for optimizing power exchanges in microgrids and larger power grids in using forward markets

and real-time controls. This approach, based on a decentralized optimization technique known as the Alternating Direction Method of Multipliers (ADMM), uses price-based coordination and independent agents in an iterative bidding procedure. Its equilibrium is a welfare-maximizing dispatch that solves the non-linear and non-convex power flow equations. This system preserves individual privacy, efficiently incorporates network congestion and voltage constraints, is highly scalable, and is robust in practice to model error. In addition to the forward market, the chapter introduces an agent-based feedback control system that continues to optimize power exchanges in real time.

The thesis concludes with a brief summary and directions for future research.

In memory of George F. Lee

As my grandfather, he made me believe I could do anything and always encouraged me to take my education as far as it would go. He taught me to strive for what is important: to use my work to serve others, to earn the respect of my colleagues and peers, and to extend a hand of care to anyone in the community who needs it. He passed in 2017, during the second year of my graduate studies. Both before and after his death, I turn to his spirit whenever I feel like I'm facing the impossible. With his strength, and the same strength shared by my mother Barbara Lee, I've been able to forge a way through the challenges I've faced as an academic and as a practitioner. This thesis is dedicated in honor of his legacy and in gratitude for supplying me with the tools to get here.

Contents

Contents	ii
List of Figures	iv
List of Tables	v
Introduction	1
1 The cost of reliability in decentralized solar power systems in sub-Saharan Africa	7
1.1 Introduction	8
1.2 A framework for the cost of reliability	9
1.3 Commodity prices and the costs of decentralized solar electricity	11
1.4 Alternative methods for estimating LCOE	15
1.5 Discussion	18
1.6 Methods	19
1.A Supplementary Note: Optimal system designs	25
1.B Supplementary Note: Sensitivity to daily load profile shape	27
1.C Supplementary Note: Electricity tariff, reliability, and electrification data . .	30
1.D Supplementary Note: Cost of reliability results by country	32
2 Non-intrusive load management under forecast uncertainty in energy constrained microgrids	35
2.1 Introduction	36
2.2 Decision problem	37
2.3 Microgrid simulation model	46
2.4 Computational experiments	47
2.5 Conclusions	51
3 Pricing and energy trading in peer-to-peer zero marginal-cost microgrids	52
3.1 Introduction	54
3.2 Centralized Welfare Maximization Approach	55
3.3 Peer-to-peer Negotiation	61

3.4	Computational Experiments and Simulations	69
3.5	Conclusions	73
3.A	Appendix: Non-ideal battery modeling	74
3.B	Appendix: Constant price-elasticity utility functions	76
4	Decentralized optimal power dispatch and secondary control with agent-based ADMM	78
4.1	Introduction	78
4.2	Agent models, notation, and basic graph-theoretic concepts	84
4.3	Economic power dispatch with ADMM	95
4.4	Economic secondary control with online ADMM	115
	Conclusion	123
	References	125

List of Figures

1.1	Average Service Availability Index (ASAI) for countries in SSA	10
1.2	Levelized cost of energy of Tier 5 decentralized systems in present and future scenarios	12
1.3	Statistical relationship of LCOE and FDS in SSA	13
1.4	Spatial distribution of reliability premium	14
1.5	Predictive power of mean insolation on LCOE	16
1.6	Cost difference between decentralized solar LCOE and grid tariffs	17
1.7	Statistical relationship between optimal system design and FDS for Tier 5 systems at current costs	26
1.8	Additional cost incurred if holding the storage to solar ratio fixed	27
1.9	Daily load profile shapes	28
1.10	Statistical relationship of LCOE and FDS for different load profiles	29
1.11	Statistical relationship between optimal storage to solar ratio and FDS for different load profiles for Tier 5 systems at current cost	30
2.1	Receding horizon control system	38
2.2	Alternative interpretations of a 2-scenario forecast with a single battery over a horizon of three time-periods.	42
2.3	Key performance metrics across trials.	49
2.4	Objective <i>ex post</i> and <i>ex ante</i> values with mean load.	50
3.1	Optimal profiles for the centralized approach	60
3.2	The standard convergent cobweb model.	61
3.3	Convergence with step-limiting constraint	62
3.4	Trial procedure flowchart.	68
3.5	Effect of γ on the number of iterations to convergence	70
3.6	Number of iterations to converge with varying battery capacity.	71
3.7	Centralized and P2P algorithm price profiles for the special instance of considerable difference in welfare.	72

List of Tables

1.1	Economic assumptions	12
1.2	Tariff, reliability, and electrification data by country	31
1.3	Mean and standard deviation of LCOE for all locations sampled in each country under the current and future economic assumptions	32
2.1	Activity parameters (time in minutes)	48
2.2	Timing results (seconds)	51
3.1	Welfare difference statistics for different time horizons.	71
3.2	Welfare difference statistics for the special instance considered in Section 3.4. . .	71

Acknowledgments

First and foremost, I want to thank my primary advisor, Professor Duncan Callaway. Duncan cares deeply about his students and is a brilliant researcher. In working with him I've felt trusted and supported every step of the way. As my career goals shifted throughout my PhD, he adapted his approach and expectations very fluidly to support me while keeping doors open and helping me to explore new options.

I am also deeply grateful to Professor Sascha von Meier who took me under her wing when I was an undergraduate, gave me my first real window into life as a researcher, and played a huge role in inspiring me to pursue a PhD. Sascha connects easily with her students and not only makes it fun to talk about the power grid, but generously shares her love of music, good food, and the mountains with everyone. A whole generation of students in the Berkeley energy and power systems community can attribute some of the best parts of their Berkeley experience to her. As a scholar and researcher, her material and physical understanding of the power grid and the people and institutions who make it work is unmatched. I will always be seeking her guidance on doing research that matters throughout my career.

I also want to thank Professor Claire Tomlin for advising me through the EECS Master's program and welcoming me into the systems and control theory community. Claire is the best technical lecturer I've had the privilege of learning from, and in every meeting about my research I was always blown away by how quickly she grasped the point I was trying to get across and responded with insights. Claire especially gave me the confidence not to shy away from some of the more theoretical work I was interested in and made it accessible to me.

Thanks to Jalel Sager for being a friend, mentor, and intellectual inspiration since we started working together in 2014. Intellectually, Jalel has guided me through critical engagement with economic thought beyond the math, understanding institutions and organizations as ecosystems, and above all what it takes to get something done in practice. As someone preceding me in ERG by about a generation, his guidance on how to approach the PhD and navigate each of the milestones has been instrumental.

My peers were the bedrock of my graduate school community. Veronica Jacome deserves special thanks for being both a senior peer mentor and an intellectual contributor to my perspective on electricity systems and energy access far beyond what is in this thesis. My conversations with her are among my most treasured from my time at ERG. I also want to acknowledge my fellow students in the electric power systems and energy access community: Keith Moffat, Isa Ferrall, Anna Brockway, Phillippe Phanivong, Mohini Bariya, Rodrigo Henriquez-Auba, Paty Hidalgo-Gonzales, Jose Daniel Lara, Ciaran Roberts, Laurel Dunn, Will Gorman, and Annelise Gill-Wiehl. Thank you from the bottom of my heart for all your time discussing ideas and collaborating on classes, projects, and papers. Thank you to my ERG cohort of 2016: Bodie Cabiyo, Esther Shears, Anusha Hossain, Hilary Yu, Isa Ferrall, Phillippe Phanivong, Gordon Bauer, Emma Tome, Alex Dolginow, and Adam Hanbury-Brown. Our discussions throughout Master's seminar and all of our potlucks, backyard BBQ's, and outdoor adventures brought the spirit that makes the experience what it is. I

also want to thank some ERGies who are my senior and both inspired me and gave me sage advice along the way: Diego Ponce de Leon Barido, Grace Wu, Ranjit Deshmukh, John Dees, Michael Cohen, and Dimitry Gershenson.

Thanks to Professor Isha Ray for guiding me through social science research and patiently working with me as I experimented with incorporating this perspective into my work and sitting on my qualifying examination committee. Thanks to Professor Daniel Kammen for being the first one to open ERG's door to me when I was an undergraduate and encouraging me to get over any hesitation and get involved with research. Dan is a tireless champion of ERG students and nobody creates as many opportunities for ERGies to get out into the world and have an impact. Thanks to Kay Burns for going to bat for the students whenever we needed her. Thanks to the "old guard" of ERG faculty, who established the department, pioneered interdisciplinary energy and environmental research, and are role models to all of us for service to our local, national, and global communities. In particular, Professors Richard Norgaard, John Harte, and the late Gene Rochlin all continued to show their love of the department and the students throughout my time at ERG, and were instrumental in fostering a close-knit and multi-generational community that welcomes in new members every year.

Lastly, I want to thank my family. Thanks to Taryn Elliott, the love of my life, for going through all of this with me and being patient and loving every step of the way. Thanks to my parents Barbara Lee and Marc Couacaud for their unwavering confidence and for raising me with curiosity. Thanks to my grandparents Kit and George Lee for always rooting for me and showing pride in what I do. And finally, thanks to all of my friends and family not mentioned by name for being with me throughout this journey.

Introduction

This thesis is motivated by two closely related imperatives: transitioning to 100% renewable electricity systems and universal access to electricity. It argues that microgrids, decentralized electricity systems at the community and household scale, should play a significant role in responding to these two challenges with resource management systems that jointly address the efficient production and use of electricity. This dissertation contributes both techno-economic analysis to support this argument and proposes mechanisms for using decentralized management and control systems to design and operate clean and accessible electricity networks.

Before proceeding, it is important to bring into focus contemporary developments in two areas of study that are essential to the context of this work. The first is how we define and understand electricity access, and the second is how advancements in electric power conversion technology are enabling and necessitating new approaches to managing electricity systems. This thesis focuses on systems that sit at the interface between these social and physical contexts, particularly how microgrids with renewable energy sources can be designed and automatically controlled to foster increased electricity access. With the exception of the first chapter, the thesis is primarily devoted to proposing novel approaches and evaluating their technical performance using engineering methods. These studies demonstrate the technical efficacy of microgrids for providing reliable electricity and show how automatic controllers can further increase the reliability and economic efficiency of decentralized electricity systems. These contributions should be thought of as a pre-requisite for comprehensive interdisciplinary evaluation and experimentation in the field and not as a substitute for it. As a whole, this dissertation stops short of a critical analysis of electricity access and leaves out some important implementation details which connect to both the social and the physical contexts. These contexts are summarized here to help complete the picture.

Electricity access has historically been understood as a binary: whether one has a formal connection to a legally recognized electricity system or not. The limitations of this binary metric became apparent as the costs of solar photovoltaics underwent massive (more than 10x) cost reductions c. 2005-2015, leading to a surge in access to informal decentralized electricity, while simultaneously the energy access literature increasingly emphasized the low reliability of many formal electricity systems in the global south. In 2015, the United Nations adopted 17 Sustainable Development Goals (SDGs) to formalize their development agenda through 2030, replacing the 8 Millennium Development Goals (MDGs) adopted in

2000. Whereas energy access itself was not directly one of the MDGs, it was specifically articulated in SDG #7, “Ensure access to affordable, reliable, sustainable and modern energy for all.” This followed from the climate crisis moving to the front and center of the MDG’s environmental sustainability goal and scholars and practitioners demonstrating the foundational and cross-cutting role of energy in other goals aimed at poverty, education, and health. Also in 2015, the Energy Sector Management and Assistance Program (ESMAP), which is governed in part by the World Bank, was established and published the Multi-Tier Tracking Framework (MTF), titled, “Beyond Connections: Energy Access Redefined,” [12]. This framework uses a concept of tiers to define energy access not as a binary, but as a spectrum with multiple criteria each for household electricity, commercial and community electricity, and cooking fuels. This framework not only adds detail and nuance to how access is defined through specific affordability, reliability, capacity, legality, safety, and other criteria, but it includes for what and how energy is used in the rationale. This marks a shift in the international development community from understanding access as something that is strictly about energy supply, to a dialectic that includes the use of energy.

This shift can be connected back to several related threads in international development, common pool resource management, and rural sociology. At a high level, the approach taken in the ESMAP MTF is consistent with part of what Ribot and Peluso argue for in their 2003 Theory of Access, which is to understand access to a resource as the *ability* to derive benefits from it rather than as a system of legal or formal rights to it, and to shift focus away from the enforcement of access rules to the means of access [99]. At the core of Ribot and Peluso’s Theory are the social power relationships that shape material access to resources. ESMAP’s MTF does not address these relationships; however, recent scholarship on electricity access is beginning to take this more seriously through a growing emphasis on gender relationships and the study of informal and illegal electricity systems, which can apply pressure for the UN and World Bank to evolve their frameworks accordingly.

Expanding the study of energy access, particularly electricity access, to a study that includes the usability of electricity at its core opens up a systems perspective that understands the supply of a resource and the use of it as existing in feedback with each other. An example of this is the common pool resource (CPR) literature, which addresses the governance and management of shared resources at the community to regional scale, and considers these as coupled “provision” and “appropriation” processes [87]. Most of the classic CPR studies c. 1990-2000 and their successors are squarely oriented towards international development and human-environmental systems under the threat of ecological collapse, such as forests, fisheries, range lands, wildlife habitats, and watersheds. The majority of the global population with low levels of electricity access live in rural areas that typically depend on CPR systems, for example the fishing settlements on the islands of Lake Victoria, which became one of the most active regions for decentralized electricity in East Africa c. 2015-2020. The potential for dialogue between researchers and practitioners of electricity and development, demographics, and the underlying natural resource systems in these rural regions is still largely untapped, but through multiple disciplines studying the same communities, some CPR methods and perspectives have begun to appear explicitly in the electricity access literature [40]. As this

perspective is adopted, the study of electricity access shifts away from least-cost methods of generating electricity to the joint study of the provision and use of the resource.

If we reflect these shifts in the international development community back onto electricity systems that are not considered to have an access issue, we see a striking parallel: there is also a movement from strictly supply-side resource management to a system with active participation by consumers with flexibility to both adjust their electricity use and also to participate in provisioning electricity. The term “prosumer” has entered the lexicon of electric power systems, and mechanisms to integrate them and their distributed energy resources (DERs) into the electricity markets and control systems of existing grids are the current focus of many applied research and pilot projects around the world. In the emerging new regime, microgrids made up of prosumers play a central role as partially autonomous systems responsible for both provisioning and appropriating the electricity resource. This is contrast to the existing regime, which takes demand as fixed and focuses solely on provision and minimizing supply-side costs.

The legacy management approach of building for high reliability with inflexible demand reflects an abundance of electricity in regions with high access, but comes at a cost. This tradeoff between costs and reliability for decentralized systems with inflexible loads in low-access regions, specifically across sub-Saharan Africa, is the topic of Chapter 1. The chapter quantifies the tradeoff and finds that if battery storage costs continue to decline, then much of the continent can meet inflexible demand and *improve* reliability with stand-alone systems at levelized costs of energy on par with retail electricity tariffs and cheaper than diesel fuel. A similar analysis focusing on the United States finds that up to 7% of households could find it economically advantageous to disconnect from the grid if they accept slightly less reliable service [41]. These were both published before increasingly frequent and intense wildfires began inducing utilities such as California’s Pacific Gas and Electric to implement public safety power shutoffs that lower reliability while simultaneously raising rates for hardening the system in response to wildfires. The studies both show that the economics are not one-sided enough to tip the balance either way towards centralized or stand-alone systems, and that the demand flexibility plays a significant role in the economics.

As we consider electricity grids with greater emphasis on flexibility from consumers, it is an open question how much demand flexibility should be used to address periods of scarcity while still being considered high-access. Beyond the low-hanging fruit of optimizing loads like electric vehicle charging and heating and cooling systems, additional flexibility sacrifices some of the benefits of electricity. In high-access regions, the discourse around this question tends to focus on what consumers will tolerate and how to incentivize participation. However, as end-use energy in high-access regions is increasingly electrified in response to the climate emergency, as the electricity supply is increasingly based on variable renewable energy, and as the physical grid is threatened by climate change and extreme weather, the dialog may shift to tradeoffs between costs and flexibility under scarcity. In low-access regions, it is widely taken as given that consumers will not only tolerate mechanisms to regulate or limit consumption, and that these are in fact best practices if not essential. This stark contrast highlights global electricity injustice, but also reveals how both high-access

and low-access regions are grappling with a similar need to engage the demand-side more actively in managing electricity as a resource system. In this thesis, Chapter 2 develops a technical approach to address the problem of appropriating electricity under scarcity using dynamically adjusted load limits in a community microgrid, as do Chapters 3 and 4 with peer-to-peer transactional systems.

Moving to the physics of electricity, recent advancements in power converters and the ongoing efforts to standardize and bring these into the mainstream are critical to enable the transition to renewable and decentralized microgrids. The issue centers around controlling many different resources in concert to balance the supply and demand of electricity in real time. The issue is that physical inertia, the backbone of real-time stability of power grids, is disappearing as we shift to renewable energy. This inertia is provided by the spinning turbines of generators in thermal power plants, which are mostly fossil-fuel based and need to be phased out of a low-carbon electricity system. On an interconnected electricity grid, the turbines in all of the generators spin synchronously, and by their inherent electromagnetic physics, act as a stabilizing force on the grid. When the net demand on the system exceeds the supply, the turbines automatically start slowing down, converting their rotational energy to electric power that supplies the load. Likewise, when supply exceeds demand, they start speeding up. This phenomenon buys time, and as a consequence, the rotational speed, or the frequency, of the grid reflects the imbalance between supply and demand. Devices on generators detect that the rotational speed of the generator drops and then increase the power input proportionally, and vice versa. In the 19th century, engineers accomplished this with a pair of spinning balls on a shaft that would mechanically open a valve as they slow down and spin more closely to the shaft, which in turn would let more steam into the generator, thus speeding it up until reaching a new equilibrium. This process, known as “droop control” and “primary control” is the basis for how grids provide stable electrical frequency and balance supply and demand for electricity.

Most renewable energy that is expected to power the low-carbon future either does not have inherent inertia (e.g. solar panels) or has some inertia that is decoupled from the grid (e.g. wind turbines) and does not act as a stabilizing force. Therefore, the question of how to maintain stability in low-inertia systems is a critical question facing power systems engineers. The leading approach at the time of writing is to use power converters to mimic the droop behavior of generators, which is called synthetic or virtual inertia. This requires some fast-acting energy storage to act in place of the kinetic energy of the turbine, which can be in the form of batteries, capacitors, or flywheels, for example. The modern approaches to designing power converters in this way were proposed c. 2005-2010 and large-scale initiatives to evaluate these in practice at national and international scales have been underway since c. 2015.

It turns out that this same fundamental issue has also limited microgrids based on solar photovoltaics and battery storage from being built in a distributed way, meaning with multiple sources distributed in space. There are some limited ways around this with high-bandwidth feedback loops, but these tend to require generation sources to be colocated with wired communication links and are not scalable. Thus, community electrification projects

in areas with low access tend to have a single generation plant supplying many users or they use multiple isolated systems. The former approach comes with additional costs in the distribution network and imposes some limits on adding capacity in response to changing demand. The latter results in excess capacity going unused. Both of these approaches limit the net value proposition of community microgrids.

The introduction of power converters that provide synthetic inertia and droop, beyond meeting a necessary hurdle for transitioning to low-carbon grids, has the potential to be game-changing in microgrids for electricity access. At minimum, it enables a more modular system design with lower network costs that will allow electricity grids to grow organically and adapt more easily to meet changing demands. But it also provides the foundation for much more granular control of how electricity flows on a network and what resources are called upon to meet demand, to the point where electricity transfer can be made transactional, and quantities of power and energy can be exchanged between individuals or groups on the network. The control systems that regulate these flows are called the “secondary” and “tertiary” controls with respect to the “primary” system described above. The tertiary level is economic, typically structured as wholesale markets in liberalized systems or cost minimization problems in vertically integrated ones, and the secondary is typically an automatic rebalancing of the primary system that attempts to track the targets of the tertiary. This thesis develops transactional and price-based tertiary systems for microgrids in Chapters 3 and 4, which are powerful in that they allows entrepreneurial individuals or community groups to monetize excess power generation. More generally, the granular ability to control flow supports a wide range of systems for provisioning and allocating electricity.

In summary, this thesis first evaluates the economics of reliability for stand-alone decentralized systems for energy access in sub-Saharan Africa in Chapter 1, and then proceeds to develop technical solutions to some of the important problems in energy management for decentralized systems in Chapters 2-4. The core contribution of Chapter 1 is to characterize the costs of reliability in detail, but it also shows that with declines in battery costs, highly reliable, fully decentralized solutions will be competitive with centralized grids. This supports the emerging consensus that decentralized systems will play, and already are playing, an important role in expanding electricity access. Chapter 2 address load management under solar power scarcity and forecast uncertainty and shows how a tertiary control system with dynamic load limits can improve access to electricity. It further compares algorithms and introduces new approaches with two-stage stochastic programming and approximate dynamic programming that outperform conventional model predictive control under uncertainty. Chapters 3 and 4 take a more forward-looking perspective towards coordinating resource supply and demand in 100% renewable energy systems. These propose specific peer-to-peer transaction systems and contribute theoretically to this area. The theoretical contributions include showing how optimal power flow for 100% renewable systems can be cast as a decentralized utility maximization problem instead of a cost minimization, uncovering some limitations of conventional pricing in the presence of battery storage, giving an iterative linearization of decentralized optimal power flow that solves the nonlinear and non-convex problem exactly under the condition that it converges, and connecting the economic

dispatch problem to frequency regulation with a novel online secondary control system. Although these chapters are motivated specifically by microgrids, they are relevant to future larger scale systems that embrace broader participation by DERs and active demand-side management.

Chapter 1

The cost of reliability in decentralized solar power systems in sub-Saharan Africa

This chapter was co-authored with Prof. Duncan Callaway, Associate Professor of Energy and Resources at the University of California, Berkeley. The text was previously published in *Nature Energy* in 2018; see [59] for the full citation. The text has been reformatted for inclusion in this thesis.

Abstract

While there is consensus that both grid extensions and decentralized projects are necessary to approach universal electricity access, existing electrification planning models that assess the costs of decentralized solar energy systems do not include metrics of reliability or quantify the impact of reliability on costs. We focus on stand-alone, household solar systems with battery storage in sub-Saharan Africa (SSA) using the fraction of demand served (FDS) to measure reliability, and develop a multi-step optimization to efficiently compute the least-cost system with FDS as a design constraint, taking into account daily variation in solar resources and costs of solar and storage. We show that the cost of energy is minimized at approximately 90% FDS, that current costs increase on average USD 0.11/kWh for each additional “9” of reliability, and that this reliability premium could be as low as USD 0.03/kWh in a plausible future price scenario.

1.1 Introduction

The United Nations’ Sustainable Development Goal #7 describes a major global task: “Ensure access to affordable, reliable, sustainable, and modern energy for all,” [119]. 1.06 billion people remain without access in 2017, and more than half of them live in sub-Saharan Africa (SSA), the geographic focus of this analysis [49]. Historically, electrification from centralized national and regional electric grids was the only path; now, interest and investment in decentralized, stand-alone options from solar home systems (SHS) to community mini-grids is growing [3]. Private capital investment in SHS providers increased from USD 3 million in 2012 to 381 million in 2015-2016, and deployment of solar energy products for electrification has grown similarly [98]. We use “decentralized” to describe systems that are self-sufficient and independent of connection to a larger grid. This definition applies in a range of applications, but here we focus on isolated, household-scale systems. Multiple studies have pointed to the efficacy of solar-diesel-battery hybrid systems that use diesel generators to increase reliability [18, 106]; however, recent work exploring future low carbon energy systems points to increasing reliance on solar-battery systems [15] and high penetrations of SHS [10], which provides motivation to study SHS at high levels of reliability.

The centralized and decentralized paradigms have advantages and disadvantages. Centralized grids in SSA utilize economies of scale [3] to achieve lower costs of energy, but often do not reach rural areas. Even where the grid is present, some households are left “under the grid” because of high connection costs and long wait times for connections [61]. When existing grid service is unreliable, some consumers use SHS as a backup [75]. Decentralized systems can be flexibly and rapidly deployed to meet basic energy needs in many geographies and communities, but high costs and the challenges of a variable resource raise questions about whether they are a viable long-term economical solution for electrification. Nonetheless, the International Energy Agency predicts that 60% of connections needed for universal electricity access will come from decentralized [48], but the eventual outcome will be influenced by cost trends, policy, and regional electrification plans that guide investment. These electrification plans rely on cost minimization models that compare grid expansions to decentralized options [79, 77, 78, 126], which are most frequently solar photovoltaics (PV) with battery storage.

Yet, we currently have limited understanding of the economics of decentralized systems and their reliability across large spatial scales. Many models assume constant per-unit costs of energy for classes of decentralized systems [79, 77, 78, 65, 64], whereas others scale system size by location-specific solar resources but assume constant storage-to-solar capacity ratios [92, 53] or a single reliability level [108, 19]. We are aware of one analysis that investigates the cost of reliability for decentralized systems, but the results are isolated to one location [62]. Ultimately, the centralized vs decentralized debate needs to be answered within the context of least-cost electricity planning tools. Although such tools have been applied in the African context [79, 77, 78, 126, 65, 64, 92, 53, 108, 26, 84, 107, 129], their computational complexity and the lack of available knowledge of decentralized system costs has precluded integration of decentralized pathways into the models.

In this study, we investigate how designing decentralized systems for different levels of reliability affects their cost, how decentralized solar systems compare in cost to national grids when designed for equivalent reliability, and how changing commodity prices can affect these relationships. We find evidence that decentralized solar-battery systems are on the cusp of reaching “grid parity” in both economic and reliability terms in many parts of SSA. Furthermore, using aggressive forecasted technology costs taken from the literature, we show that in the future a large fraction of the continent could be served by decentralized systems with better economics and reliability than the existing grid. We propose a method for identifying optimal PV and battery capacities for decentralized systems using 11 years of location-specific daily solar resource data (at one-degree latitude and longitude resolution) and across a large range of reliabilities (measured in terms of the Fraction of Demand Served, or FDS). We also rigorously quantify the cost to improve reliability at each location and show that this “reliability premium” has strong spatial variability, increasing levelized cost of energy (LCOE) by USD 0.05-0.15/kWh for each order of magnitude improvement in reliability (e.g. from 99 to 99.9%). These results point to the potential for decentralized solar systems to provide very high reliability service at costs that are competitive with existing, often highly unreliable, grid infrastructure. The method we use is computationally scalable and can be used for rapid analysis of different cost assumptions (an open-source implementation is available at <https://emac.berkeley.edu/reliability>). The model can also be integrated into planning models to capture the trade-offs between a spectrum of technology pathways ranging from centralized grid options to fully decentralized systems. We compare our method to existing approximations, and find that simpler methods provide a good – though not perfect – estimate at FDS below 99%, but that these methods are inadequate to quantify LCOE variation at higher FDS.

1.2 A framework for the cost of reliability

To quantify reliability, we define and use the Fraction of Demand Served (FDS), which is, over the analysis period, the sum of all energy delivered divided by the sum of all energy demanded. An FDS of 1 indicates perfect reliability. FDS is similar to the more common Average Service Availability Index (ASAI) [13], which describes the fraction of time that service is available, and can be estimated by the World Bank Enterprise Surveys for many countries, shown in Figure 1.1 [113]. FDS is effectively the ASAI weighted by demand: ASAI does not distinguish between outages when demand is high vs. low, but outages during periods of high demand will have a greater impact on FDS. Note also that $FDS = 1 - ESP$, where ESP is the Energy Shortfall Probability used by Lee, Soto, and Modi [62]. Studies employing the optimization software HOMER [108, 62] use a single year of measured radiation to construct estimates of FDS; we use eleven years of remote sensing data to increase the robustness of the estimate. Our approach could alternatively use ASAI or the System Average Interruption Duration Index (SAIDI) [13], but we use FDS as it relates more directly to energy demanded instead of time. We only model solar resource

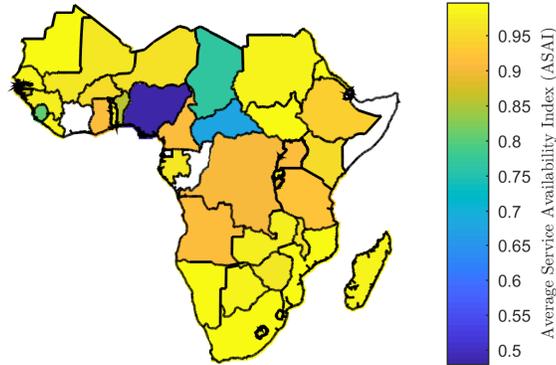


Figure 1.1: Average Service Availability Index (ASAI) for countries in SSA

ASAI for national grids as reported by World Bank Enterprise surveys [113] (see data in Supplementary Table 1.2). Data are unavailable for blank countries.

driven outages; evidence from the field suggests technical failures are much less significant than resource outages [36].

We quantify FDS because it has an important impact on the cost and utility of decentralized solar systems, and because there is significant variation in the reliability of grids in SSA, as illustrated in Figure 1.1. A 2014 simulation for a micro-grid in Mali showed that as the fraction of demand served approaches 100% there is a log-linear relationship between reliability and added cost to achieve that reliability, and that the optimal ratio of battery to photovoltaic capacity changes significantly for different target reliabilities [62]. The optimal battery-to-PV ratio also depends on local weather patterns and demand patterns. For example, additional battery capacity might be the best choice at a location that has frequent but intermittent cloud cover, but would be ineffective in regions with prolonged rainy seasons. The software HOMER uses simulation and searches over a large parameter space, making it adequate for single-location studies [108, 62] but unsuitable for evaluating costs on continental scales or for use as a sub-problem in a larger optimization.

Our approach preserves the nature of the solar array vs battery bank capacity trade-off by computing an isoreliability curve (equivalent to a Pareto optimal frontier), i.e., the set of all system designs that achieve a desired FDS for a particular location. With an isoreliability curve, the cost-minimizing capacity of storage and solar given their costs (battery, PV module, racking, charge controller, etc.) can be found by simple line search. We construct isoreliability curves through simulation for each location in SSA at 1 degree latitudinal and longitudinal resolution and for the FDS of interest and store the results. This approach enables modelers to then analyse a large number of cost scenarios with minimal computation,

enabling detailed optimization at high spatial-temporal resolution and geographic scales, bridging the gap between detailed local models and wide area studies (see Methods: “Cost optimization” for how to adjust economic assumptions).

An isoreliability curve depends on hourly consumption patterns; night-time load requires more storage than daytime load. However, because the curves are constructed per-unit daily load, they are independent of average load (kWh/day). This independence enables the method’s scalability. Because location-specific load shape data are not available for the regions we study, and because our solar insolation database provides daily (rather than hourly) data, we present our results assuming constant load throughout the day (or equivalently, with a load factor of 1). To test this assumption, we performed sensitivity analyses using different load curves: 1) constant load, 2) all load is during 6 PM to 6 AM and constant during that time, 3) all load is during 6 AM to 6PM and constant during that time, and 4) load follows a representative profile containing an evening peak that was empirically measured on a rural micro-grid in Uganda. These tests show that the qualitative results we present here are robust to other load shapes. Specifically, we find that across FDS, constant load yields approximately the same costs as a measured load profile from Uganda with a night-time peak, and that concentrating demand at night raises LCOE by up to USD 0.15/kWh while concentrating demand during the day lowers LCOE by USD 0.10/kWh (see Supplementary Note 1.B). We also assume that load is identical every day; incorporation of stochastic load models into this framework is an open technical challenge.

1.3 Commodity prices and the costs of decentralized solar electricity

We developed scenarios from electricity access “Tier 5” defined by the World Bank’s Energy Sector Management Assistance Program (ESMAP) and computed the LCOE across SSA using the economic assumptions in Table 1.1. This tier is the highest level of access for household and productive uses of electricity and includes explicit metrics for reliability, capacity, and consumption [12]. In particular, assuming FDS and ASAI are equivalent because of constant load, Tier 5 requires FDS of 95% by specifying that service is available 159 out of 168 hours per week. Figure 1.2 shows current LCOE varies by about USD 0.15/kWh across SSA, and that potential future cost reductions are greater in magnitude than current spatial variation. Lower component costs in the future scenario reduce the average LCOE, and also lower its coefficient of variation, showing that cost declines have disproportionate impact on higher cost locations. The future scenario entails aggressive cost reductions that could plausibly be realized c. 2025, but we do not forecast the exact time-frame of cost reduction (see Methods). The location of high- and low-cost areas is similar to earlier work with simplified cost models [108], but we find generally higher costs in the current scenario and that cost reduction effects outweigh current spatial variation (Figure 1.2).

Table 1.1: Economic assumptions. See Methods for discussion and justification

	2017	Future (c. 2025)
<i>Solar costs</i>		
Modules plus DC Balance of System	1.00	0.50
Charge controller	0.20	0.10
Total \$/W	1.20	0.60
<i>Battery costs</i>		
Total \$/kWh	400	100
<i>Load costs</i>		
Inverter	0.30	0.15
Soft costs plus AC Balance of System	1.00	0.50
Total \$/W	1.30	0.65
<i>Additional economic assumptions</i>		
O&M costs	\$100/kW peak load/year	
Project length	20 years; battery replacement at 10 years	
Annual discount rate	10%	

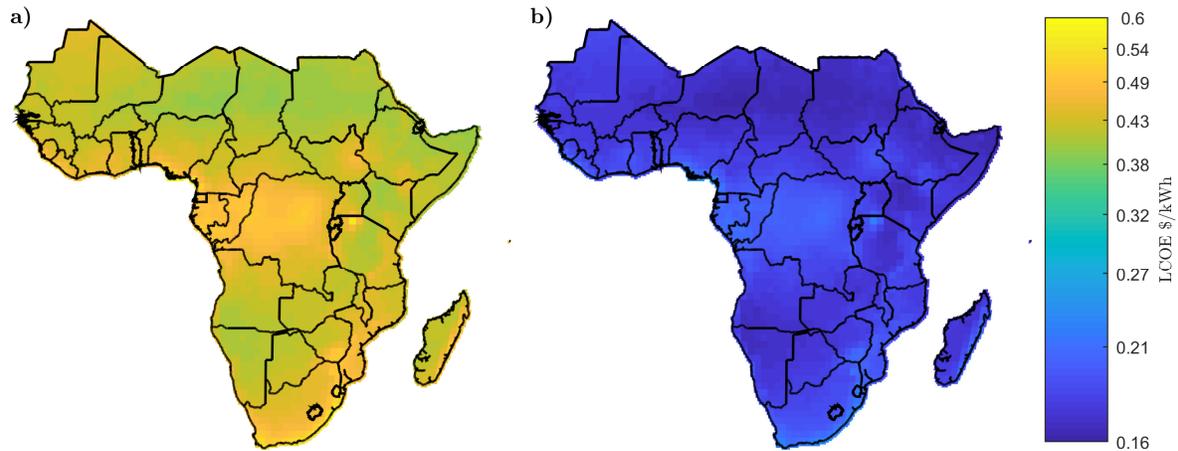


Figure 1.2: Levelized cost of energy of Tier 5 decentralized systems in present and future scenarios

Panel a) shows the present LCOE, while panel b) shows a future scenario entailing a 75% reduction in battery and 50% reduction in solar module and balance of system costs. Tier 5 refers to systems serving 8.2 kWh/day at 95% FDS with 2 kW peak capacity. Please see the Methods for additional economic assumptions. A nonlinear colour scale is used to better show spatial variation in the future scenario. Data are presented aggregated by country in tabular form in Supplementary Table 1.3.

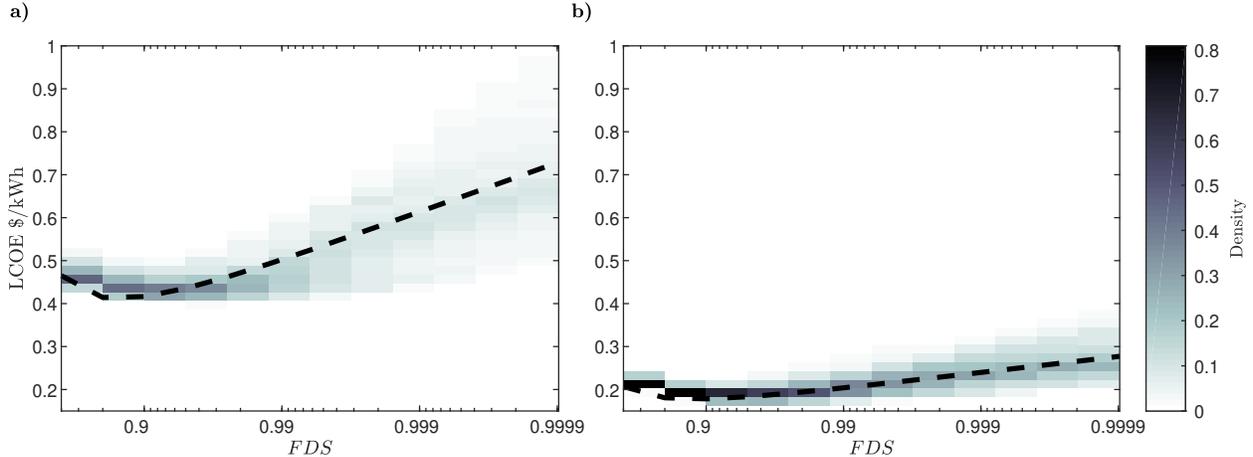


Figure 1.3: Statistical relationship of LCOE and FDS in SSA

The plots show the density of computed LCOE at different FDS for each location in SSA sampled at 1 degree latitudinal and longitudinal resolution. As in Figure 1.2, panel a) shows present costs, while panel b) shows a future scenario. Moving to the right approaches perfect reliability and the black line shows the general trend through a least-squares fit for all locations in SSA to a single regression model given in equation (2). The regression yields $a = -0.11$, $b = 0.18$, $c = 0.088$ with $R^2 = 0.61$ for the current costs, and $a = -0.037$, $b = 0.081$, $c = 0.047$ with $R^2 = 0.57$ for the future costs.

To understand the trends in costs incurred to achieve a desired reliability, we show the density of LCOE across a range of reliabilities using all locations under current and future cost scenarios (Figure 1.3). The figure indicates three important results: 1) LCOE increases logarithmically as FDS approaches one, but has a minimum slightly above 90% reliability; 2) reducing component costs has a disproportionate impact on reducing the premium for high levels of reliability and serves to flatten the LCOE to reliability curve; 3) spatial variance in LCOE increases at higher reliability. This logarithmic scaling, which was previously predicted by a probabilistic model that approximates isoreliability curves [16], simplifies decision-making for sizing reliability and estimating the cost of reliability. The LCOE minimum arises because there are constant fixed costs associated with the system (e.g. inverter and wiring), but they are spread over fewer kWh consumed as FDS declines. This implies that it is not economical – on an LCOE basis – to design systems with less than 90% FDS, given these assumptions; this minimum is approximately stationary in different cost scenarios. Synthesizing Figure 1.1 with Figure 1.3, the current LCOE of decentralized solar varies approximately USD 0.2/kWh in the range of FDS observed on the grid.

The possible future cost scenario indicates that the reliability premium declines and has less spatial variability as component costs decline. The decline in variance in LCOE is greater than what we would expect as a statistical implication of the lower mean LCOE, so the reduction in component costs causes a disproportionate reduction in the reliability premium for high cost areas.

We find that the logarithmic growth in LCOE as FDS approaches one, as well as the

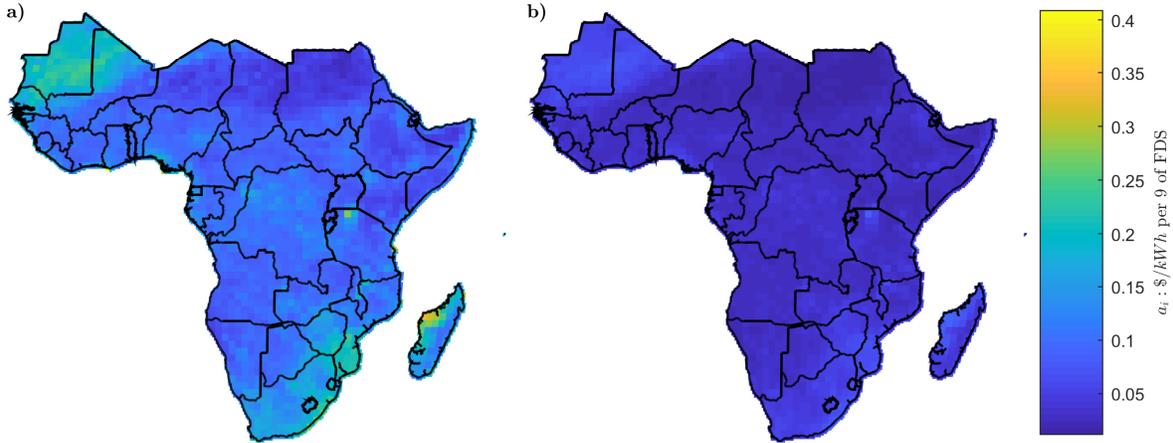


Figure 1.4: Spatial distribution of reliability premium

The premium is given by the coefficient a_i in the least-squares fit to a model for each individual location in SSA, and represents the cost in USD per kWh per additional “9” of FDS at each location i . Note that there is a small area with high premiums currently, and that future premiums are more uniform in space. Data are presented aggregated by country in tabular form in Supplementary Table 3.

existence of the reliability minimum, are well-captured by the parameterized relationship for each location in (1.1) and across locations in (1.2).

$$LCOE_i = -a_i \frac{\log_{10}(1 - FDS)}{FDS} + b_i \frac{1}{FDS} + c_i \quad (1.1)$$

$$LCOE = -a \frac{\log_{10}(1 - FDS)}{FDS} + b \frac{1}{FDS} + c \quad (1.2)$$

In (1.1), i indicates a particular location sampled at 1 degree latitudinal and longitudinal resolution across SSA. The parameters a_i , b_i and c_i are calculated using least-squares regression to each location. For high reliabilities, a_i gives the reliability premium at location i : for every “9” of reliability, LCOE increases by approximately a_i USD/kWh (Figure 1.4). The reliability premium is driven by additional storage and solar capacity needed to ensure capacity during low solar resource (see Supplementary Note 1.A for additional discussion on the optimal system size), however we did not observe any patterns in the distribution of outages or component costs across locations (for example storage cost dominating in one region and solar costs dominating in another). At current costs, most of SSA has a reliability premium of USD 0.05-0.15/kWh per 9 of FDS, though with high variation. The future reliability premium could be as low as USD 0.03/kWh in most of SSA.

1.4 Alternative methods for estimating LCOE

A motivating hypothesis of this analysis is that we can improve upon methods that estimate the cost of decentralized solar using a fixed ratio of solar PV capacity to battery storage capacity. To test this, we check how well mean annual insolation alone predicts our calculated LCOE (Figure 1.5). Mean insolation explains roughly 80 percent of the variation in LCOE for FDS of 0.9 (panel a) of Figure 1.5) and further loses its predictive power relative to our model as FDS increases. Intuitively, this loss in predictive power is because reserve capacity to account for temporal variation in solar resource is more significant at higher FDS and starts to determine costs more than mean insolation. This suggests that methods such as the one we propose are increasingly necessary at high FDS. At lower FDS, one could use a simple linear model driven by annual insolation. However, that model would need to be parameterized with output data from a model such as the one presented here, and it could still incur significant errors, especially at low annual insolation levels (see Figure 1.5).

We also find that the optimal capacity of storage and solar is more variable across locations at higher FDS, implying that it is especially important to compute the optimal system design at higher FDS. We calculate the cost penalty from using a sub-optimal system design by first computing costs holding the ratio of storage to solar capacity fixed as desired FDS changes. We then compare this cost to the optimal cost at different FDS. The penalty incurred for the current cost scenario increases with FDS to around USD 0.10/kWh in many regions, though the spatial variation is significant (see Supplementary Note 1). We conclude that, while the inaccuracies of a simple model (based on mean annual insolation and an approximate storage to solar ratio) may be acceptable for estimating the cost of decentralized solar systems at FDS below 99%, at higher FDS it is necessary to use an optimization that accounts for local weather patterns on a daily scale. Referring to Figure 1.1, we see that many countries in SSA have reliability less than 99%. This is the range in which simpler estimates of LCOE, based on constant storage to solar ratios, are relatively close to the estimates from our model. Our more detailed model remains important for several reasons. First, results from our detailed model are needed to validate simpler ones at different reliabilities, and this contributes an understanding of the threshold at which it is appropriate to use one model versus another. Second, and perhaps more importantly, the model is needed in cases where reliability higher than current national grids is desired, for example in the planning of future power systems that have reliability on par with the rest of the world (e.g. the United States reports ASAI at greater than 99.9% in 2015 [120]).

Decentralized solar LCOE and grid tariff comparison

We now compare the LCOE from decentralized solar to the grid using equivalent performance metrics, namely FDS. We computed ASAI for most countries in SSA using the World Bank Enterprise Surveys [113] (Figure 1.1). These surveys record the frequency and duration of outages reported by businesses, though we note that these numbers are typically 6-7 times higher than those reported by utilities [111]. Treating FDS and ASAI as equivalent

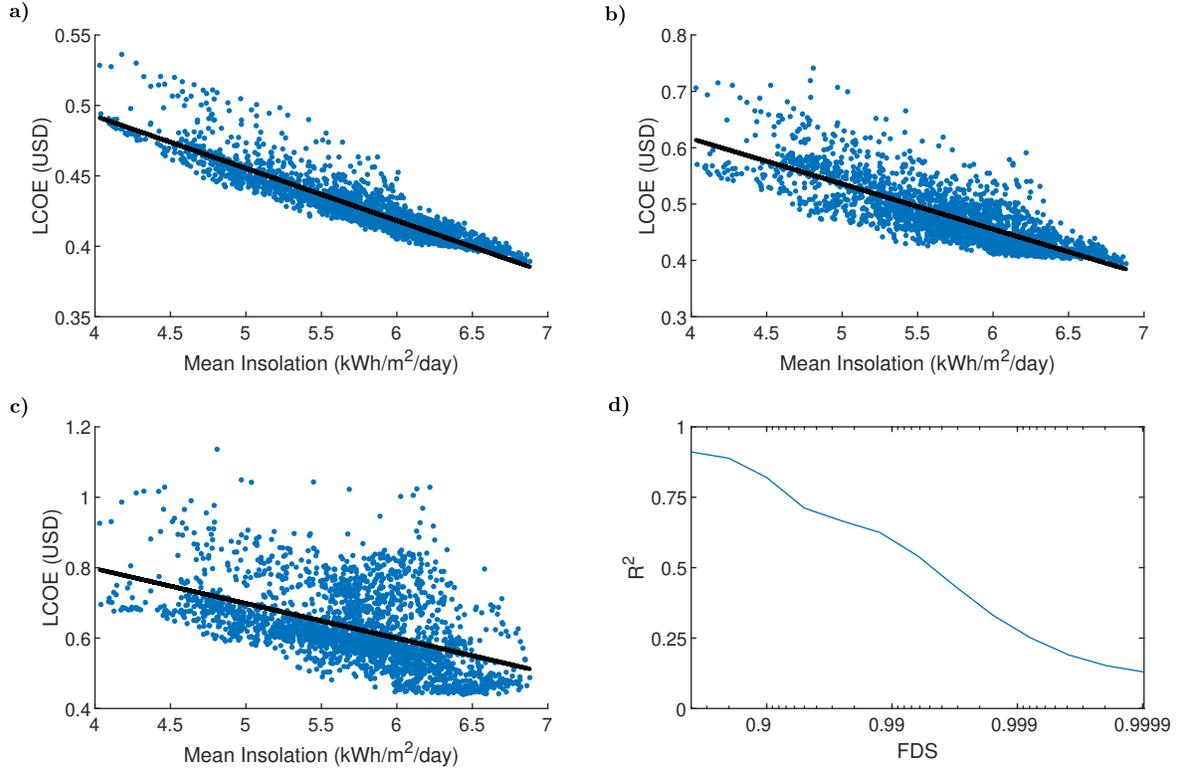


Figure 1.5: Predictive power of mean insolation on LCOE

Panel a) displays our calculated LCOE and mean insolation at each location for FDS of 0.9, with a least-squares, linear regression model of LCOE onto mean insolation indicated by the black line, while b) and c) show the same for FDS of 0.99 and 0.999, respectively. Panel d) shows how the coefficient of determination for this regression deteriorates at high FDS. At low FDS, mean insolation is a good predictor of the LCOE calculated by the isoreliability optimization model. At high FDS, temporal variability in the solar resource starts to drive costs more than the mean insolation, and estimates from mean insolation no longer predict the results of the isoreliability optimization.

lent under our constant load assumption, we compute the LCOE of a decentralized system that provides the ASAI reported for the grid. This enables us to compare the cost differential between decentralized solar LCOE and grid tariffs [114] at approximately equivalent quality of service (Figure 1.6). Though grid tariffs are often both directly subsidized and cross-subsidized between customers, and do not accurately reflect the cost of service [64], comparison to tariffs facilitates understanding where utilities might face competition from decentralized solar and for understanding how decentralized solar compares to business as usual. Additionally, the grid reliability may be unacceptably low in some regions; however, adjusting for reliability provides a more appropriate comparison for countries with reliability above 90% (see Figure 1.1). Grid tariffs and outage rates are reported in Supplementary Note 1.C.

Figure 1.6 shows that in some areas, particularly in West and East Africa, the costs of

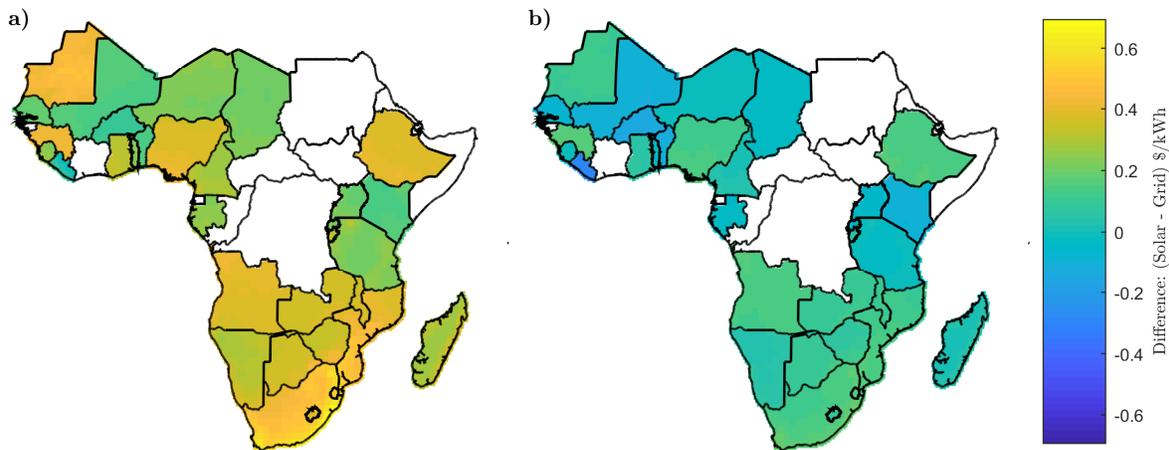


Figure 1.6: Cost difference between decentralized solar LCOE and grid tariffs

Panel a) shows the current cost difference, while panel b) shows the difference under a future cost scenario (as in Figure 1.2 and Figure 1.3). Blank countries are those for which grid tariff or grid reliability data was unavailable. Under current cost structures (see Supplementary Table 1.2 for grid tariffs [114]), only 0.2% of the shaded area is cheaper than the grid, and 0.3% is less than USD 0.05/kWh more expensive than the grid; but under possible cost declines, 28% of the area becomes cheaper than the grid, and 35% is less than USD 0.05/kWh more expensive than the grid. Data are presented aggregated by country in tabular form in Supplementary Table 4.

decentralized solar are approaching grid parity at equivalent reliability, and in the future, could become cheaper than current grid tariffs. A benchmark study for comparing the current cost of solar against the grid [108] finds that including the estimated cost of grid extension in addition to grid tariffs results in one third of the population of Africa being most cost-effectively served by solar, but with the vast majority of those in rural areas. Our results show a current cost difference much less favourable to solar because we do not estimate a cost of grid extension; however, we project that future low costs for solar and storage could enable decentralized systems to threaten the utility business model in many countries, even in urban areas where grid extension is not necessary. Many of the countries where decentralized solar is competitive are also those with relatively low rates of electrification, suggesting that they could be ideal locations for decentralized electricity solutions. In particular, Mali, Liberia, Uganda and Rwanda all have low electrification rates [112] and relatively low cost difference between decentralized solar and the grid. Results by country are available in Supplementary Note 1.D.

There are many factors not analysed here that determine the cost of grid service to the customer, including subsidy, connection fees, and usage; and there are also complex factors determining the cost for the utility to extend service. However, this suggests that particular national utilities could face increasing competition from decentralized solar on the individual household scale, and that certain countries have grid tariff structures and solar

resource characteristics that make decentralized solar a competitive option.

1.5 Discussion

Our first major conclusion is that as solar and battery costs decline, decentralized systems with high reliability meeting ESMAP Tier 5 criteria are likely to become cost-competitive with the grid in a large portion of SSA. This cost parity with the centralized grid is within reach both because of declining costs in solar and storage, and also because centralized systems have significant costs to offset losses (both technical and non-technical) and to build and maintain transmission and distribution infrastructure. The magnitude of cost declines are of course strongly dependent on assumptions input to the model; these assumptions can be explored in an online, open-source version of the model (<https://emac.berkeley.edu/reliability>).

Many countries where cost parity is likely also have low rates of electrification. Our results therefore highlight the risk that the standard aid agency pathway of funding large grid development projects (which have long lead times and complex regulatory process) could be stymied by private decisions to build decentralized systems (whose lead times are extremely low and require little to no government involvement).

Low reliability premiums present an additional challenge to the centralized grid paradigm – while centralized grid reliability is outside customers’ control, in decentralized systems customers can influence their reliability with both upfront decisions as well as real-time curtailment decisions. For example, in many parts of our study region, a customer considering a system with 7900 hours of service per year (90% reliability) could add another 790 hours of service (to 99% reliability) for less than USD 0.10/kWh. This cost could lower to USD 0.03/kWh if aggressive cost declines in PV batteries are realized. From a planning perspective, it would be useful to compare these reliability premiums to the grid to direct investment; however, the premium for the grid is not well understood because of its complexity. This is an important area for further research. While low reliability premiums for decentralized systems coincide with low LCOE in parts of East Africa and regions just south of the Sahara Desert, in general the premium is heterogeneously distributed and does not always coincide with regions with low LCOE, thus it is important to consider the metrics separately.

There are a few important directions for additional investigation. First, different load profiles result in different LCOE estimates; we used a constant load profile for our study, and though we demonstrated that system costs are similar for real customers on a system in Uganda with a night-time peak demand, other realistic load profiles may lead to different costs. We report further analysis on the impact of different load profiles in the Supplementary Note 1.B, where we find that more consumption at night increases the cost of decentralized systems. We do note, however that our general conclusions about the reliability premium are robust to load shape. However, large commercial and industrial systems will have high costs associated with the inverter, and in this case some level of power sharing (to leverage load diversity) or direct grid connections may be essential. Second, our financial models focus on

social cost and as such they do not include overhead for managing payment schemes, such as pay-as-you-go systems³, which bring energy access within reach in cases where the up-front costs are too large. These costs will need to be better understood and eventually factored into analyses such as ours. Third, there is wide variation in reported solar installation costs across SSA [50]; the drivers for this heterogeneity needs further understanding and incorporation into planning models. Fourth, further research is needed into how declining costs of solar and storage will impact grid tariffs; modelling these costs could push grid parity further into the future. Finally, though preliminary analysis shows that technical failures in modern SHS are insignificant compared to resource outages in FDS [36], technical failure modes need significantly more investigation.

In the long run, if the costs we modelled can be paired with rigorous assessments of the societal benefits to electricity and reliability [17], our work enables aid agencies and governments to make informed decisions about if, when and where they should rely on decentralized electrification pathways to meet reliability and development goals. However, detailed models are no substitute for the input of the end users of electricity, and it may well be that markets for decentralized solar systems will blossom long before rigorous conclusions about their benefits can be made. Centralized planning models must take into account the informal process of energy decentralization that will likely emerge in the coming years and decades.

1.6 Methods

Solar and load modelling

To model solar production, we use eleven years of daily average insolation incident on a horizontal surface at each location with a one degree latitudinal and longitudinal resolution. These data were obtained from the NASA Langley Research Center Atmospheric Science Data Center Surface meteorological and Solar Energy (SSE) [121] web portal supported by the NASA LaRC POWER Project. The data span January 1, 1995 to December 31, 2005. To compute the mix of solar and battery storage necessary to supply power throughout the day, we model the system dynamics on an hourly time scale, which requires up-sampling the daily average insolation to hourly average insolation. To do this, we calculate the hourly extraterrestrial horizontal insolation as a function of day of year and location and scale it so that the sum of hourly insolation equals the measured daily insolation using the definitions and solar geometry equations from Chapters 1 and 2 of Duffie and Beckman [31]. This scaling factor, representing cloud cover and atmospheric attenuations is called the clearness index, and our method introduces the assumption that this is constant throughout the day, which in general is not correct [31]; however, given the data available, it is the simplest assumption and suffices to capture the approximate daily profile and account for seasonal variation in sunrise and sunset times.

Computation of isoreliability curves

The isoreliability curves, or the set of solar and storage capacities that meet a specified FDS, is computed through simulation with an hourly time step. We use the isoreliability curve to represent the technical constraints imposed by the physics of the system. At a high level, the approach is to compute these constraints offline through simulation, so that they may be used in an online cost minimization.

The simulation uses the following dynamics, where C_s and C_b are solar array and battery storage capacities in units of kW and kWh per kWh of daily load. I_n and L_n are insolation and load at hour n , in units of per unit full sun and kW per kWh daily load, respectively. which are input vectors described above, ΔP_n is the excess power (or deficit if negative) that is charging or discharged from the battery at hour n , SOC_n is the energy stored in the battery at hour n , U_n is the unmet load at hour n , N is the number of hours in the simulation, and FDS is the Fraction of Demand Served. Note that with an hourly time step, kW and kWh are interchangeable units.

$$\Delta P_n = C_s I_n - L_n \quad (1.3)$$

$$SOC_{n+1} = \max(0, \min(C_b, SOC_n + \Delta P_n)) \quad (1.4)$$

$$\begin{aligned} U_n &= \max(SOC_{n+1} - SOC_n - \Delta P_n, 0) \\ &= \max(L_n - C_s I_n - SOC_n, 0) \end{aligned} \quad (1.5)$$

$$FDS_N = 1 - \frac{\sum_{n=1}^N U_n}{\sum_{n=1}^N L_n} \quad (1.6)$$

The simulation dynamics compute the FDS over a time horizon for a particular insolation, load shape, solar capacity, and battery storage capacity. We can represent these dynamics as the map (1.7) where $I, L \in \mathbb{R}^N$ and $C_s, C_b \in \mathbb{R}$.

$$FDS_N = f_N(I, L, C_s, C_b) \quad (1.7)$$

To construct an isoreliability curve for a given location and load, we iterate across possible storage capacities and find the solar capacity that yields the target FDS. Finding this solar capacity can be written as an optimization problem, where \hat{C}_b and \hat{C}_s are a pair of battery and solar capacities that yield the target FDS over time horizon N , denoted $F\hat{D}S_N$:

$$\hat{C}_s = \arg \min_{C_s} \|F\hat{D}S_N - f_N(I, L, C_s, \hat{C}_b)\| \quad (1.8)$$

This problem is convex and can be solved with Newton's method, yielding a point (\hat{C}_b, \hat{C}_s) on the isoreliability curve. We denote this isoreliability curve by the set of points $C_{FDS,I,L}$, noting that it is parameterized by the Fraction of Demand Served, the insolation profile (given by the location), and the normalized load profile. We define (1.9) as the function that gives a solar capacity on the isoreliability curve associated with a battery capacity C_b .

$$C_s = G_{FDS,I,L}(C_b) \quad (1.9)$$

We use a variable step size iteration to select storage capacity points. First, we compute the minimum storage necessary if there was an effectively infinite solar capacity, and use this as the minimum storage. Starting with this storage capacity, we compute the solar capacity, and then increase the storage capacity and repeat. We expand the step size in storage capacity to compensate for the fact that as storage gets large, it has a diminishing effect on reducing storage capacity. This reduces the number of iterations while still constructing a smoothly sampled curve. Finally, we check that the isoreliability curve was calculated with sufficient precision so that it is convex, and that our variable step size yielded sufficiently many points for the subsequent cost minimization to yield accurate results; if these criteria are not met, we repeat the process with greater precision in Newton’s method and a smaller step size in storage capacity. This results in a discretized isoreliability curve for a particular FDS at a particular location that is stored in a lookup table.

Recall that load is normalized to 1 kWh/day. Thus, points on the isoreliability curve are solar and storage capacities that serve 1 kWh/day load at a specific FDS. We can see by inspecting the dynamics above, that if load was scaled by some factor, we would achieve the same reliability by scaling solar and battery capacity by the same factor. Thus, the points on the isoreliability curve can be scaled to satisfy an arbitrary average load. The curve is dependent on the shape of the load profile, but is independent of the average load.

We calculate and store isoreliability curves at each degree longitude and latitude across SSA. We construct curves for values of FDS: 0.6, 0.8, 0.9, 0.95, 0.975, 0.9875, 0.9938, 0.9969, 0.9984, 0.9992, 0.9996, 0.9998, 0.9999. These numbers result from the sampling expression (1.10) and is constructed to sample evenly from the logarithm of FDS. As we show in our results, we find the cost of electricity scales linearly with the logarithm of FDS.

$$FDS_k = 1 - 0.1 \times 2^{-k}, \quad k \in \{-2, -1, \dots, 10\} \quad (1.10)$$

This pre-processing technique allows the computationally expensive part of the algorithm – computing the isoreliability curves – to be done offline and stored for future use.

Cost optimization

The minimum cost system design can be computed from an isoreliability curve and the costs of storage and solar respectively. The specific costs used for different scenarios are described in the economic assumptions in the next section and listed in Table 1.1. We include the discounted battery replacement costs in the total price of storage (1.11), where P_b is the total price of the battery including replacement, r is the discount rate, n is the project term, T is the battery lifetime, and p_b is the battery costs per kWh from Table 1.1.

$$P_b = \frac{1 - (1 - r)^n}{1 - (1 - r)^T} p_b \quad (1.11)$$

The total battery capital cost is denoted K_b as in (1.12).

$$K_b = C_b P_b \quad (1.12)$$

We also include solar derating; the simulation is performed with being the de-rated capacity of the solar. Thus the total capital cost of solar K_s is given by (1.13) where p_s and p_c are the prices per kW of solar modules plus hardware and the charge controller, respectively, from Table 1.1, and α is the de-rating factor of 0.85 to account for dirt, wiring and conversion losses, and panel mismatch.

$$K_s = C_s \left(\frac{p_s}{\alpha} + p_c \right) \quad (1.13)$$

The total capital cost of the system K is given by (1.14), where \bar{L} is the average daily load and K_l are the capital costs per unit peak load (these are the inverter and soft costs assumed to scale with the peak load), and L_{peak} is the system peak load capacity.

$$K = (K_s + K_b)\bar{L} + K_l L_{peak} \quad (1.14)$$

Note that L_{peak} is not necessarily the peak given by the load profile use to generate the isoreliability curve; rather it is the peak capacity of the system specified by design. These differ in the case that a representative load profile is used in the isoreliability curve, such as an average load profile, but the system is required to meet intermittent atypical peaks in load. We calculate the Levelized Cost of Energy in (1.15) using the methodology of the United States National Renewable Energy Laboratory (NREL), where CRF is the Capital Recovery Factor and O is the fixed annual operations and maintenance costs from Table 1.1.

$$LCOE = \frac{K \times CRF + O \times L_{peak}}{365 \times \bar{L} \times FDS} \quad (1.15)$$

The CRF is given by (1.16) where r is the annual discount rate, and m is the project term. We assume no variable operations and maintenance costs.

$$CRF = r \frac{(1+r)^m}{(1+r)^m - 1} \quad (1.16)$$

Our optimization problem is to minimize LCOE over the decision variables C_s and C_b , subject to the constraints of the isoreliability curve. The project term, discount rate, average load, peak load, solar derating, and prices of solar and batteries are all parameters that can be adjusted independently of the isoreliability curve. This flexible parameterization without needing to perform additional simulation is what allows this approach to scale as a module of a higher-level optimization. Note that this problem is equivalent to minimizing the sum of the solar and battery capital costs (1.17) subject to the isoreliability constraint (1.18).

$$K_s + K_b = C_s \left(\frac{p_s}{\alpha} + p_c \right) + C_b P_b \quad (1.17)$$

$$(C_b, C_s) \in C_{FDS,I,L} \quad (1.18)$$

For a continuously differentiable isoreliability curve, this problem has an analytical solution given by (1.19)

$$\frac{\partial C_s}{\partial C_b} = - \frac{P_b}{\frac{p_s}{\alpha} + p_c} \quad (1.19)$$

Given that in the application of this algorithm, the isoreliability curve is discretized with on the order of one hundred sample points, it is trivial to perform an integer search over the reliability curve to find the cost minimizing point. We use this integer search rather than approximating the analytical solution.

Component costs and economic assumptions

For the purposes of describing large scale penetration of decentralized solar technologies, this study assumes that competition will drive costs to the best practice rates. IRENA provides the self-reported cost breakdown of solar home systems (SHS) greater than 1 kW installed across Africa, including a representative “best-practice” scenario that we use to construct present-day cost estimates [50]. This scenario indicates an installed price of USD 2.3/W installed, excluding the cost of storage, which is broken down into approximately USD 1/W for solar modules and DC balance of system costs such as racking, wiring, and circuit protection, USD 0.3/W for the inverter, and USD 1/W for AC balance of system and soft costs. These estimates are consistent with IRENA’s reported values for grid-tied PV systems of a similar scale, and we use them as our best estimate of current costs.

Estimating battery costs is more complicated as the choice of technology involves trade-offs between capital cost, efficiency, lifetime, ease of maintenance, required space, and transportation logistics. IRENA reports in 2016 that most SHS use lead-acid (LA) batteries, but that lithium-ion (LI) batteries are beginning to appear on the market [50]. Based on these market trends, and techno-economic analyses from the literature [58, 32, 28] we consider LI and LA to be the leading battery technology choices for small scale decentralized solar. Though the use of LI is still in its infancy for SHS applications larger than 1 kW, it appears that we are approximately at a parity point where a solar developer would be indifferent to the two technologies if both are available with training and support. Diouf et al [27] stated in 2015 that because of longer lifetime, deeper depth of discharge, and better performance at high temperatures, there are cases where LI are more economical. IRENA echoed this position in 2016 [50]. A simple calculation taking these factors into account with a 2012 LI price of USD 600/kWh and a LA price of USD 120/kWh (this LA price is consistent with estimates from the literature [106, 50]) shows that at temperatures found in Africa, LI resulted in a lower cost over its lifetime [2]. In 2017, for example, one could purchase a 13.5 kWh LI Tesla Powerwall for USD 5,500 (resulting in USD 400/kWh) that is warranted for 10 years with a 100% depth of discharge, indicating that LI is increasingly competitive with LA.

Given the present approximate lifetime cost parity between LI and LA batteries, and the expectation that spill-over effects from vehicle electrification will soon make LI the dominant battery technology [27], we use LI battery models for our analysis. This choice simplifies comparison between present and future scenarios; and at present, the modelled costs can be interpreted as roughly equivalent to the costs for LA systems that are more widely available. We assume the costs of LI batteries are currently USD 400/kWh. We also assume a maximum power point tracking charge controller will be used with a cost of USD 0.2/W, which is

similar to those reported [50] and offered internationally by online retailers. We assume that all significant \$/kW costs associated with the battery are captured by the charge controller and inverter costs.

To calculate the levelized cost of electricity (LCOE), we assume a discount rate of 10% and a project term of 20 years. Systems may not last 20 years if they are not well maintained, but we assume best practices. The discount rate is higher than that used by Szabó et al [108], but consistent with other studies [20, 8], and reflect a perception of high project risk. We assume battery replacement occurs at 10 years, which is conservative relative to an estimate of 15 years found in the literature [23]. We use the estimate to reflect uncertainty in the lifetime in potentially challenging technical environments found in decentralized systems in SSA, and to reflect the warranty term of 10 years of the Tesla battery. We also assume replacement at initial costs so as not to embed precise assumptions of the rate of price decreases into the model. All of these approximations are designed to be conservative, meaning actual LCOE are likely to be lower than our estimates. This analysis includes a look ahead at what the costs of decentralized solar could be in the future. The numbers used are meant to be suggestive, not predictive, but are drawn from market studies that are based in learning-curve methods. We construct a future scenario based on predictions for 2025. IRENA predicts that utility-scale PV, including soft costs and balance-of-system will be reduced by something in the neighbourhood of 50% by 2025 [51]. We suggest that decentralized costs in Africa might fall by a similar percentage of their current costs; that is, our future scenario examines a 50% reduction in module, soft costs, and balance of systems (both AC and DC). We also use the “low-cost” scenario published by researchers at the (United States) National Renewable Energy Laboratory which predicts a reduction to 25% of 2015 costs in 2025 [23]. When calculating the trade-off between solar and storage capacity, we assume that the inverter and soft costs are proportional to the peak power demand of the system, rather than the solar or storage capacity. The charge controller cost is added to the cost of solar specific components, because its power rating is dependent on the output of the solar module. The costs used in the optimization model are shown in Table 1.1. Soft costs include installation costs but do not include profit or overhead for a business providing systems. The AC balance of system costs are assumed to be proportional to the peak load capacity of the system along with the inverter, while the DC balance of system costs are proportional to the array capacity.

Tiers of access

Additional technical parameters are dictated by the ESMAP Tiers of Access [12]. We analyse Tier 5, which is considered the highest level of access. Tier 5 requires an average load of 8.2 kWh per day, a peak capacity of 2 kW, and service available for 159 out of 168 hours in a week, which translates to an FDS of 94.6% [12]. These parameters are used to calculate the LCOE using the optimization methods described above.

Empirical grid and electrification metrics

We use the World Bank Enterprise Surveys on Infrastructure to estimate the ASAI for countries in SSA [113]. These surveys estimate the number of grid outages per month and the average duration of the outage. From this, we can compute the average number of hours in a year when the grid power out for each country, and then divide by the number of hours in a year to give the ASAI. At the time of writing, the years surveyed range from 2006 to 2017 by country

Electricity tariff data is obtained at a national level from the World Bank and is available for most countries [114]. We use retail rates. Where tiered rate structures are in effect, we use the rate at 250 kWh per month for ESMAP Tier 5.

National electrification rates are available through the World Bank data tables [112]. Please see the Supplementary Table 1.2 for copies of the tariff, reliability, and electrification metrics.

Code Availability

Analysis was conducted using MATLAB R2017a. Code is available upon request from Jonathan Lee at jlee@berkeley.edu and at <https://github.com/leejt489/solar-reliability-cost-matlab>. A Python version is also available upon request and at <https://github.com/leejt489/solar-reliability-cost-python>. An implementation of the model is available at <https://emac.berkeley.edu/reliability>.

Data Availability

All data used in this study with the exception of the sample micro-grid load profile from Uganda is publicly available and referenced below. The sample micro-grid load profile is owned by New Sun Road, P.B.C., and can be made available upon reasonable request. Tables on national electrification rates [112], grid reliability [113], and electricity tariffs [114] used to generate Figure 1.1 and Figure 1.6 are included in the Supplementary Table 1.2 for convenience. Solar insolation data is available from the [United States] National Aeronautics and Space Administration [121], and code to download and analyse this data is available at the repositories listed above.

1.A Supplementary Note: Optimal system designs

Supplementary Figure 1.7 shows the range of optimal system designs – that is, the lowest cost capacities of solar and storage that ensure a given FDS – for ESMAP Tier 5 systems across each location as FDS changes. There is wide variation in the storage to solar ratio at high FDS, and though solar capacity increases fairly regularly with FDS, storage capacity does not increase much from 90% to 99% FDS, causing the ratio to decrease in this range of FDS. At FDS above 99%, storage capacity takes a bimodal distribution, causing the storage to

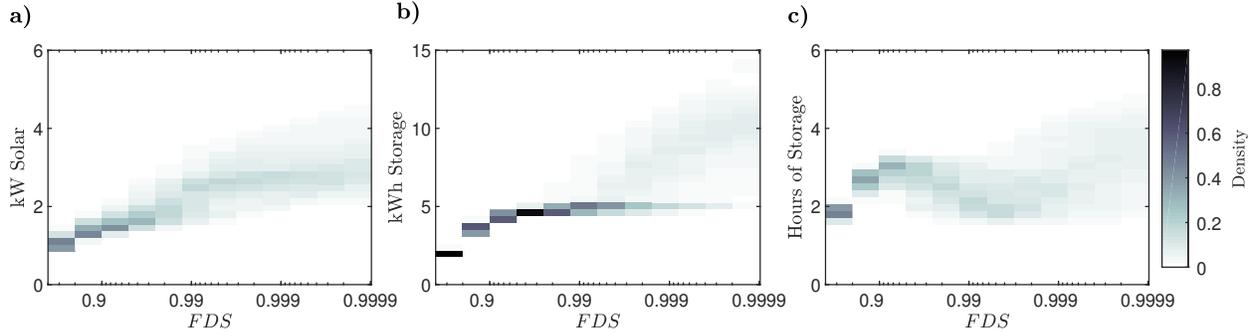


Figure 1.7: Statistical relationship between optimal system design and FDS for Tier 5 systems at current costs

Panel a) shows the density of the optimal solar capacity at different FDS for each location, while panel b) and c) show the density of optimal storage capacity and the ratio of storage to solar capacity at optimum.

solar ratio to scatter nearly uniformly between 2 and 4 hours from 99.9% to 99.99% FDS. The key results from Supplementary Figure 1.7 are that optimal system design is more variable at higher levels of reliability, which implies that considering how the optimal system design changes with location is especially important at high levels of reliability, and is consistent with Figure 1.5 of the main text and Supplementary Figure 1.8. We find also that the optimal storage to solar ratio is mostly confined to between 2 and 4 kWh of storage for every kW of solar, though this range is dependent on the shape of daily electricity load profile (see Supplemental Information). We note that this recommends a higher storage capacity than is reported by Levin and Thomas in a survey of installed SHS, who find most SHS use a ratio in the range of 1 to 2 kWh of storage per kW solar [64].

When repeating this analysis for the possible future cost scenario with factor of 4 reductions in battery costs and factor of 2 reductions in solar costs, there is a shift to higher storage capacity deployment. This shift in optimal system design as costs change will be important to include in multi-year planning tools that forecast changing prices of materials; i.e. it is not appropriate to assume that future systems will employ the same capacities of solar and storage. These results are also of practical importance to project developers and system designers in industry who must select the optimal system design to be built.

We have shown that there is wide variation in optimal system design at high FDS. We now examine the impact of choosing a suboptimal system design on costs to support the use of our more detailed model. To do this, we draw on the approach used by Szabó et al [108], where we calculate the optimal system design at each location for 95% FDS given current costs. We then hold that storage to solar ratio fixed, and calculate the solar capacity and resulting LCOE necessary to achieve different FDS using the fixed ratio. Thus, we can calculate the difference in cost between the optimal choice and the fixed ratio solution for each location (Supplementary Figure 1.8). The additional cost, or penalty for the suboptimal solution, increases with FDS, and the spatial variation is significant. At FDS at or exceeding 99.9%, there is a cost penalty around USD 0.10/kWh in many regions if a

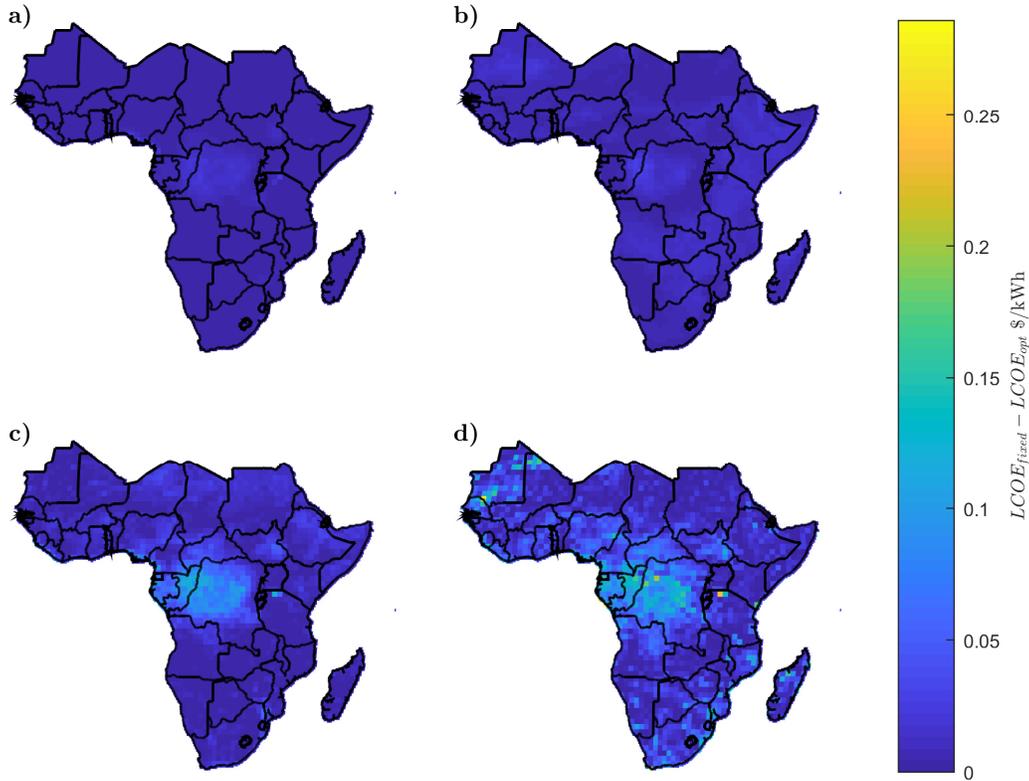


Figure 1.8: Additional cost incurred if holding the storage to solar ratio fixed
 Panels a), b), c), d) show the difference for FDS of 0.9, 0.99, 0.999, and 0.9999, respectively.

suboptimal system design is used. From these results we conclude that it is appropriate to use a simple model for estimating the cost of decentralized systems at FDS below 99%, but at higher FDS, it is necessary to use an optimization that accounts for local weather patterns on a daily scale.

1.B Supplementary Note: Sensitivity to daily load profile shape

As mentioned in the main text, the isoreliability curve is dependent upon the shape of the daily electric load profile: the shape of the load profile is part of the simulation to ensure that a particular system satisfies the target FDS. In the analysis presented in the main text, we use the simple assumption of a constant load profile. Ideally, we would have large datasets of electricity consumption across locations in SSA, but this is not available. Therefore, to test the robustness of our results we compare different load profiles. We find that having

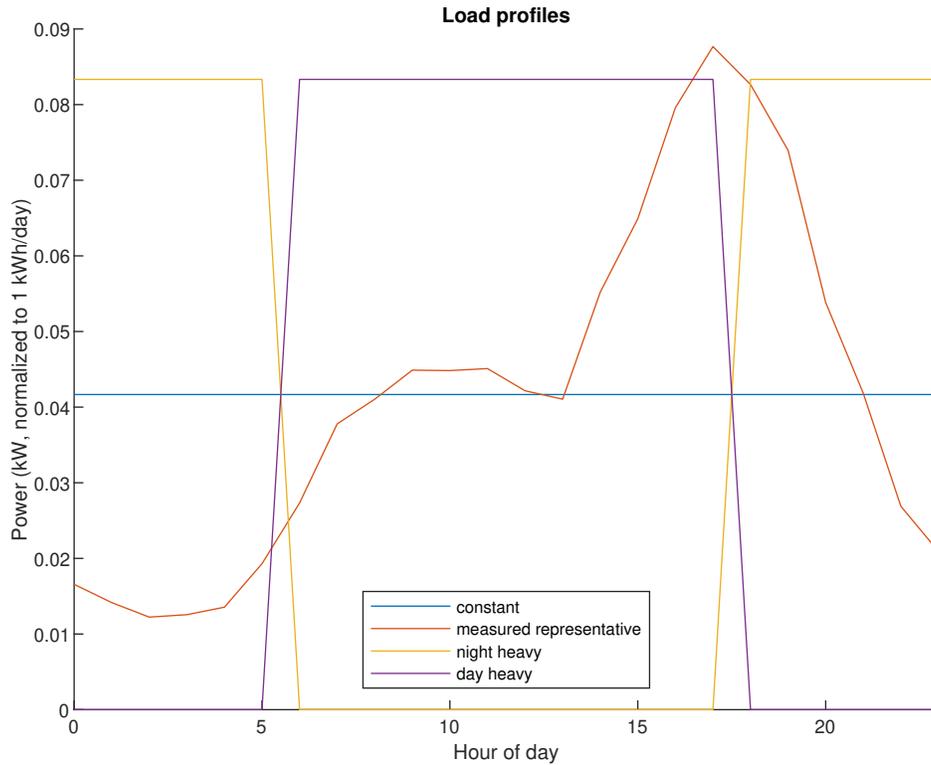


Figure 1.9: Daily load profile shapes

more demand at night increases LCOE due to additional storage capacity required, but that the logarithmic scaling and cost of incremental FDS are unchanged.

We compare four different load profiles: constant, day-heavy, night-heavy, and an empirical sample from a microgrid in Uganda shared by New Sun Road, Public Benefit Corporation. The shapes of the different load profiles are shown in Supplementary Figure 1.9.

We compare the relationship between FDS and LCOE for different load profiles in Supplementary Figure 1.10. Demand concentrated during the day causes lower LCOE for the same FDS, while demand concentrated at night causes higher LCOE; however, the shape of the reliability premium and variance around the best-fit line are approximately the same in all cases. From this we conclude that changes in the daily load profile shift the LCOE by an amount that is independent of FDS.

It is also interesting to note that the constant demand assumption yields very similar results to the empirical load profile, which has a spike in the evening and low demand in the pre-dawn hours. This supports the use of the constant demand assumption in our analysis. In Supplementary Figure 1.10, we see a difference of approximately USD 0.25/kWh between the day-heavy and night-heavy scenarios. This can roughly be interpreted as an upper bound on the benefits that could be realized by shifting load from the night to the day through

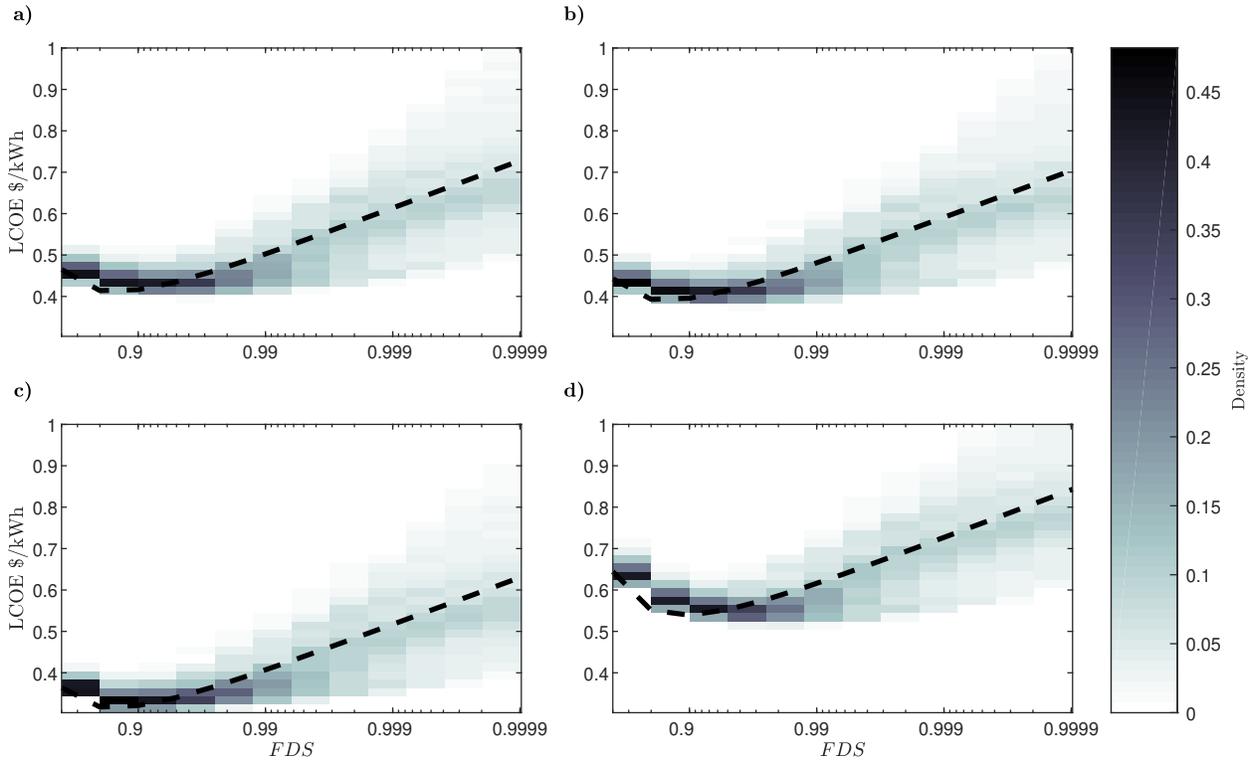


Figure 1.10: Statistical relationship of LCOE and FDS for different load profiles

The plots show the density of computed LCOE at different FDS for each location in SSA sampled at 1 degree, as in Figure 1.3 of the main text. Panels a), b), c), and d) show the relationship for the constant, measured representative, day heavy, and night heavy load profiles of Supplementary Figure 1.3, respectively. Note that panel a) is identical to Figure 1.3, panel a), of the main text on a different vertical axis.

demand-response type programs. It is important to observe in Supplementary Figure 1.10 that the constant and sample measured demand have approximately identical costs. This is because the sample demand has an evening peak that is approximately balanced out by low consumption late at night, resulting in approximately the same amount of battery capacity necessary as the constant case. This similarity in results validates the assumption of constant demand used in the main text. It is also important to point out that the inverter capacity is specified by the ESMAP Tier 5 requirements³, and is not dependent on the shape of the load profile.

We see in Supplementary Figure 1.11 that the demand profile affects the optimal storage to solar ratio. Unsurprisingly, having more demand at night requires more storage capacity relative to solar. In the case where all demand is at night, the optimal ratio is in the range of 3 to 5 hours of storage.

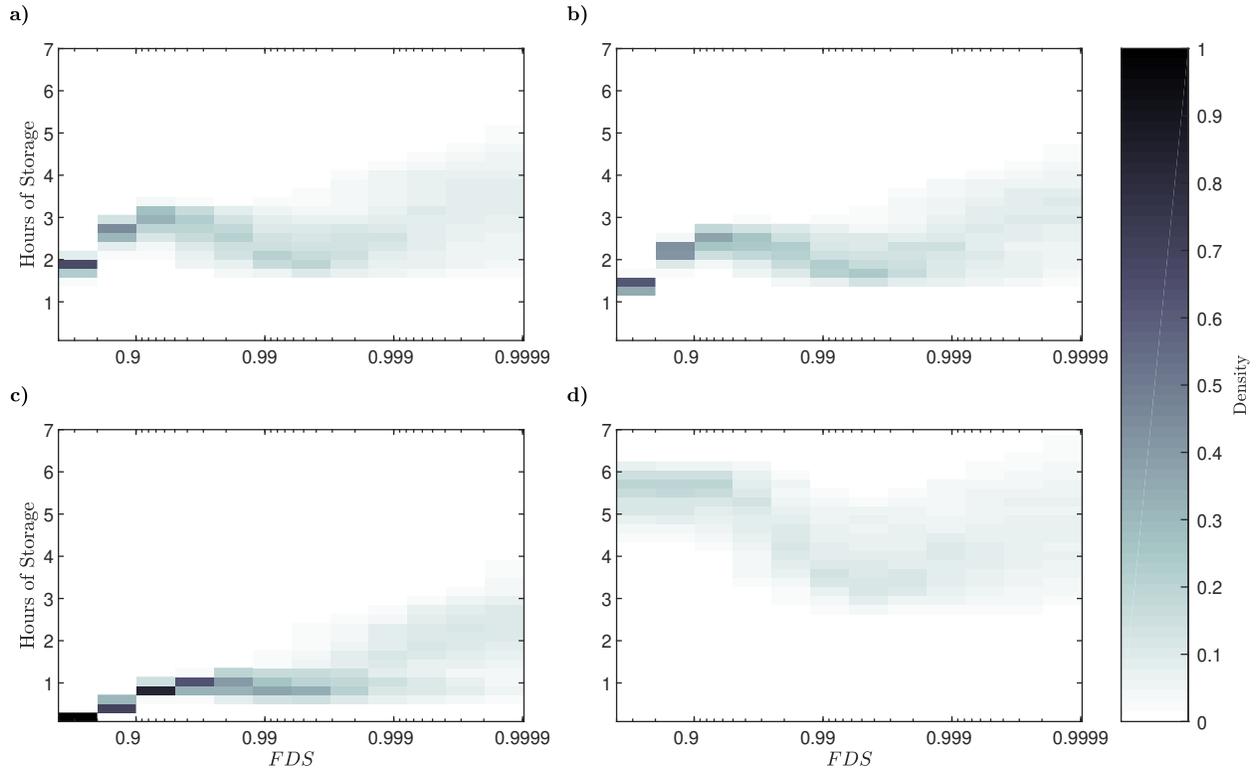


Figure 1.11: Statistical relationship between optimal storage to solar ratio and FDS for different load profiles for Tier 5 systems at current cost

Panels a), b), c), and d) show the relationship for the constant, measured representative, day heavy, and night heavy load profiles of Supplementary Figure 1.9, respectively. Note that panel a) is identical to panel c) of Supplementary Figure 1.7.

1.C Supplementary Note: Electricity tariff, reliability, and electrification data

The data used in the main text for the electricity tariffs, reliability, and electrification are shown in Supplementary Table 1.2. Electricity tariff data was made available through a World Bank report [114], and at the time of writing can be accessed online at <http://databank.worldbank.org/data/download/avpfa/avpfa-tariff.xlsx>. For comparison to ESMAP Tier 5 systems, the tariff at 250 kWh per month was used as it corresponds to approximately 8 kWh per day, which is closest to the ESMAP Tier 5 requirement of 8.2 kWh per day of consumption. Reliability data was computed from the World Bank Enterprise Infrastructure Surveys [113], which at the time of writing can be viewed at <http://www.enterprisesurveys.org/data/exploretopics/infrastructure#sub-saharan-africa>. The outage time is computed from the number of outages per month and the average duration. Electrification rates were recorded by the Sustainable Energy for All

(SE4All) initiative and made available at the time of writing by the World Bank [112] at <https://data.worldbank.org/indicator/EG.ELC.ACCS.ZS>.

Table 1.2: Tariff, reliability, and electrification data by country

Country	Retail price @ 100 kWh / month	Retail price @ 250 kWh / month	Monthly outages	Average outage duration	ASAI	Electrification Rate
Angola	0.04	0.04	4.7	13.5	8.7%	33%
Benin	0.3	0.29	28	3.7	14.2%	29%
Botswana	0.1	0.09	4.1	2.7	1.5%	53%
Burkina Faso	0.31	0.32	9.8	3.3	4.4%	18%
Burundi	0.07	0.13	16.6	4.8	10.9%	5%
Cameroon	0.1	0.16	7.6	8.7	9.1%	62%
Cape Verde	0.55	0.54	3.2	9.2	4.0%	96%
C.A.R.			29	8.1	32.2%	3%
Chad	0.18	0.21	19.6	8.8	23.6%	4%
Comoros	0.36	0.36				69%
Congo DRC			12.3	5.6	9.4%	18%
Congo						42%
Côte d'Ivoire	0.27	0.2	3.5	5.5	2.6%	62%
Djibouti						42%
Equatorial Guinea						66%
Eritrea			0.5	2.8	0.2%	32%
Ethiopia	0.02	0.03	8.2	5.8	6.5%	25%
Gabon	0.25	0.25	4.6	5.4	3.4%	89%
Gambia	0.23	0.23	21	6.9	19.8%	45%
Ghana	0.12	0.12	8.4	7.8	9.0%	72%
Guinea	0.03	0.03	4.5	3.2	2.0%	26%
Guinea-Bissau			5.2	17.9	12.8%	21%
Kenya	0.23	0.27	6.3	5.6	4.8%	20%
Lesotho	0.11	0.11	2.2	6.6	2.0%	17%
Liberia	0.51	0.51	1.7	4.7	1.1%	10%
Madagascar	0.17	0.17	6.7	1.9	1.7%	13%
Malawi	0.08	0.08	6.7	4.3	3.9%	12%
Mali	0.27	0.28	4.2	5.1	2.9%	26%
Mauritania	0.13	0.12	1.2	3.2	0.5%	29%
Mauritius	0.15	0.19	1.2	3.2	0.5%	100%
Mozambique	0.11	0.09	1.6	4.3	0.9%	40%

Namibia	0.18	0.15	0.6	5.8	0.5%	32%
Niger	0.17	0.18	18.5	1.6	4.1%	15%
Nigeria	0.14	0.1	32.8	11.6	52.1%	45%
Rwanda	0.23	0.23	4	4.3	2.4%	27%
Sao Tome & Principe	0.1	0.14				59%
Senegal	0.24	0.27	6	1.8	1.5%	61%
Seychelles	0.11	0.11				98%
Sierra Leone	0.19	0.21	13.7	10.2	19.1%	14%
Somalia						15%
South Africa	0.1	0.1	0.9	4.5	0.6%	86%
South Sudan			1.5	4.7	1.0%	1%
Sudan			3.4	2.5	1.2%	40%
Swaziland	0.09	0.09	3.7	3.7	1.9%	65%
Tanzania	0.12	0.2	8.9	6.3	7.7%	30%
Togo	0.24	0.24	5.5	2.1	1.6%	27%
Uganda	0.22	0.23	6.3	10.1	8.7%	19%
Zambia	0.1	0.08	5.2	2.8	2.0%	28%
Zimbabwe	0.07	0.1	4.5	5.2	3.2%	52%

1.D Supplementary Note: Cost of reliability results by country

The results by country for the costs of Tier 5 solar systems (Figure 1.2 of the main text), the reliability premium (Figure 1.4 of the main text), and the difference between decentralized solar and the grid (Figure 1.6 of the main text) are presented in Supplementary Table 1.3, Supplementary Table 3, and Supplementary Table 4, respectively. To be clear, by “results by country” we mean the mean and standard deviation of the values falling within a country’s borders sampled at 1-degree latitude by 1-degree longitude. The spatial variation can be seen in the figures in the main text. Note that Cape Verde, Comoros, Sao Tome and Principe, Mauritius, Swaziland, and Togo are omitted because there is no point on the 1-degree latitude by 1-degree longitude sample grid that falls within those countries’ borders. Note also that in the case of Nigeria, the high cost of decentralized solar is because the grid reliability is so low (48%), and that designing a decentralized solar system for FDS below 90% results in increasing costs, as shown in Figure 1.3 of the main text.

Table 1.3: Mean and standard deviation of LCOE for all locations sampled in each country under the current and future economic assumptions

Country	Current decen- tralized Tier 5 LCOE (mean)	Current decen- tralized Tier 5 LCOE (std)	Future decen- tralized Tier 5 LCOE (mean)	Future decen- tralized Tier 5 LCOE (std)
Angola	0.437	0.021	0.182	0.010
Benin	0.433	0.018	0.181	0.009
Botswana	0.433	0.017	0.181	0.008
Burkina Faso	0.419	0.009	0.174	0.004
Burundi	0.448	0.009	0.189	0.004
Cameroon	0.451	0.036	0.189	0.015
C.A.R.	0.431	0.014	0.180	0.007
Chad	0.406	0.009	0.168	0.004
Congo DRC	0.488	0.009	0.205	0.005
Congo	0.460	0.024	0.193	0.010
Côte d'Ivoire	0.446	0.015	0.187	0.007
Djibouti	0.406	0.007	0.168	0.003
Equatorial Guinea	0.494	0.025	0.209	0.008
Eritrea	0.411	0.011	0.171	0.005
Ethiopia	0.421	0.015	0.175	0.007
Gabon	0.492	0.015	0.207	0.005
Gambia	0.423	0.002	0.176	0.001
Ghana	0.445	0.007	0.187	0.004
Guinea	0.435	0.018	0.181	0.008
Guinea-Bissau	0.441	0.016	0.184	0.007
Kenya	0.420	0.013	0.175	0.006
Lesotho	0.481	0.029	0.201	0.012
Liberia	0.496	0.040	0.207	0.013
Madagascar	0.441	0.025	0.184	0.012
Malawi	0.440	0.015	0.184	0.007
Mali	0.426	0.010	0.178	0.005
Mauritania	0.434	0.011	0.182	0.005
Mozambique	0.453	0.023	0.190	0.010
Namibia	0.428	0.018	0.179	0.008
Niger	0.408	0.010	0.170	0.005
Nigeria	0.453	0.050	0.189	0.020
Rwanda	0.452	0.000	0.191	0.000
Senegal	0.420	0.012	0.175	0.006
Seychelles	0.436	0.005	0.183	0.002

Sierra Leone	0.499	0.023	0.209	0.008
Somalia	0.402	0.012	0.166	0.006
South Africa	0.487	0.036	0.205	0.016
South Sudan	0.440	0.015	0.184	0.007
Sudan	0.409	0.010	0.170	0.005
Tanzania	0.433	0.025	0.181	0.011
Uganda	0.422	0.008	0.176	0.004
Zambia	0.426	0.005	0.177	0.002
Zimbabwe	0.441	0.021	0.184	0.010

Chapter 2

Non-intrusive load management under forecast uncertainty in energy constrained microgrids

This chapter was co-authored with Sean Anderson, then an engineer with New Sun Road, P.B.C., and subsequently a doctoral student at the University of California, Santa Barbara; Dr. Claudio Vergara, Chief Architect at ZOLA Electric; and Prof. Duncan Callaway, Associate Professor of Energy and Resources at the University of California, Berkeley. The text was previously published in *Electric Power Systems Research* in 2021. See [60] for the full citation. The text has been reformatted for inclusion in this thesis.

Abstract

This paper addresses the problem of managing load under energy scarcity in islanded microgrids with multiple customers and distributed solar generation and battery storage. We explore an understudied, practical approach of scheduling customer-specific load limits that does not require direct control of appliances or a market environment. We frame this as a stochastic, model-predictive control problem with forecasts of solar resource and electricity demand, and develop alternative solutions with two-stage stochastic programming and approximate dynamic programming. We test the efficacy of the alternative solutions against heuristic and deterministic controllers in an environment simulating the customers' responses to load limits. We show that using forecasts to schedule limits can significantly improve power availability and the customers' benefits from consumption, even without the controller having a full model of the customers' responses.

2.1 Introduction

Without measures for microgrid operators to manage load or communicate scarcity, customers in energy-constrained microgrids will experience suboptimal interruptions. For example, in an islanded microgrid with multiple customers sharing limited photovoltaic generation and battery storage capacity, high daytime loads on cloudy days might lead to evening interruptions of low-power / high-value loads such as lighting. This problem could exacerbate inequity across customers, for example, if some are only able to consume electricity in evening hours when interruptions are more prevalent.

We seek to improve the allocation of energy services in time by establishing dynamic *load limits* based on forecasts that allow customers to consume energy over a time window in quantities up to, but not in excess of, the limit. This control problem is related to other flavors of microgrid Energy Management Systems (EMS) and connected methodologically to recent work on Stochastic Unit Commitment (SUC). The classic unit commitment problem schedules generators to minimize startup, shutdown, and variable fuel costs while meeting an estimate of inflexible load. The stochastic extension typically minimizes a measure of the expectation of costs over a set of uncertain *scenarios* while satisfying constraints [110, 127, 101, 88].

Solutions to the stochastic microgrid EMS problem in the literature typically employ the same scenario approach as its SUC counterpart, but in different contexts with varying models of physical systems, points of control, and objectives. Generally, the microgrid has local intermittent renewable generation and energy storage, can be either grid-connected or islanded, may contain dispatchable generation, and may have controllable loads. If the microgrid is grid-connected, the main grid is treated as an unconstrained resource, but with a time-varying price entering the optimization problem [90, 105, 97]. In islanded or off-grid microgrids, dispatchable generation or flexible load [103] is used to balance supply and demand.

Our system of interest can be classified as an islanded EMS where supply-demand balance is met by flexible demand and storage dispatch, and lost load is assigned a cost in the EMS optimization problem. Prior related studies assume load is directly controllable [86, 91], or that customers respond to a pricing signal [122, 97]. Direct load control and time-varying prices are promising pathways; however, they have some limitations. Direct load control requires ubiquitous remotely controllable appliances and is intrusive to customers, particularly if very large demand shifts are required during periods of scarcity. Time-varying pricing requires carefully designed price formation rules and sufficiently responsive load.

In contrast, load limits require only broadcasting a limit to customers and the ability to disconnect load at the meter if the limit is exceeded. Although this approach is more blunt than direct load control or pricing, it is simple and inexpensive to implement. In the simplest case, the load limit can be sent to the customer directly via a mobile interface, in which case they would manually adjust their consumption. More sophisticated smart appliances could automate the adjustment for the customer, but are not required. In either case, using the total load limit preserves privacy and a degree of customer autonomy without distributed

automation or a structured market. Load limits are advocated in [74], although that study works within the context of market-based solutions.

In the framework we present in this paper, an operator is held accountable implicitly for unreliable service and chooses load limits that maximize a simplified value metric of each customer’s energy consumption. We show that this choice can be formulated mathematically as a *sequential decision problem*, which is well-known to have significant computational complexity [11]. This complexity is intensified by non-convexities in the model of demand subjected to load limits.

We develop two approximations and reformulations to reduce the complexity of the problem. The first uses two-stage stochastic programming with assumptions about the forecast to cast the problem as a less complex mixed integer quadratic program (MIQP). The second uses approximate dynamic programming in conjunction with two-stage stochastic programming to reduce the problem to a sequence of smaller MIQPs. We compare the performance of the approaches in simulation against a baseline model with no control, a heuristic, and a predictive controller that uses only the mean forecast.

The paper contributes a framework for developing stochastic, predictive models for controlling load through consumption limits under forecast uncertainty. The framework is novel in separating the decisions of the customer to respond to load limits from those of the operator to set them, providing a mechanism for evaluating controller performance in the face of model mismatch. We show how stochastic forecasts can be combined with approximate models of the customer response into an optimal decision model that can be solved with out-of-the-box numerical solvers. Our computational experiment results show significant benefits from using forecasts in a receding-horizon control framework, but more modest and variable benefits from using stochastic formulations in place of deterministic forecasts, with the conclusion that model mismatch limits the additional benefit from stochastic approaches. The paper provides a mathematical and computational foundation for exploring different formulations of value and mechanisms to allocate scarce electricity supply.

2.2 Decision problem

We develop a method to set customer-specific load limits in a microgrid where multiple customers share distributed solar and battery storage with limited capacity. The load limit sets a maximum amount of energy that a customer can consume over a window of time. We assume a receding-horizon control (RHC) framework where a microgrid controller acts on behalf of the operator to compute both load limits and power injection setpoints at a fixed time interval, and then transmits these to customers, metering devices, and the distributed energy resources (DERs), as depicted in Figure 2.1.

The essential states are the state of charge of each battery belonging to each customer and the status of the loads and activities that each customer requires electricity for. The evolution of these states are affected by both the decisions of the microgrid controller and the customer. We assume the microgrid controller cannot control individual loads directly,

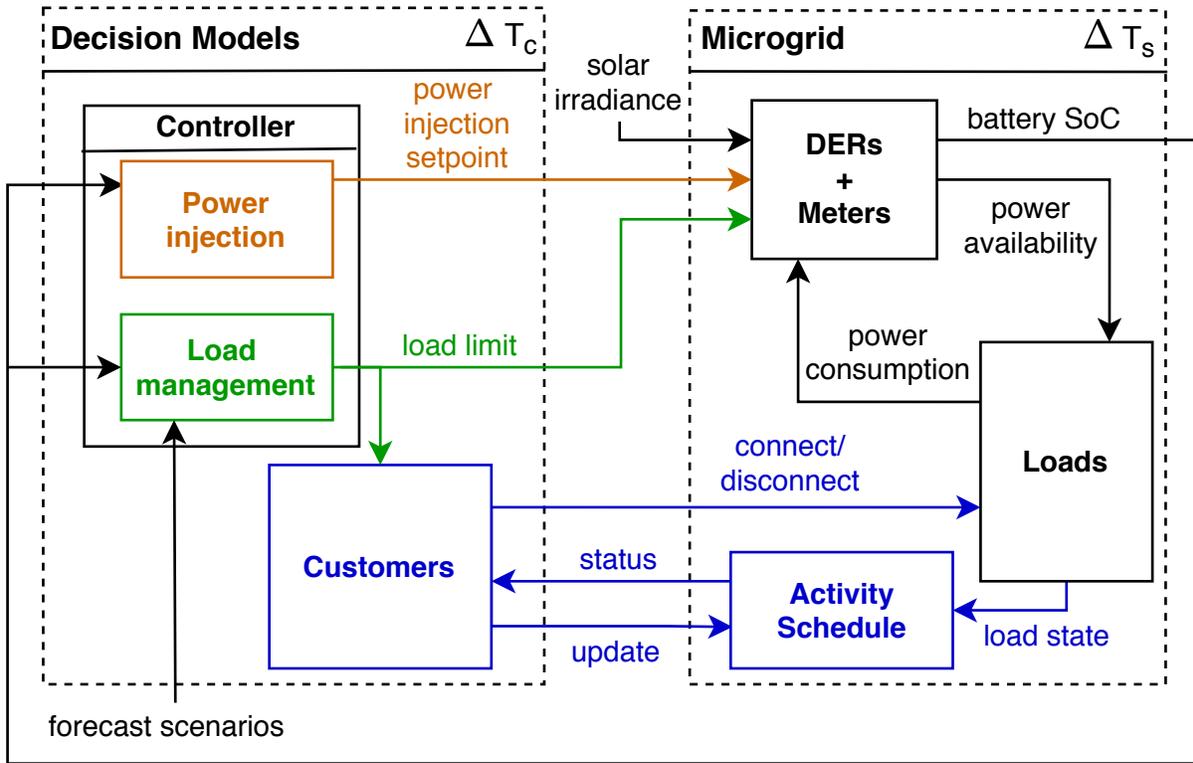


Figure 2.1: Receding horizon control system

but that customers can be sent a load limit that is enforced at their meter. We also assume the controller has no knowledge of the customer’s decision model, activities, or individual loads, so its decision is to set an upper bound on uncertain consumption. However, we assume the controller is given an exogenously determined forecast of solar power potential and electricity demand in the absence of consumption. We treat these forecasted variables as stochastic, which given the dynamic nature of the system, presents the controller with a sequential decision problem under uncertainty.

In the following subsections, we present first a relatively simple model of the customer’s decision to adjust consumption given a load limit. The purpose of this model is both to capture model uncertainty from the controller’s limited information and to define performance metrics for evaluating the control strategy from the perspective of the customer. We then formulate the controller’s decision model to set load limits and propose specific approximate methods to make the problem tractable. This is the core contribution of the paper. Lastly, we briefly define a simple feedback controller to compute power injection setpoints to balance state-of-charge between batteries. The purpose of this component is mainly to facilitate simulating power-sharing among DERs and to provide a placeholder for future work to integrate the load-limiting control with optimal power flow models.

Customer decision and consumption model

We assume customers use their loads to conduct a set of activities that they schedule stochastically around a daily pattern. For example, lights are more likely to be used at night for several hours at a time, and microwaves around meal times for a few minutes. Based on assumptions about appliance ownership and usage patterns which are qualitatively consistent with our field experience, we randomly generate a schedule of activities for each customer that they would carry out if not subjected to limits. Customers derive a value when activities are completed without interruption, but incur an interruption cost otherwise. A customer can cancel an activity before it begins with zero cost but also zero gain.

When a customer is sent a load limit, we assume they cancel or interrupt activities and disconnect the associated loads to maximize their value of completing activities minus any interruption costs from activities already in progress. We introduce this model to emulate behavior to the first order and capture model error in the controller when evaluating performance. We considered models of thermostatically controlled loads but determined this complexity did not provide additional insight, and recommend future work to comprehensively examine the effects of different types of shiftable and state-dependent loads.

Formally, we assume an activity a has a start time T_a^s , time to complete T_a^c , completion value v_a , interruption cost c_a , and a power consumption P_a when its associated load is on. The activity has two states: its remaining time to completion t_a^r and its status σ_a . The status evolves as a finite state machine with states: $\{0 = \text{queued}, 1 = \text{in progress}, 2 = \text{completed}, 3 = \text{interrupted while in progress}, 4 = \text{cancelled before commencing}\}$. We omit the formal transition rules as they are intuitive. Activities are initialized to $t_a^r = T_a^c$ and $\sigma_a = 0$. When the start time is reached, $\sigma_a \rightarrow 1$ and t_a^r decrements as time passes. Unless the activity is interrupted by either the customer or loss of power in the microgrid, $\sigma_a \rightarrow 2$ when t_a^r reaches zero. Statuses 2, 3, and 4 are terminal and the customer receives v_a for $\sigma_a = 2$ and pays c_a for $\sigma_a = 3$.

At a time t , when the customer is faced with a load limit of average power l over ΔT in the future, the sets of relevant activities are those that are already in progress $\mathcal{A}_1 := \{a \mid \sigma_a = 1\}$, and those that are queued but will start within the time window $\mathcal{A}_0 := \{a \mid \sigma_a = 0 \wedge T_a^s < t + \Delta T\}$. For each $a \in \mathcal{A}_0 \cup \mathcal{A}_1$, the customer chooses either $u_a = 0$ to cancel (for $a \in \mathcal{A}_0$) or interrupt (for $a \in \mathcal{A}_1$) the activity, or $u_a = 1$ to proceed as planned. The energy consumed by each activity over the time window is $P_a \min(t_a^r, \Delta T)$ for $a \in \mathcal{A}_1$ and $P_a \min(t_a^r, \Delta T - \max(T_a^s - t, 0))$ for $a \in \mathcal{A}_0$. For activities that will not be completed within the window, we assume the customer expects no load limit in the next window and effectively receives the completed value for activities still in progress. This allows us to represent their decision $u = \{u_a\}$ as an integer linear program to maximize their utility:

$$\max_u \sum_{a \in \mathcal{A}_0 \cup \mathcal{A}_1} u_a v_a + \sum_{a \in \mathcal{A}_1} u_a c_a \quad (2.1)$$

$$\text{s.t.} \sum_{a \in \mathcal{A}_0} P_a \min(t_a^r, \Delta T - \max(T_a^s - t, 0)) + \sum_{a \in \mathcal{A}_1} P_a \min(t_a^r, \Delta T) \leq l \Delta T \quad (2.2)$$

Operator load-limit decision model with forecasts

We assume the microgrid uses RHC with fixed time-step ΔT_c over a horizon T . The controller decides on an action u_t to take on behalf of the operator at time t , based on the current state x_t and a probabilistic forecast \mathcal{W}_t of exogenous disturbances w_t . In our problem, x_t is a vector of the stored energy $E_{n,t}^{\text{stor}}$ in each customer n 's battery, u_t is the vector of load limits $l_{n,t}$, and w_t is the solar generation potential $P_{n,t}^g$ and electricity demand $P_{n,t}^l$ for each customer. We assume \mathcal{W}_t is a finite set of S scenarios consisting of distributed generation and demand values at each time over the horizon for each customer. Each scenario has a probability of occurrence p_s , which we assume to be uniformly $\frac{1}{S}$, but could be given explicitly by the forecast algorithm or tuned to hedge against particular outcomes. We assume the scenarios can be derived from historical measurements, but do not present algorithms for doing so in this paper. The dynamics f are given by:

$$E_{n,t+1}^{\text{stor}} = E_{n,t}^{\text{stor}} + P_{n,t}^c \Delta T_c, \quad 0 \leq E_{n,t}^{\text{stor}} \leq E_n^{\text{max}} \quad (2.3)$$

where $P_{n,t}^c$ is the average net charge power into each customer's battery. $P_{n,t}^c$ is determined implicitly by the controller's action, the state variables and disturbances across *all* customers, and constraints defined subsequently, such as the capacity of each battery E_n^{max} and conservation of energy.

A critical detail in RHC is that the operator makes the next decision *after* observing a realization of the forecast w_t , the new state x_{t+1} , and given a new forecast \mathcal{W}_{t+1} ; however, to make the optimal decision u_t at time t , they have to compute what decision they *would* make at the next time-step given all possible outcomes, and so on over the horizon. This requires assuming how the forecast will be updated as realizations are observed, which we denote with the function g . The proper definition of g is ambiguous without additional information about the forecasting process, but has implications for the decision model; we discuss this in detail after stating the decision model in its general form.

The objective is to maximize the expected benefit of using electricity in the current time period plus the expected future benefit in subsequent time periods. This multi-stage decision problem can be represented mathematically in general with (2.4)-(2.8), where u'_t denotes hypothetical actions to take at the present time t and $\tau \in [t, t + T - 1]$ denotes time-steps over the horizon. Note that the variables defined for $\tau > t$ are predicted future trajectories. Similarly f and g are models and do not necessarily match the physical or simulated system dynamics exactly. Q_t determines the expected benefit over the forecast horizon for any state and action, and is defined recursively as a sum of the expected present benefits b and the future benefits V_{t+1} given the new state and new forecast. V_τ is the maximum value from time τ assuming the operator acts optimally given state x_τ and forecast \mathcal{W}_τ .

$$u_t = \underset{u'_t}{\operatorname{argmax}} Q_t(x_t, u'_t, \mathcal{W}_t) \quad (2.4)$$

$$Q_\tau(x_\tau, u_\tau, \mathcal{W}_\tau) = \mathbb{E}_{\mathcal{W}_\tau} [b(x_\tau, u_\tau, w_\tau) + V_{\tau+1}(x_{\tau+1}, \mathcal{W}_{\tau+1})] \quad (2.5)$$

$$V_\tau(x_\tau, \mathcal{W}_\tau) = \max_{u_\tau} Q_\tau(x_\tau, u_\tau, \mathcal{W}_\tau) \quad (2.6)$$

$$x_{\tau+1} = f(x_\tau, u_\tau, w_\tau) \quad (2.7)$$

$$\mathcal{W}_{\tau+1} = g(\mathcal{W}_\tau, x_\tau, u_\tau, w_\tau) \quad (2.8)$$

$$V_{t+T}(x, \mathcal{W}) \equiv 0 \quad (2.9)$$

We assume for simplicity with (2.9) that the future benefit at the end of the horizon is zero regardless of the final battery state, but this can be replaced with any linear or quadratic function. We define the benefit b as a quadratic function of the *actually used* load power P^u averaged over a time-step. P^u is not directly controllable, but is a stochastic variable influenced nonlinearly by the load limit l , whose realization depends on the customer decision and information not available to the controller. To formulate the controller's decision, we model it as (2.11), which is an overestimate of consumption because the customer is unlikely to be able to adjust *exactly* to the limit.

$$b(x_t, u_t, w_t) = \frac{1}{N} \sum_n \left(P_{n,t}^u - \frac{1}{2P_n^{l,\max}} P_{n,t}^{u,2} \right) \quad (2.10)$$

$$P_{n,t}^u = \min(l_{n,t}, P_{n,t}^l) \quad (2.11)$$

The appropriate choice of b in different contexts is an important topic that requires careful study beyond the scope of this paper. We select the quadratic form for the common case where there is diminishing marginal value of consumption. In contrast, a linear function would value all consumption equally and effectively not steer the operator to take any actions to “keep the lights on” by reducing the usage of a few high power loads, which is our qualitative objective. We show in the results that using this form yields desired behavior despite b not representing any direct value. Eq. (2.10) can be modified in several ways while preserving the same structure: it can be shaped for different rates of diminishing marginal value, and weighted differently for particular customers over times of day. These parameters can be functions of past load limits or consumption. Note also that b is increasing up to the maximum possible load, $P_n^{l,\max}$, which is the power rating of their meter.

To specify g , one must assume whether each scenario represents a single trajectory, or a Markov process where the possible values at each moment in time are independent of prior values. The former implies up to S possible trajectories and final states, while the latter implies S^T , effectively leading to two different *scenario trees* after time $t + 1$, illustrated in Figure 2.2. Assuming for the illustration that the scenarios are unique over the first time-step, the two interpretations respectively imply that the operator assumes either, after observing w_t , that 1) they will know with certainty what trajectory they are on and then act optimally with perfect information, or 2) they will again face an uncertain forecast with no gained information.

Both interpretations are approximations of the optimal decision because the forecast itself is an approximation of reality via a finite number of scenarios.¹ Here, we focus not

¹We refer the reader to [11, Ch. 6] for additional discussion showing how the trajectory interpretation can in fact be cast as an approximate solution to the Markov interpretation.

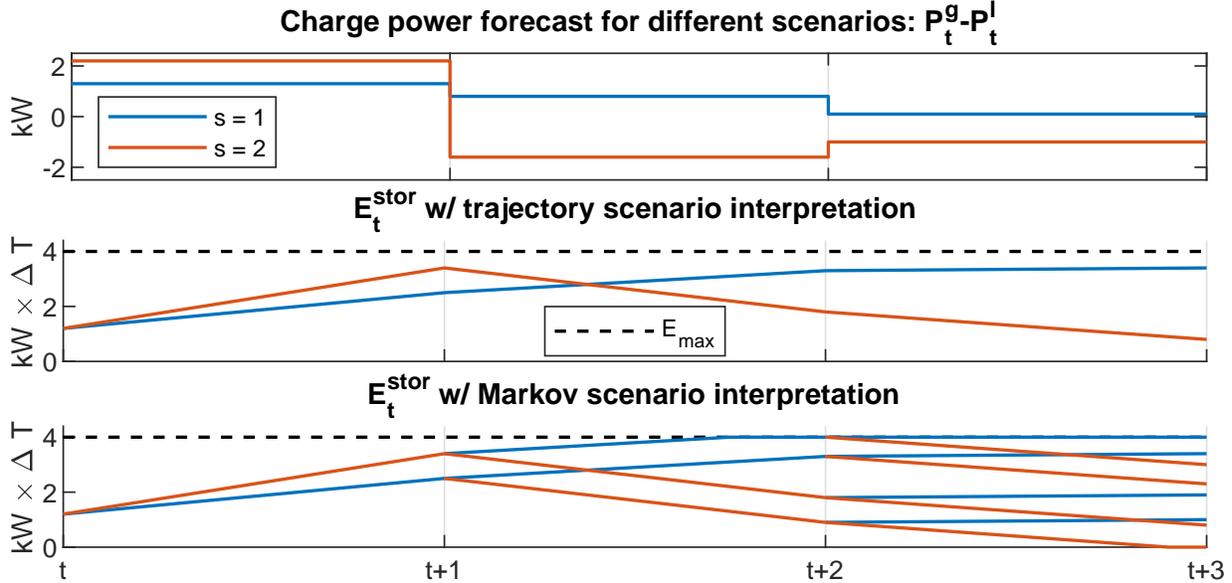


Figure 2.2: Alternative interpretations of a 2-scenario forecast with a single battery over a horizon of three time-periods.

on which is correct – it depends entirely on the details of the forecasting algorithm – but develop solutions for both and compare their performance in simulation. We show that due to having fewer trajectories, the trajectory interpretation can be computed with two-stage stochastic programming, while the Markov interpretation requires additional approximate dynamic programming techniques to solve.

Two-stage stochastic programming solution with trajectory forecast

The key insight and distinction of our model from others is that the operator cannot directly control load, but can only indirectly influence it via a non-convex, piecewise-linear constraint as in (2.11). Otherwise, the problem employs the standard two-stage stochastic model by assuming each scenario is a distinct trajectory [86, 90, 97]:

$$Q_t(x_t, u_t, \mathcal{W}_t) = \sum_s p_s \left(b(x_t, u_t, w_{t,s}) + \sum_{\tau=t+1}^{T-1} b(x_{\tau,s}, u_{\tau,s}, w_{\tau,s}) \right) \quad (2.12)$$

where the single recourse decision $u_{\tau,s}$ for each scenario s is a trajectory with a corresponding state $x_{\tau,s}$ from time $t+1$. The optimization problem includes the constraints eqs. (2.13-2.26) with variables specified by customer n , scenario s , and over time τ as in (2.12). $E_{n,t,s}^{\text{stor}}$ is fixed at the initial condition $E_{n,t}^{\text{stor}}$ for each scenario. P^w is wasted solar (i.e. curtailed when batteries are full), and P is net flow into the network. $P_n^{c,\text{max}}$ is the maximum charge power

of a battery, assumed for simplicity to be the same as discharge power.

$$P_{n,\tau,s}^c = P_{n,\tau,s}^g - P_{n,\tau,s}^w - P_{n,\tau,s}^u - P_{n,\tau,s} \quad (2.13)$$

$$E_{n,\tau+1,s}^{\text{stor}} = E_{n,\tau,s}^{\text{stor}} + P_{n,\tau,s}^c \Delta T_c \quad (2.14)$$

$$0 = \sum_n P_{n,\tau,s} \quad (2.15)$$

$$0 \leq P_{n,\tau,s}^w \leq P_{n,\tau,s}^g \quad (2.16)$$

$$0 \leq P_{n,\tau,s}^u \leq P_{n,\tau,s}^l \quad (2.17)$$

$$-P_n^{c,\max} \leq P_{n,\tau,s}^c \leq P_n^{c,\max} \quad (2.18)$$

$$-P_n^{\max} \leq P_{n,\tau,s} \leq P_n^{\max} \quad (2.19)$$

$$\forall n \in [1, N], \forall \tau \in [t, t + T - 1], \forall s \in [1, S]$$

$$0 \leq E_{n,\tau,s}^{\text{stor}} \leq E_n^{\max} \quad (2.20)$$

$$\forall n \in [1, N], \forall \tau \in [t + 1, t + T], \forall s \in [1, S]$$

To cast the problem in a generic form for standard numerical optimization solvers, we replace (2.11) with the equivalent set (2.21)-(2.26) using binary variables $q_{n,s}$ and the constant $M_n := \max_s P_{n,t,s}^l$ [9]. These constraints, along with (2.17), give two disjoint cases for whether or not the load limit is binding in scenario s : $q_{n,s} = 1 \implies P_{n,t,s}^u = l_{n,t}$, and $q_{n,s} = 0 \implies P_{n,t,s}^u = P_{n,t,s}^l$. Note that the constraints only include the decision $l_{n,t}$ at the first time-step, and that only one decision is made for all scenarios, reflecting that the action must be taken before a scenario is realized. In contrast, the operator assumes they will be taking actions with certainty for $\tau \geq t + 1$, meaning they can set a load limit exactly to the desired consumption in that scenario. In the case where the optimal load limit is the maximum over the forecast, i.e. $l_{n,t} = M_n$, then any $l_{n,t} \geq M_n$ is optimal, so the controller selects no load limit with $l_{n,t} = \infty$.

$$q_{n,s} \in \{0, 1\} \quad (2.21)$$

$$P_{n,t,s}^u \leq l_{n,t} \quad (2.22)$$

$$l_{n,t} \leq P_{n,t,s}^u + (1 - q_{n,s})M_n \quad (2.23)$$

$$l_{n,t} \leq P_{n,t,s}^l + (1 - q_{n,s})M_n \quad (2.24)$$

$$P_{n,t,s}^l \leq P_{n,t,s}^u + q_{n,s}M_n \quad (2.25)$$

$$P_{n,t,s}^l \leq l_{n,t} + q_{n,s}M_n \quad (2.26)$$

This is a mixed integer quadratic program (MIQP) with NS binary variables, and $\mathcal{O}(NST)$ continuous variables and constraints. This scaling in dimension is not to be confused with the complexity of solving the MIQP, which itself scales nonlinearly with the number of variables and constraints.

Approximate dynamic programming solution with Markov forecast

If the forecast is considered Markov, Eqs. (2.4)-(2.8) can be solved with backwards recursion, which in practice requires computing and storing values of $V_\tau(x)$ for each possible x . Computing this if each of N batteries is approximated with X discrete state-of-charge regions requires $(X + 1)^N$ samples, which is intractable. We address this by employing state-space aggregation, approximating the state by the sum of energy stored in all batteries \hat{x}_τ and sampling it uniformly at $X + 1$ points indexed by i . We denote samples of the aggregated state and value function $\tilde{x}(i)$ and $\tilde{V}_\tau(i)$. The continuous and sampled forms are related by piecewise linear interpolation in (2.31)-(2.32), with weights r_i satisfying SOS2 constraints defined for each scenario in (2.36)-(2.42).²

$$\hat{x}_\tau := \sum_n E_{n,\tau}^{\text{stor}} \quad (2.27)$$

$$\hat{x}^{\text{max}} := \sum_n E_n^{\text{max}} \quad (2.28)$$

$$\hat{x}_{\tau+1} = \hat{x}_\tau + \sum_n P_{n,\tau}^c \quad (2.29)$$

$$\tilde{x}(i) = \frac{i}{X} \hat{x}^{\text{max}} \quad \forall i \in \{0, 1, \dots, X\} \quad (2.30)$$

$$\hat{x}_\tau = \sum_i r_i \tilde{x}(i) \quad (2.31)$$

$$\hat{V}_\tau(\hat{x}_\tau) := \sum_i r_i \tilde{V}_\tau(i) \quad (2.32)$$

Given the above, we can now define the optimization problem with objective (2.33) for computing the value function $\tilde{V}_\tau(i)$ at a sample of the state space i at time τ , given values of the next step value function at all samples of the state space $\tilde{V}_{\tau+1}(j) \forall j$, and forecast scenarios $w_{\tau,s}$:

$$\tilde{V}_\tau(i) = \max_{u'} \sum_s p_s(b(\tilde{x}(i), u', w_\tau, s) + \hat{V}_{\tau+1,s}) \quad (2.33)$$

The constraints are the same as the previous two-stage stochastic formulation $\forall s \in [1, S]$ and $\forall n \in [1, N]$, except that only one time-step τ is considered (the load limit constraints (2.22)-(2.26) are defined for time τ), the individual state-of-charge dynamics (2.14) are replaced with the aggregate dynamics (2.34) and likewise for battery capacity (2.35), and the SOS2 constraints are included:

$$\hat{x}_{\tau+1,s} = \tilde{x}(i) + \Delta T_c \sum_n P_{n,s}^c \quad (2.34)$$

²SOS2 refers to “special ordered sets of type 2” constraints [9], which have a structure that can be exploited for better performance by some solvers.

$$0 \leq \hat{x}_{\tau+1,s} \leq \hat{x}^{\max} \quad (2.35)$$

$$\hat{x}_{\tau+1,s} = \sum_j r_{s,j} \tilde{x}(j) \quad (2.36)$$

$$\hat{V}_{\tau+1,s} = \sum_j r_{s,j} \tilde{V}_{\tau+1,s}(j) \quad (2.37)$$

$$\sum_j r_{s,j} = 1 \quad (2.38)$$

$$\sum_j y_{s,j} \leq 2 \quad (2.39)$$

$$y_{s,j} \in \{0, 1\} \quad (2.40)$$

$$0 \leq r_{s,j} \leq y_{s,j} \quad \forall j \in [0, X] \quad (2.41)$$

$$y_{s,j} + y_{s,k} \leq 1 \quad \forall j \in [0, X-2], \forall k \in [j+2, X] \quad (2.42)$$

In general, b should be redefined on the aggregated state space, but our form in (2.10) does not directly depend on state, so we use the same b . Note that (2.33)-(2.42) define an optimization problem only over one time-step. The solution process consists of starting at time $\tau = t + T - 1$ with $\tilde{V}_{\tau+1}(j) = \tilde{V}_{t+T}(j) \equiv 0$, solving the above problem to determine $\tilde{V}_{\tau}(i)$ for each $i \in \{0, X\}$, repeating for $\tau = \tau - 1$, and stopping after solving for $\tau = t + 1$. This entails solving $(X + 1)(T - 1)$ MIQPs, each with a dimension on the order of NS . Once \tilde{V}_{t+1} has been determined, we solve the problem again, but only given the initial state x_t to determine the optimal action u_t to take at time t using \tilde{V}_{t+1} as an approximation of V_{t+1} .

Alternative deterministic solutions

The two controllers of primary interest are described above, but we also define three alternative controllers for use in the computational experiments. The first trivially sets no load limit, the second sets limits according to the piecewise-linear feedback rule (2.43), using only the aggregated state of charge (2.27)-(2.28) and no forecasts, and the third uses a single forecast, computed as the mean over all scenarios, without considering uncertainty. The single forecast formulation is actually equivalent to the stochastic trajectory forecast with $S = 1$, making the binary variables extraneous and reducing the problem to a QP.

$$l_{n,t} = \begin{cases} 0.01P_n^{l,\max} & 0 \leq \hat{x}_t < 0.1\hat{x}^{\max} \\ 0.05P_n^{l,\max} & 0.1\hat{x}^{\max} \leq \hat{x}_t < 0.2\hat{x}^{\max} \\ 0.1P_n^{l,\max} & 0.2\hat{x}^{\max} \leq \hat{x}_t < 0.3\hat{x}^{\max} \\ \infty & 0.3\hat{x}^{\max} \leq \hat{x}_t \end{cases} \quad (2.43)$$

Power dispatch model

In a microgrid with DERs, a dispatch mechanism is required to maintain power balance and coordinate the charge power of each individual batteries. We model solar generation

and batteries interfaced with grid-forming converters, where each group n tracks a setpoint $P_{n,t}^{\text{inj},*}$ of power to inject into the network and the total imbalance is shared by an automatic generation control described below. The primary control objective here, given assumptions to ignore network constraints, is to keep states of charge equally balanced to each other to prevent losing instantaneous power capacity if some were to become drained before others. This is an open research area, but we achieve sufficient balancing with a simple, centralized, proportional feedback controller with gain $K = 2$ and ΔT_c the time-step between control action:

$$P_{n,t}^{\text{inj},*} = \frac{1}{K\Delta T_c} \left(E_{n,t}^{\text{stor}} - \frac{1}{N} \sum_n E_{n,t}^{\text{stor}} \right) \quad (2.44)$$

Integrating more sophisticated predictive power dispatch models with load-limiting to account for network constraints and losses is an important area for future work that becomes increasingly relevant in larger microgrids.

2.3 Microgrid simulation model

To evaluate controller performance, we develop a simulation model of a distributed microgrid to capture interruption events and the evolution of battery states. We use a quasi-static simulation of the steady-state behavior of the the primary and secondary controls of the DER power converters, which govern power sharing and the availability of supply. We introduced grid frequency Δf_t as a state variable in the simulation model to maintain instantaneous power balance. The DERs act as synchronous interconnected areas that maintain power balance using classic droop and automatic generation control subject to constraints on the solar availability and battery charge [39]. We assume the charge and discharge capacity is constrained by the battery inverter rating $P_n^{c,\text{max}}$, the free capacity of the battery, and a linear power derating when the battery state-of-charge is within 10% of its limits. These dynamic constraints are captured respectively by the three terms in the min functions defining the maximum charge $P_{n,t}^{c,+}$ and discharge $P_{n,t}^{c,-}$:

$$\begin{aligned} P_{n,t}^{c,+} &= \min \left(P_n^{c,\text{max}}, \frac{E_n^{\text{max}} - E_{n,t}^{\text{stor}}}{\Delta T_s}, \frac{P_n^{c,\text{max}}(E_n^{\text{max}} - E_{n,t}^{\text{stor}})}{0.9E_{n,t}^{\text{stor}}} \right) \\ P_{n,t}^{c,-} &= \min \left(P_n^{c,\text{max}}, \frac{E_{n,t}^{\text{stor}}}{\Delta T_s}, \frac{P_n^{c,\text{max}} \cdot E_{n,t}^{\text{stor}}}{0.1} \right) \end{aligned} \quad (2.45)$$

The net injection $P_{n,t}^{\text{inj}}$ of each “area” n of DERs tracks the setpoint $P_{n,t}^{\text{inj},*}$ with a frequency response stiffness β_n subject to the charge and solar generation capacity constraints as well as conservation of energy given the loads $P_{n,t}^u$:

$$P_{n,t}^{\text{inj}} = \min(P_{n,t}^g + P_{n,t}^{c,-}, \max(-P_{n,t}^{c,+}, P_{n,t}^{\text{inj},*} - \beta_n \Delta f_t))$$

$$0 = \sum_n P_{n,t}^{\text{inj}} - P_{n,t}^u \quad (2.46)$$

We set the stiffness of area n as proportional to the total inverter capacity: $\beta_n = \beta(P_n^{c,\max} + P_n^{g,\max})$ where $P_n^{g,\max}$ is the PV inverter capacity, and we choose $\beta = 4$. The above system has either a unique solution for Δf_t or no solution; in the latter case, a blackout is implied. In the event of a blackout, meters disconnect all load (thus interrupting customer activities) until the aggregate state of charge reaches 10%, and the DERs come back online automatically. When there is no blackout, the solar generation, curtailment, and battery charge are recovered from Δf_t and $P_{n,t}^{\text{inj}}$ by minimizing curtailment, and the battery stored energy is updated incremented by $P_{n,t}^c \Delta T_s$.

As shown in Fig. 2.1, the control system sets limits for each customer and a power injection setpoint for each DER every $\Delta T_c = 4$ hours. Within that window, the DERs, loads, meters, and customer activity states are simulated on a $\Delta T_s = 2$ minute time-step. We assume the customer updates their activity schedule whenever they receive a new limit and that individual meters enforce load limits by disconnecting load if the limit is exceeded.

2.4 Computational experiments

We conducted two computational experiments with multiple trials to assess the efficacy and computational tractability of the proposed algorithms using the experimental methodology and terminology proposed in [57]. All modelling code and data are available on GitHub, including the complete implementation of the models above, all experimental parameters, and additional data visualization.³ We ran the experiments on a personal computer using an Intel i7-7600U CPU Dual Core, 2.80 GHz CPU with 16 GB of memory. We used CVX version 2.1, build 1127, with MATLAB 2018a to develop the optimization problems with Gurobi 9.0.1 as the solver [43, 42]. In the simulation and timing results, we used MATLAB compiled binaries and the Gurobi API directly instead of CVX to improve performance.

In each trial, we simulate a microgrid of N customers by randomly distributing 300 W PV units and 2 kWh battery units with 1.2 kW charge power. We set the total solar capacity to produce the average unconstrained demand of 330 W, which was computed by simulating users' activities, and 3 kWh of total battery capacity per kW of PV. This results in an average 1.5 kWp PV and 4.5 kWh of storage per user, but variable distributed, and ensures energy scarcity. Each customer is assumed to have a maximum possible load of $P_n^{l,\max} = 10$ kW. Customers are assumed for simplicity to have the same activities and loads with parameters given in Table 2.1, but multiple types are supported in the simulation. The tables dictating the probability of a customer scheduling an activity to start in each hour of the day are not shown for space reasons, but are available in the repository.

We used satellite-measured solar irradiance from a location on Lake Victoria, Uganda, spanning 2004 to 2019 at one minute resolution, to generate irradiance forecasts and real-

³Code: <https://github.com/Energy-MAC/pscc2020-load-limiting>.

Table 2.1: Activity parameters (time in minutes)

Activity	Watts	Min Time	Max Time	Compl. Val.	Int. Cost
Electronics 1	50	5	15	0.5	1
Electronics 2	75	30	180	4	2
TV	50	30	240	1	5
Lighting 1	300	5	260	2	10
Lighting 2	450	5	30	2	6
Microwave	650	2	10	2	5
Hair dryer	1800	2	17	2	5
Clothes Washer	500	30	60	3	5
Clothes Dryer	2500	45	60	3	5
Dishwasher	1200	60	90	3	5

izations [24]. This region has active development of energy-scarce, isolated microgrids and exhibits daily variation in irradiance. In each experimental trial, we randomly select one year to use as realization, and draw S times with replacement from the remaining fifteen years for forecasts. We created sample load forecasts by simulating the customer load model with random activity schedules S times.

Controller Efficacy

In this experiment, we use $N = 7$ customers and $S = 15$ forecast scenarios with 48 hour horizons and simulate the RHC for 28 days. These, and the parameters defined in previous sections, comprise the experiment parameters. For each trial, we draw a random start day, random DER configuration, random customer activity schedule, and random forecasts and realizations as confounding variables. For these confounding variables, we compare each of the five controllers as independent variables: no load limit, proportional feedback, deterministic forecast, the two-stage model, and the approximate dynamic programming model. For each of these, we simulate the RHC and define three key performance metrics on the outputs: the value of the quadratic objective function (2.10) applied to realized consumption averaged over the 4 hour decision interval, the net customer utility derived from their successful completion and interruption of the loads, and the per unit average service availability index (ASAI), which is the fraction of time power was available averaged across customers (ASAI [13]). The objective values and customer utility are average per user per 4 hour time-step. These results are shown in Fig. 2.3, where the bar height is the median and the range shows the 5th and 95th percentile values across trials. We conducted 150 trials, observing that the coefficients of variation across trials for the performance metrics stabilize by 100 trials.

The predictive controllers significantly improve customer utility and power availability,

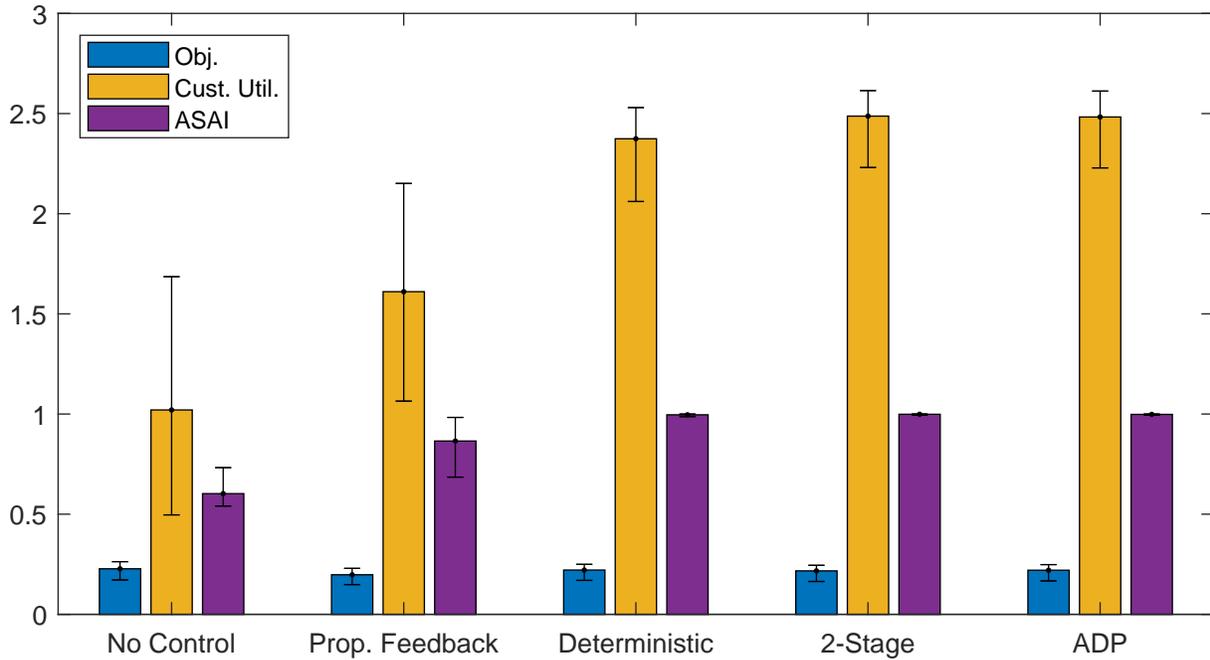


Figure 2.3: Key performance metrics across trials.

but they do not improve the quadratic objective measure they explicitly maximize. We expect this is due to model mismatch where the controller assumes customers adjust load *exactly* to the limits, but they in fact reduce load *below* the limits. This is consistent with Fig. 2.4, which shows that the predictive controllers overestimate the objective even when accounting for forecast uncertainty. As expected, the no control case has the highest mean load because there is no curtailment. The objective values correspond closely to the mean load and are only slightly lower because the quadratic term is small, especially at normally low load, which leads us to conclude that the greater consumption drives the higher objective value.

The key result is that despite the model mismatch, optimizing for the simple quadratic value of consumption produces an outcome that allows customers to respond to scarcity with lower interruption costs and greater utility. This may not be the outcome in some cases, for example if customers have very high-value and high-power, daytime loads, but if this is known to the microgrid operator, this can be addressed with weights in the benefit function. Further, gains on the feedback controller could be tuned to give better performance in particular cases, although it would likely be challenging to set gains that are effective across a wide variety of cases.

Among the predictive controllers, the two stochastic approaches yield similar results to each other; however, Fig. 2.3 shows they tend towards slightly higher utility and minimally higher ASAI than the deterministic. The deterministic overestimates the objective relative

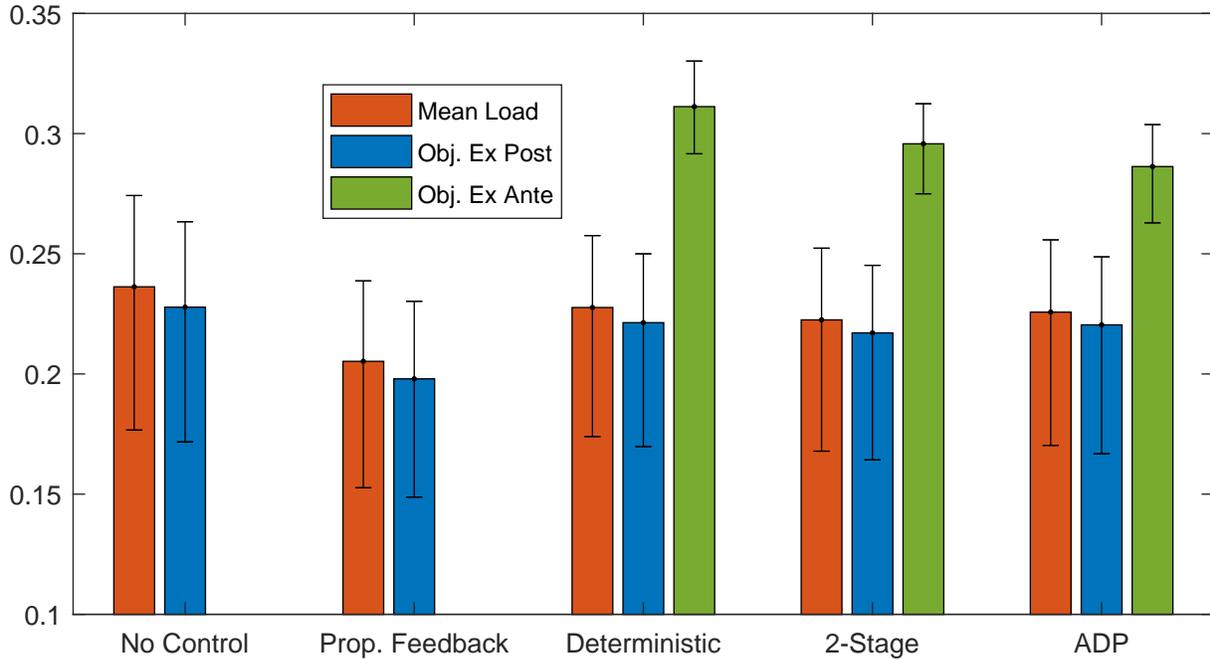


Figure 2.4: Objective *ex post* and *ex ante* values with mean load.

to the 2-stage, which also overestimates relative to the ADP formulation if the forecast values are independent in time. This can be shown theoretically and is supported empirically in Fig. 2.4. In supplemental figures provided in the GitHub repository, we show that the stochastic approaches impose load limits more of the time but at higher and less restrictive levels when they are imposed, resulting in reductions in interruption costs. Essentially, they perform some effective hedging, but the benefits are small and the deterministic approach provides satisfactory performance.

Computational Tractability

To test computational performance, we varied the number of customers $N \in \{5, 15\}$, the number of scenarios $S \in \{5, 15\}$, and the time-steps in the forecast $T \in \{12, 24, 36\}$, and recorded the time for each formulation of the decision algorithm to converge to a solution. Table 2.2 shows the median time over 20 trials with random forecasts and initial states. We observed that for larger products of NS approaching the range of 300, the solver does not reliably converge within an hour, so we do not show results for problems of this size. We observed that for these problems that do not converge, approximate solutions are reached relatively quickly, but that thousands of successive iterations in the branch-and-bound algorithm continue with minimal improvements.

The results show that the approximate dynamic program generally takes longer to solve, but that the two-stage solution exhibits poor scaling with the forecast horizon. Both formu-

Table 2.2: Timing results (seconds)

N	S	T	Time: Det.	Time: 2 Stage	Time: DP
5	5	12	0.01	0.01	0.97
5	5	24	0.01	0.02	1.9
5	5	36	0.01	0.03	2.6
5	15	12	0.01	0.03	5
5	15	24	0.01	0.06	9.3
5	15	36	0.01	0.09	13
15	5	12	0.01	0.04	1.7
15	5	24	0.01	0.07	3.5
15	5	36	0.02	0.11	4.7
15	15	12	0.01	0.11	30
15	15	24	0.01	0.24	47
15	15	36	0.02	0.39	61

lations are tractable for a real-time control scheme for products of NS up to around 100 with a forecast horizon of 24-36 hours. The tractability for larger products of NS requires more research into solver customization, appropriate solution tolerance, and convex relaxations.

2.5 Conclusions

This paper develops a mathematical framework for managing electricity consumption in energy-constrained microgrids by scheduling load limits to improve the availability and value of electricity service. We propose two techniques for incorporating stochastic forecasts into the decision to schedule load limits, and show how these can be modelled as mixed-integer programs. We find that both improve metrics of the value of electricity service and are tractable with an out-of-the-box MIQP solver for microgrids on the order of 15 customers, but that a deterministic approach, using only a single forecast, yields comparable performance improvements in our particular test case but with much lower computational complexity. Our modelling approach and simulation environment contribute a foundation for exploring different formulations of value and mechanisms to allocate scarce electricity supply.

Chapter 3

Pricing and energy trading in peer-to-peer zero marginal-cost microgrids

Abstract

Efforts to utilize 100% renewable energy in community microgrids require new approaches to energy markets and transactions to efficiently address periods of scarce energy supply. In this paper we contribute to the promising approach of peer-to-peer (P2P) energy trading in two main ways: analysis of a centralized, welfare-maximizing economic dispatch that characterizes optimal price and allocations, and a novel P2P system for negotiating energy trades that yields physically feasible and at least weakly Pareto-optimal outcomes. Our main results are 1) that optimal pricing is insufficient to induce agents with batteries to take optimal actions, 2) a novel P2P algorithm to addresses this while keeping private information, 3) a formal proof that this algorithm converges to the centralized solution in the case of two agents negotiating for a single period, and 4) numerical simulations of the P2P algorithm performance with up to 10 agents and 24 periods that show it converges on average to total welfare within 0.1% of the social optimum in on the order of 10s to 100s of iterations, increasing with the number of agents, time periods, and total storage capacity.

Nomenclature

Sets and indices

\mathcal{C} Set of agents/consumers indexed by n .

- \mathcal{U} Subset of agents that propose quantities (q -agents), indexed by k .
- \mathcal{X} Subset of q -agents that have exited the negotiation.
- \mathcal{Y} Subset of q -agents that are still negotiating.
- \mathcal{V} Subset of agents that respond with price (π -agents), indexed by v .
- \mathcal{T} Set of time periods, indexed by t .
- \mathcal{B} Set of batteries, indexed by i .
- \mathcal{G} Set of generators, indexed by g .

Variables

- $d_{n,t}$ Local power consumption of agent n at time t .
- $p_{g,t}^s$ PV power production from generation g at time t .
- $p_{i,t}^b$ Discharge power from battery i at time t .
- $s_{i,t}$ State of charge of battery i at time t .
- $\lambda_{x,t}^y$ Dual variable of constraint y , where $x \in \{n, i, g\}$.
- $q_{k,t}$ Proposed quantity by q -agent k to receive from the π -agent at time t .
- $q'_{k,t}$ Projection of $q_{k,t}$ to the feasible quantity set by the π -agent at time t .
- β Auxiliary variable used to project $q_{k,t} \rightarrow q'_{k,t}$.
- α_n Binary state, true iff agent n prefers current offer to no trade at all.
- $o_{k,t}$ Binary state, true iff a proposed quantity from q -agent k at time t is “oscillating”, i.e., not monotonically increasing or decreasing over 3 iterations.
- η_k Binary state, true iff q -agent k requests to settle.

Parameters

- $\bar{P}_{g,t}^s$ Max. PV power from generation g at time t .
- $\bar{P}_{i,t}^b$ Max. rate of charge/discharge of battery i at time t .
- $\bar{S}_{i,t}^b$ Max. energy capacity of battery i at time t .
- $\delta_k^{(0)}$ Initial size of step-limiting constraint on a q -agent k .
- γ Shrinking rate of step-limiting constraint $\in (0, 1)$.

Notation: The utility functions of the agents with respect to their local demand are denoted by $U_n(d_n)$, marginal utility $\partial U_n/\partial d_n$ by $g_n(d_n)$, and its inverse $g_n^{-1}(d_n):=h_n(\pi_n)$. Bold symbols represent a vector or a collection of points, e.g., $\mathbf{q}_k \equiv \{q_{k,t}\}_{t \in \mathcal{T}}$. The power and energy units are kW and kWh, and ΔT is the time step duration in hours. The symbols \neg , \vee , and \wedge denote logical negation, OR, and AND respectively.

3.1 Introduction

Due to declining technology costs and a drive to reduce carbon emissions, 100% renewable electricity grids systems are receiving increasing attention. California’s 2018 Senate Bill 100, for example, sets a large-scale 100% renewables target for 2045. At community scales, 100% renewable microgrids for resilience and energy access in rural areas have become competitive with hybrid solutions with fuel-based generators, and can be preferable in cases where emissions, fuel logistics, or generator maintenance are strong concerns.

Novel pricing mechanisms for 100% renewable systems are not yet well-developed, but we contend they will become increasingly important for policy-makers and practitioners. For example, extending the current paradigm where load serving entities procure electricity at the lowest cost to meet inflexible demand, implies that zero (short-run) energy costs would lead to a zero (short-run) price [104]. However, Fripp *et al* [47] have shown by including demand-side participation in a capacity expansion model for Hawaii that dynamic electricity pricing to consumers results in non-zero prices and is increasingly important for maximizing welfare in 100% renewable systems.

In this paper we contribute a general theoretical analysis of pricing in 100% renewable plus storage systems, characterizing optimal price dynamics and the challenges energy storage presents for standard bidding mechanisms. We focus on zero marginal cost renewables, with a solar generation case in simulations. We then propose an approach for community microgrids where individual “prosumers” with solar and storage could interact informally in a “peer-to-peer” (P2P) system to negotiate energy trades and form dynamic electricity prices and allocations that approximate optimal outcomes.

In the next two paragraphs we briefly review the relevant literature in this space, arguing specifically that the case of peer-to-peer trading in 100% renewable plus storage systems requires more analysis and innovation. P2P systems are valuable for grid resiliency, renewables integration, electricity access in less developed regions, and individual participation in electricity systems [55, 44, 102, 116, 115]. Early work proposed centrally coordinated energy trading between distributed energy resources (DERs) where the generation and battery storage are fully controllable [72, 22]. In [76], the authors lay the foundation for defining the physical and virtual layers required for a pooling-based system, but the paper does not develop bidding strategies for agents and assumes the microgrid remains connected to the main grid; [85] describes stochastic P2P methods to match prosumers with consumers and share profit, however no storage is considered in the model; similarly, [93, 83, 118, 117] propose ADMM-based methods to determine dispatch and/or pricing in P2P settings without

storage; [5] proposes a cooperative coalition scheme, based on energy reduction that can achieve savings for their participants, but the method requires a pricing scheme that varies if the demand surpasses a given threshold to encourage power reduction and; [54] describes a P2P architecture accounting for network charges, but does not consider the temporal aspects of storage, while [130] uses comfort constraints for the next time step as a limited approach to address this.

Energy storage is fundamental to 100% renewable systems, and some papers do incorporate it into P2P algorithms. For example, [94] proposes a game-theoretic model, while a few papers define specific rules for battery charge/discharge cycles based on the traded quantity at each time step [68, 66]. An additional group of algorithms consider storage and either exchange shadow prices, employ ADMM-based bilateral trading mechanisms or additional cost-sharing methods [45, 124, 67, 63, 1, 69]. However, none of these address the context of scarce, zero marginal-cost renewables coupled with storage.

We address the gap in the literature with a novel P2P approach that can describe informal interactions between prosumers negotiating trades for electricity in a finite time horizon setting. We assume P2P agents individually derive private utility from energy use, and iterate on price and quantity bids with their peers until convergence. In contrast to a centralized approach, the P2P approach maintains the privacy of individual utility functions and addresses the complexity of bidding storage while converging to the centralized solution in special cases. The main contributions are:

- (1) In Section 3.2, to characterize the optimal price and explain the challenge of coordinating storage through centralized pricing, we formulate a centralized optimization model and highlight several non-trivial observations, such as optimal prices not uniquely determining battery dispatch decisions.
- (2) In Section 3.3, we define a novel P2P system with minimal prescriptive rules through which agents with private information exchange offers to arrive at a trade, and theoretically prove convergence for 2-agents with a single time period. In the proposed algorithm, each agent can easily manage storage and state of charge constraints in its private decision.
- (3) In Section 3.4, we find that the P2P approach converges to a solution in all of 1200 general cases that were simulated, with welfare outcomes on average within 0.1% of the centralized while maintaining the privacy of utility functions, with a worst-case divergence of 8% that can arise from longer time horizons and relatively large storage capacities.

3.2 Centralized Welfare Maximization Approach

In this section, we define a model for optimal energy dispatch over a finite time horizon, analyze the solution for relevant insights into P2P electricity markets, and illustrate its dependence on energy storage through example. The model applies a utility maximization framework. For analysis of the optimal dispatch, we take the perspective of a benevolent

central operator and assume knowledge of the individual utility functions. In practice this could be the perspective of a DER aggregator or a distribution system operator; however, it is difficult to know utility functions in practice, and this issue is a fundamental motivation for exploring peer-to-peer markets in the first place. We use a deterministic approach, where decisions are made off of a single, expected forecast of solar generation without hedging for uncertainty, and note this can be suboptimal to stochastic approaches. We also present an idealized battery model for simplicity, but show in Appendix 3.A that the key insights still hold when we incorporate constant charge and discharge inefficiencies, self-discharge, and asymmetric power constraints. We also assume the battery is not required to achieve a final state-of-charge, but the model can easily include this constraint without loss of generality as long as the final state-of-charge is feasible. This model provides a baseline for comparing decentralized approaches, and could be extended to other DERs such as electric vehicles in the context where the generation is zero marginal-cost and energy constraints are relevant. In the case where the DERs introduce additional costs or utility functions, such as fuel costs or non-concave utility functions, the theoretical results may not hold.

The key theoretical insight we provide is that in the presence of energy storage, the dispatch cannot be controlled by price alone. Specifically, we show that if individuals act independently to maximize their utility in the presence of an optimal price, there is no guarantee that their corresponding target power injections will be feasible and satisfy power balance. This highlights that ensuring feasibility is an important requirement of decentralized mechanisms. We describe why this is not trivial in the presence of storage, and also derive equations describing the optimal power and price trajectories.

Utility maximization model

The model (3.1) is similar in structure to a standard discrete-time, centralized energy management system. The central constraint is matching supply and demand on the time scale of hours, while we assume that droop-like control of power converters is necessary and sufficient to adjust any power imbalance in the short-term.¹ We include operational constraints on energy storage, but not the network constraints,² and assume strictly concave utility functions $U_{n,t}$.

$$\min_{\mathbf{p}, \mathbf{d}, \mathbf{s}} - \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{C}} U_{n,t}(d_{n,t}) \quad (3.1a)$$

$$\text{s.t. } \pi_t : \sum_{n \in \mathcal{C}} d_{n,t} = \sum_{i \in \mathcal{B}} p_{i,t}^b + \sum_{g \in \mathcal{G}} p_{g,t}^s, \forall t \in \mathcal{T} \quad (3.1b)$$

¹This technology has been extensively studied, and while important questions remain for large system stability with high penetrations of converter-interfaced generation, a variety of techniques have been validated for microgrids, see e.g. [100].

²The model can be extended to include linearized power flow and line loading constraints, which would add some complexity without affecting the main results; however, full AC power flow equations would destroy the constraint linearity (and convexity) that the analysis relies on.

$$\lambda_{g,t}^s : 0 \leq p_{g,t}^s \leq \bar{P}_{g,t}^s, \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (3.1c)$$

$$\lambda_{n,t}^{d,-} : -d_{n,t} \leq 0, \forall n \in \mathcal{C}, \forall t \in \mathcal{T} \quad (3.1d)$$

$$\lambda_{i,t}^b : -\bar{P}_{i,t}^b \leq p_{i,t}^b \leq \bar{P}_{i,t}^b, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.1e)$$

$$\lambda_{i,t}^c : 0 \leq s_{i,t} \leq \bar{S}_{i,t}, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.1f)$$

$$s_{i,t} = s_{i,t-1} - p_{i,t}^b \Delta T, \forall i \in \mathcal{B}, \forall t \in \mathcal{T}. \quad (3.1g)$$

This allows battery constraints to be time-varying but typically \bar{P}^b and \bar{S} are static. The dual variables of the respective constraints are indicated before the colon. For compactness, we use a single variable to represent the difference in upper and lower bound duals, $\lambda := \lambda^+ - \lambda^-$. The initial state of charge $s_{i,0}$ is a parameter. We eliminate the constraint (3.1g) and decision variables $s_{i,t}$ by solving for it as $s_{i,t} = s_{i,0} - \Delta T \sum_{\tau \leq t} p_{i,\tau}^b$ and substituting this into (3.1f).

Theoretical analysis

Firstly, note that all constraints in (3.1) are affine, thereby satisfying the linearity constraint qualification (LCQ). This implies that for a locally optimal primal solution, there exists a set of dual variables satisfying the Karush-Kuhn-Tucker (KKT) conditions. Secondly, as all $U_{n,t}$ are concave, the problem is convex. Any point satisfying the KKT conditions is thus globally optimal and strong duality holds.

Remark 1 (*Dual decomposition into private decisions*): *The Lagrangian dual of the centralized problem (3.1) is separable and equivalent to the sums of Lagrangian duals for constrained individual welfare maximization for a price equal to π_t . This allows interpretation of π_t as the electricity price. Assuming the utility functions are concave, the Lagrangian dual problem gives the optimal price and total welfare.*

The Lagrangian of (3.1) can be written as:

$$\begin{aligned} \mathcal{L}(d, p^s, p^b, \pi, \lambda) = & \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{C}} \left(-U_{n,t}(d_{n,t}) + (\pi_t - \lambda_{n,t}^{d,-}) d_{n,t} \right) \\ & + \sum_{g \in \mathcal{G}} \left((\lambda_{g,t}^s - \pi_t) p_{g,t}^s + \lambda_{g,t}^{s,+} \bar{P}_{g,t}^s \right) + \sum_{i \in \mathcal{B}} \left((\lambda_{i,t}^b - \pi_t) p_{i,t}^b \right. \\ & \left. + \lambda_{i,t}^c \left(s_{i,0} - \Delta T \sum_{\tau \leq t} p_{i,\tau}^b \right) - (\lambda_{i,t}^{b,+} + \lambda_{i,t}^{b,-}) \bar{P}_{i,t}^b - \lambda_{i,t}^{c,+} \bar{S}_{i,t} \right). \end{aligned} \quad (3.2)$$

We define individual utility/profit-maximization problems for each of the consumers, PV, and battery operators for an electricity price as in (3.3)-(3.5).

$$W_n(\pi) := \min_{d_n \geq 0} \sum_t -U_{n,t}(d_{n,t}) + \pi_t d_{n,t} \quad (3.3)$$

$$W_g(\pi) := \min_{p_g^s} \sum_t -\pi_t p_{g,t}^s \quad \text{s.t.} \quad (3.1c) \quad (3.4)$$

$$W_i(\pi) := \min_{p_i^b} \sum_t -\pi_t p_{i,t}^b \quad \text{s.t.} \quad (3.1e) - (3.1g) \quad (3.5)$$

Denoting their Lagrangians by $\mathcal{L}_n, \mathcal{L}_g, \mathcal{L}_i$, one can show that

$$\mathcal{L}(d, p^s, p^b, \pi, \lambda) = \sum_{n \in \mathcal{C}} \mathcal{L}_n(d_n, \pi) + \sum_{i \in \mathcal{G}} \mathcal{L}_g(p_g^s, \lambda_g^s, \pi) \sum_{i \in \mathcal{B}} \mathcal{L}_i(p_i^b, \lambda_i^b, \lambda_i^c, \pi). \quad (3.6)$$

As W_g and W_i are linear programs, strong duality holds for these subproblems, and the Lagrangian dual problem is

$$\begin{aligned} & \max_{\pi, \lambda} \inf_{d, p^s, p^b} \mathcal{L}(d, p^s, p^b, \pi, \lambda) \\ &= \max_{\pi} \sum_{n \in \mathcal{C}} W_n(\pi) + \sum_{g \in \mathcal{G}} W_g(\pi) + \sum_{i \in \mathcal{B}} W_i(\pi). \end{aligned} \quad (3.7)$$

By strong duality (3.7) gives the optimal objective value with its maximizer π^* equal to the optimal price. However, as we establish later, the optimal p_i^b for (3.5) is not necessarily unique, meaning that broadcasting an optimal price to individual agents does not necessarily satisfy constraint (3.1b) and clear the market; i.e., primal feasibility is not guaranteed.

Remark 2 *For all $t \in \mathcal{T}$, the following relations hold true at optimum and characterize the optimal price*

$$\pi_t^* = \partial U_{n,t}(d_{n,t}^*) / \partial d_{n,t} + \lambda_{n,t}^{*,d,-}, \quad \forall n \in \mathcal{C} \quad (3.8a)$$

$$= \lambda_{i,t}^{*,b} - \Delta T \sum_{\tau \geq t} \lambda_{i,\tau}^{*,c}, \quad \forall i \in \mathcal{B} \quad (3.8b)$$

$$= \lambda_{g,t}^{*,s}, \quad \forall g \in \mathcal{G}. \quad (3.8c)$$

Each of the equalities follow from the stationarity conditions of (3.1). We interpret the dual variable π_t^* as the price by **Remark 1** and note from (3.8b) that it depends on the cumulative future shadow prices of the storage capacity constraint. Eq. (3.8a) requires $U_{n,t}$ to be differentiable for equality but can be replaced by the subdifferential of $U_{n,t}$ otherwise.

Remark 3 *If at time t , a utility function for at least one customer is differentiable and strictly increasing on \mathbb{R}^+ , then at optimum, the price is strictly positive and solar production is at its maximum.*

This follows from **Remark 2** and the properties of strictly increasing functions:

$$\exists n \ni \partial U_{n,t}(d_{n,t}^*) / \partial d_{n,t} > 0 \quad \forall d_{n,t} \Rightarrow \pi_t^* > 0 \Rightarrow \lambda_{g,t}^{*,s} > 0.$$

By complementary slackness, $\lambda_{g,t}^{*,s} > 0 \Rightarrow p_t^{*,s} = \bar{P}_t^s$. This is intuitive as it is better to supply any benefiting consumer than curtailing available solar. This also implies that solar generation can be removed as a decision variable and set to the available resource in this case.

Remark 4 *The optimal price evolves as*

$$\pi_{t+1}^* - \pi_t^* = \lambda_{i,t+1}^{*,b} - \lambda_{i,t}^{*,b} + \Delta T \lambda_{i,t}^{*,c}. \quad (3.9)$$

This follows from **Remark 2** by expanding the expression $\pi_{t+1}^* - \pi_t^*$. This captures the price trajectory, from which price volatility can be analyzed. Note that both $\lambda_{i,t}^{*,b}$, $\lambda_{i,t}^{*,c}$ can be less than 0. We will use (3.9) for our analysis in **Remark 5**.

Remark 5 (*Non-uniqueness of decentralized battery dispatch*): *There are non-trivial optimal prices π^* such that the optimal individual battery dispatch $W_i(\pi^*)$ is not-unique.*

This can be observed in a simple example. Suppose $T=5$, $\Delta T=1$, $\bar{P}_{i,t}^b \equiv 3$, $\bar{S}_{i,t}=10$, $s_{0,i}=5$, and $\pi^* = [1, 1, 2, 3, 1]$. One can verify that $p_i^b = [-0.5, -0.5, 3, 3, 0]$, $p_i^c = [0, -1, 3, 3, 0]$, and $p_i^d = [-1, -1, 3, 3, 1]$ are all optimal solutions with a net cost of -14 . Here, equal prices imply there is no change in cost to shift energy from one period to another and the constraints allow this shift. More formally, if an optimal solution is not on any of the constraint boundaries (3.1e)-(3.1f) at t and $t+1$, then it will not be unique because not being on the boundary implies 1) that $\lambda_{i,t}^{*,b}$, $\lambda_{i,t+1}^{*,b}$, $\lambda_{i,t}^{*,c}$, and $\lambda_{i,t+1}^{*,c}$ are all 0, so $\pi_{t+1}^* = \pi_t^*$ by (3.9), therefore $p_{i,t}^{\prime,b} = p_{i,t}^{*,b} + \varepsilon$ and $p_{i,t+1}^{\prime,b} = p_{i,t+1}^{*,b} - \varepsilon$ have equivalent net cost $\forall \varepsilon$ without affecting the solution at other times; and 2) that this perturbation is feasible for sufficiently small $|\varepsilon| > 0$. Note that this condition is overly restrictive and not necessary for non-uniqueness; in particular energy may be shifted between non-consecutive time periods, and only particular combinations of constraints between the time periods need to be non-binding rather than all constraints. Equal prices between time periods may indicate non-uniqueness, but the optimal solution may still be unique if the constraints do not allow a perturbation to remain feasible. In Appendix 3.A, we show that this applies with a variation in the optimal price profile even when battery inefficiencies and self-discharge are considered.

The consequence of **Remark 5** is that in general, an optimal price is not sufficient to yield individual battery dispatch decisions with the optimal quantity, meaning that a system operator cannot control dispatch outcomes solely by broadcasting a price signal or adequately forecast the decentralized response to price. Even when all utility functions are strictly concave so that the solution to the centralized problem is unique and corresponds to an optimal price π^* , there are (likely common) conditions whereby an individual battery operator's decision in response to π^* does not satisfy the constraint (3.1b). Intuitively, if the price is constant between two successive periods, a battery operator would be indifferent to selling more energy in one period versus the next, so their dispatch is not unique and there is no guarantee that the dispatch will meet demand.

The previous observation implies that extending standard centralized market mechanisms to systems with energy storage faces limitations. If a centralized energy market limits entities with storage to submitting a single curve of price and quantity for each time-period, it is likely to result in suboptimal outcomes to the utility maximization problem and even in infeasibility. Although not shown here, we expect this result extends to load that can be shifted without cost, and to storage models with constant charge or discharge inefficiencies.

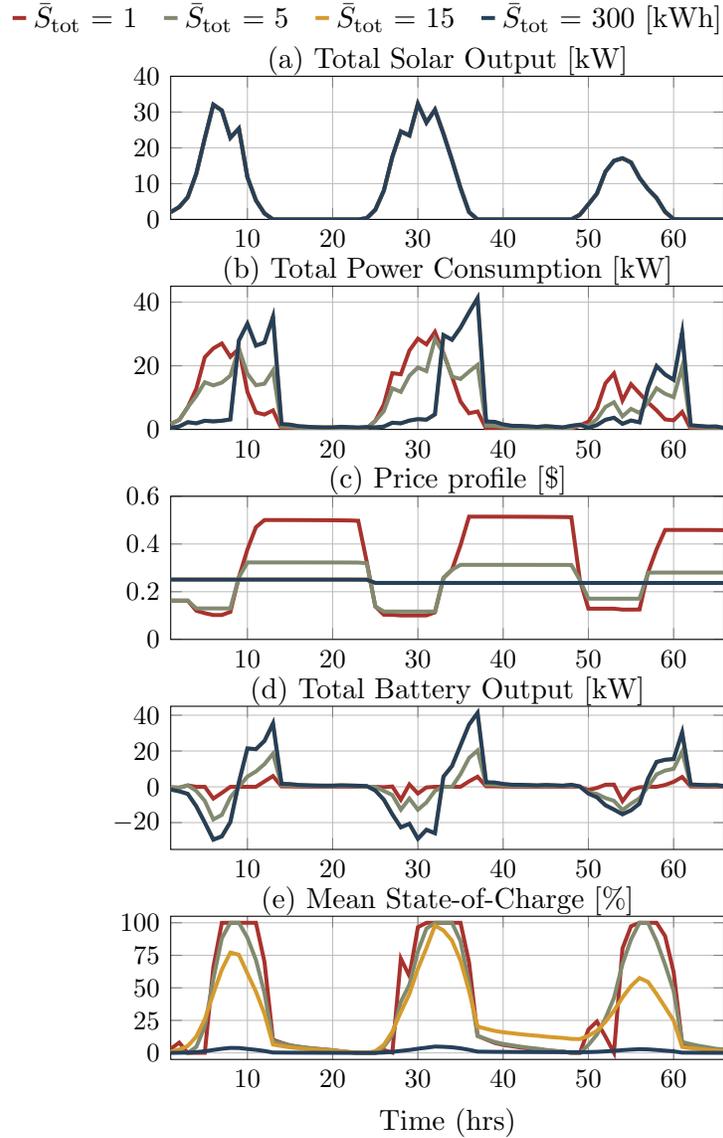


Figure 3.1: Optimal profiles for the centralized approach with 10 agents, $\bar{P}_i^b = 10$, and demand elasticity $\in [-3, -2]$ for all agents. Note: The plots for $\bar{S}_{\text{tot}}=15, 300$ kWh overlap in (b), (c), (d).

P2P approaches where agents explicitly agree on quantity are a potential opportunity for addressing this challenge.

Example optimal trajectories and the effects of storage

To show how the PV profile and storage capacity affect the optimal trajectories of (3.1), we simulate scenarios with total storage capacity varying $\bar{S}_{\text{tot}} \in \{1, 5, 15, 300\}$ kWh and

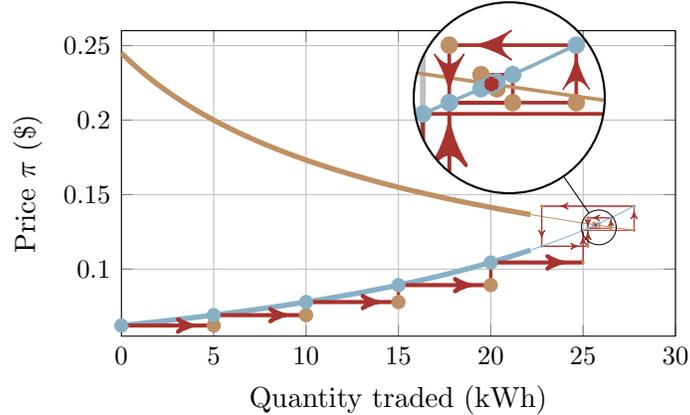


Figure 3.3: Convergent trajectory under the dynamic step-limiting constraint where the standard cobweb model would diverge.

[52] and observe that classical results show the process can diverge. We therefore, consider an additional *dynamic step-limiting constraint* on the process to ensure convergence, which could be thought of as a behavioral tendency of agents or an explicit rule to be imposed by a bidding platform. We assume agents are matched *a priori* and that offers are synchronized so as to simplify the analysis and presentation, but posit that the process can be generalized to capture more informal interaction between agents.

As a starting point, consider an interaction between two agents who are “prosumers” with private solar and storage systems and who individually derive private value from energy use. Most likely, there exists a trade that makes both agents better off. An intuitive way for the agents to find such a trade is for one to start by proposing a quantity (either positive or negative) and for the other to respond with a price. The first agent would likely reassess the quantity they would seek at that price, propose a new quantity, and so on. This iterative process is described by the cobweb model illustrated in Fig. 3.2. The equilibrium is the intersection of supply and demand curves arising from the utility functions. This is the optimum of the utility maximization model but the process converges to this point if and only if the magnitude of the slope of the demand curve exceeds that of the supply curve at the equilibrium [52].

We modify the cobweb model to ensure convergence even when this condition is not met by including a step-limiting constraint, illustrated in Fig. 3.3. This constraint assumes (or enforces) that agents will not adjust their quantity offers by more than some threshold each iteration, and that this threshold shrinks if the quantity is “oscillating.” We generalize to consider multiple agents proposing quantities (called q -agents) to agents who respond with price (called π -agents). The agents exchange vectors of quantity and price for each period over a finite time horizon. To simplify the analysis, we assume a single π -agent interacts with multiple q -agents. In practice, there would likely be multiple π -agents, and q -agents would select one or more π -agents to negotiate with, based on their expectation of the outcome of

the negotiation, but this matching problem is beyond the scope of this paper.

We present formal decision models for the q -agents and the π -agents, and define an iterative process that guarantees physically feasible and at least weakly Pareto-optimal outcomes (i.e., no agents are worse off). We prove theoretically that the process converges to within a tolerance of the centralized solution for the 2-agent, single time step case, and demonstrate convergence using simulations for the general case in the next section. These results show that an informal, decentralized, peer-to-peer negotiation process is capable of approximating the centralized welfare maximization problem, and offers a specific approach that could be implemented on a software platform and evaluated in practice. In addition, and contrary to the centralized approach, this negotiation process does not require full information exchange between agents, since private utility functions are hidden and only trading of quantities and prices are required for the execution of the algorithm.

We denote the set of π -agents with \mathcal{V} and q -agents with \mathcal{U} , such that $\mathcal{C} = \mathcal{U} \cup \mathcal{V}$ and $\mathcal{U} \cap \mathcal{V} = \emptyset$. We index the q -agents by $k \in \mathcal{U}$ and the single π -agent as v , $\mathcal{V} = \{v\}$. The q -agents may exit the process early, which we track by partitioning \mathcal{U} into exited agents \mathcal{X} and negotiating agents \mathcal{Y} , and updating these dynamically. All constraints are implicitly defined $\forall t \in \mathcal{T}$.

Optimization problem for the π -agent

The π -agent receives a set of requested quantities $\mathbf{q} = \{\mathbf{q}_k\}$ from each $k \in \mathcal{Y}$ (positive means k receives energy), with $\mathbf{q}_k = \{\mathbf{q}_{k,t}\}$, which may not be feasible. The π -agent first projects \mathbf{q} to a feasible \mathbf{q}' by keeping reference to a quantity $\hat{\mathbf{q}}$ known to be feasible to all agents; \mathbf{q}' is restricted to lie on the line connecting \mathbf{q} and $\hat{\mathbf{q}}$ defined by (3.10d), where $\beta=0$ yields the requested \mathbf{q} and $\beta=1$ the known feasible $\hat{\mathbf{q}}$. Thus, minimizing $\beta \geq 0$ selects the closest point to \mathbf{q} satisfying the constraints:

$$\min_{\mathbf{d}_v, \mathbf{p}_{v,t}^s, \mathbf{s}_v, \mathbf{q}', \beta} \quad \beta \quad (3.10a)$$

$$\text{s.t.} \quad d_{v,t} + \sum_{k \in \mathcal{Y}} q'_{k,t} + \sum_{k \in \mathcal{X}} q_{k,t} = p_{v,t}^s + p_{v,t}^b \quad (3.10b)$$

$$0 \leq \beta \leq 1 \quad (3.10c)$$

$$q'_{k,t} = \beta \hat{q}_{k,t} + (1 - \beta) q_{k,t} \quad (3.10d)$$

$$\text{and constraints (3.1d) - (3.1g)}. \quad (3.10e)$$

We maintain that $\hat{\mathbf{q}}$ is feasible for all agents. Before any agents exit, $\mathcal{X} = \emptyset$ and $\hat{\mathbf{q}} = 0$ is feasible, so we initialize with $\hat{\mathbf{q}} = 0$ and update $\hat{\mathbf{q}}$ as agents exit at feasible points. As shown below, \mathbf{q} is necessarily feasible for each q -agent, and their constraints are convex, so any point connecting two feasible points is feasible, and in particular \mathbf{q}' . It is also possible to include additional constraints in this optimization problem, for example, to ensure power flow feasibility. Extending this model to include a network feasibility validation as constraints over $q'_{k,t}$ is a promising direction for future work.

Next, the π -agent solves their utility maximization problem to obtain $\boldsymbol{\pi} = \{\pi_t\}$ and their utility from these proposed trades. A key assumption is that they set $\boldsymbol{\pi}$ at their marginal utility; i.e., they bid according to a competitive market strategy and cannot exercise market power. This is likely to hold in practice if there are sufficiently many π -agents the q -agents can access; however, we recommend a more careful analysis of market power in the scope of a “many-to-many” extension to this work. The maximization problem is:

$$\min_{\mathbf{d}_v, \mathbf{p}_v^b, \mathbf{s}_v} - \sum_{t \in \mathcal{T}} U_{v,t}(d_{v,t}) \quad (3.11a)$$

$$\text{s.t.} \quad \pi_t : d_{v,t} + \sum_{k \in \mathcal{Y}} q'_{k,t} + \sum_{k \in \mathcal{X}} q_{k,t} = p_{v,t}^s + p_{v,t}^b \quad (3.11b)$$

$$\text{and constraints (3.1d) – (3.1g).} \quad (3.11c)$$

As in the centralized model, the price is given directly by the stationarity condition with $\lambda_v^{d_v^-} = 0$:

$$\pi_t = \partial U_{v,t}(d_{v,t}^*) / \partial d_{v,t}.$$

Lastly, the π -agent checks whether its utility from this potential trade is at least as high as its optimal utility from no trade (specifically solving the same problem with $\mathbf{q}' = 0$), and sets a binary variable α_v true if so, and false otherwise. This α_v signals whether v would prefer \mathbf{q}' to no trade. We denote the entire decision as $\mathcal{P}_v^\pi : (\mathbf{q}, \hat{\mathbf{q}}) \mapsto (\mathbf{q}', \boldsymbol{\pi}, \alpha_v)$.

Optimization problem for q -agents

The k -th q -agent makes the decision $\mathcal{P}_k^q : (\boldsymbol{\pi}, \mathbf{q}'_k, \boldsymbol{\delta}_k) \mapsto (\mathbf{q}_k, \alpha_k, \eta_k)$, where α_k carries the analogous meaning to α_v , η_k signals whether they are “satisfied”, \mathbf{q}'_k is the subset of \mathbf{q}' for k , and $\boldsymbol{\delta}_k$ is the step-limiting constraint restricting the q -agent to select something close to the offer \mathbf{q}' . The decision is:

$$\min_{\mathbf{d}_k, \mathbf{q}_k, \mathbf{p}_k^b, \mathbf{s}_k} \sum_{t \in \mathcal{T}} -U_{k,t}(d_{k,t}) + \pi_t q_{k,t} \quad (3.12a)$$

$$\text{s.t.} \quad d_{k,t} - p_{k,t}^s - p_{k,t}^b - q_{k,t} = 0 \quad (3.12b)$$

$$|q_{k,t} - q'_{k,t}| \leq \delta_{k,t} \quad (3.12c)$$

$$\text{and constraints (3.1d) – (3.1g).} \quad (3.12d)$$

Agent k requests to finalize the trade and exit if their (not necessarily unique) optimal \mathbf{q}_k is close enough to the offer \mathbf{q}'_k , where the distance is determined by a small ε :

$$\eta_k = \begin{cases} \text{True} & \text{if } |q_{k,t} - q'_{k,t}| \leq \gamma \varepsilon \\ \text{False} & \text{otherwise.} \end{cases} \quad (3.13)$$

The exit condition includes the constant $\gamma \in (0, 1)$ to simplify the statement of Theorem 6, but could be modified with an update to the bound in the theorem. An alternative criterion based on whether the utilities from these offers are close enough could also be used but would affect the bound.

Iterative Algorithm

Algorithm 1: Bounded cobweb iteration for a single π -agent and multiple q -agents.

Result: Energy trades $(\boldsymbol{\pi}_k^*, \mathbf{q}_k^*)$ for each agent $k \in \mathcal{C}$.

Initialization: Define the π -agent $v \in \mathcal{C}$ and the parameters $\gamma \in (0, 1)$, $\varepsilon > 0$, initial step-limit $\delta^{(0)} > \gamma\varepsilon$ and max iterations M ;

Set $i \leftarrow 1$, $(\mathbf{q}^{(1)}, \hat{\mathbf{q}}) \leftarrow (0, 0)$, $\{\delta_{k,t}^{(1)}\} \leftarrow \delta^{(0)}$, $\mathcal{X} \leftarrow \{0\}$, and $\mathcal{Y} \leftarrow \mathcal{C} \setminus \{v\}$;

while $\mathcal{Y} \neq \emptyset$ **and** $i \leq M$ **do**

$(\mathbf{q}'^{(i)}, \boldsymbol{\pi}^{(i)}, \alpha_v) \leftarrow \mathcal{P}_v^\pi(\mathbf{q}^{(i)}, \hat{\mathbf{q}})$;

for $k \in \mathcal{Y}$ **do**

$(\mathbf{q}_k^{(i+1)}, \alpha_k, \eta_k) \leftarrow \mathcal{P}_k^q(\boldsymbol{\pi}^{(i)}, \mathbf{q}'^{(i)}, \boldsymbol{\delta}_k^{(i)})$;

$\mathbf{o}_k^{(i)} \leftarrow f^o(\mathbf{q}_k^{(i+1)}, \mathbf{q}_k^{(i)}, \mathbf{q}_k^{(i-1)})$;

if η_k **then** $\boldsymbol{\delta}_k^{(i+1)} \leftarrow \boldsymbol{\delta}_k^{(i)}$ **else** $\boldsymbol{\delta}_k^{(i+1)} \leftarrow f^\delta(\boldsymbol{\delta}_k^{(i)}, \mathbf{o}_k^{(i)})$;

end

if $\alpha_j \forall j \in \mathcal{Y} \cup \{v\}$ **then**

$\hat{\mathbf{q}} \leftarrow \mathbf{q}'^{(i)}$;

for $k \in \mathcal{Y}$ **where** η_k **do**

$\mathcal{Y} \leftarrow \mathcal{Y} \setminus \{k\}$, $\mathcal{X} \leftarrow \mathcal{X} \cup \{k\}$, $(\boldsymbol{\pi}_k^*, \mathbf{q}_k^*) \leftarrow (\boldsymbol{\pi}^{(i)}, \mathbf{q}'^{(i)})$

end

end

$i \leftarrow i + 1$

end

The negotiation algorithm is presented in Algorithm 1. At each iteration, q -agents submit their energy quantity requests to the π -agent based on the last price and quantity offered by the π -agent. The q -agents are allowed to exit only when all agents have declared $(\boldsymbol{\pi}, \mathbf{q}')$ preferable to no trade through α (i.e., $\alpha_k = \text{True} \forall k \in \mathcal{U}$), guaranteeing that trades are at least weak-Pareto improvements. Importantly, the step-limit $\boldsymbol{\delta}$ is shrunk by $\gamma \in (0, 1)$ if the quantity is “oscillating” (see Fig. 3.2), defined by the binary state $\mathbf{o}^{(i)}$ as the quantity not monotonically increasing or decreasing over 3 iterations, with $\mathbf{o}^{(1)} = 1$ and update maps f^o and f^δ :

$$f^o : o_{k,t}^{(i)} = \neg(q_{k,t}^{(i+1)} > q_{k,t}^{(i)} > q_{k,t}^{(i-1)} \vee q_{k,t}^{(i+1)} < q_{k,t}^{(i)} < q_{k,t}^{(i-1)}),$$

$$f^\delta : \delta_{k,t}^{(i+1)} = (1 - o_{k,t}^{(i)}) \delta_{k,t}^{(i)} + o_{k,t}^{(i)} \gamma \delta_{k,t}^{(i)}.$$

This shrinking step-limit prevents the divergent case of the cobweb model [52].

Optimality of the two-agent, single time step case

In this subsection we prove that Algorithm 1 converges within an ε tolerance in finite iterations to the socially optimal quantity in the case of only two agents with single time

horizon. We ignore storage in this case, as it can equivalently be treated as solar production for $\mathcal{T} = \{1\}$, and drop the time index t for brevity. We assume the the solar production is greater than zero for at least one agent, and that each agent's marginal utility of consumption $\partial U_n(d_n)/\partial d_n$ is strictly monotonically decreasing on $[0, \infty)$ and decreasing asymptotically to zero.

Note that $q = d_k - p_k^s = -d_v + p_v^s$, and the *unconstrained* demand and supply curves are defined as $g_k \equiv \partial U_k(q)/\partial d_k$ and $g_v \equiv \partial U_v(q)/\partial d_v$. Thus, g_k is monotonically decreasing and g_v is monotonically increasing. Without the step-limiting constraint (3.12c), the problem \mathcal{P}_k^q for the q -agent has a closed form solution:

$$q^\dagger = g_k^{-1}(\min(g_k(-p_k^s), \pi)) \equiv h_k(\pi), \quad (3.14)$$

where g_k^{-1} denotes the inverse of g with domain $(0, g_k(-p_k^s)]$. With the step-limiting constraint, the solution is

$$q = \begin{cases} q^\dagger & \text{if } |q^\dagger - q'| \leq \delta \\ q' + \delta & \text{if } q^\dagger > q' + \delta \\ q' - \delta & \text{if } q^\dagger < q' - \delta. \end{cases} \quad (3.15)$$

The projection step reduces to $q' = \min(p_v^s, q)$, and the π -agent's price is given by $\pi = g_v(q')$.

The optimal quantity of the centralized problem q^* is the unique fixed point of the iteration if $q^* < p_v^s$ or if $q^* = p_v^s$ and $g_k(p_v^s) = g_v(p_v^s)$. Indeed, note that $-p_k^s \leq q^* \leq p_v^s$ by the constraints, and that $h_k(\pi^*) \equiv q^*$. If $q^{(i)} = q^*$, then $q' = q^*$ and $\pi^{(i)} = g_v(q^*) = \pi^* - \lambda_s^{*,d,-}$ by (3.8a). When $q^* < p_v^s$ or $g_k(p_v^s) = g_v(p_v^s)$, then we have $\lambda_s^{*,d,-} = 0$ and $\pi^{(i)} = \pi^*$, and hence $q^\dagger = q^*$ with $q^{(i+1)} = q^*$. Otherwise, $\lambda_s^{*,d,-} > 0$ and $\pi^{(i)} < \pi^*$, so $q^\dagger > q^*$ by the strict monotonicity of h_k , and $q^{(i+1)} > q^*$, so it is not a fixed point. In other words, the fixed point is the intersection of the curves g_v and g_k , as shown in Fig. 3.2. Since both curves are strictly monotonic, this fixed point is unique. If they do not intersect on $[-p_k^s, p_v^s]$, then q^* is only a fixed point if it is $-p_k^s$.

Lemma 1 (Movement towards equilibrium) *At any iteration i , if $q^* < p_v^s$, then $q' \leq q^* \Leftrightarrow q^{(i+1)} \geq q'$ and $q' \geq q^* \Leftrightarrow q^{(i+1)} \leq q'$. Moreover, $q' \leq q^* \Leftrightarrow q^{(i+1)} \geq q^{(i)}$ and $q' \geq q^* \Leftrightarrow q^{(i+1)} \leq q^{(i)}$.*

Proof: We prove this by showing the forward direction of the first set of statements $q' \leq q^* \Rightarrow q^{(i+1)} \geq q'$ and $q' \geq q^* \Rightarrow q^{(i+1)} \leq q'$. Each of these statements implies the converse of the other is true, establishing the reverse direction. We use the same approach for the second set of statements.

Let $\pi = g_v(q')$. If $q^* < p_v^s$, then $g_v(q^*) = \pi^*$ and if $q' \leq q^* \Rightarrow \pi \leq \pi^* \Rightarrow h_k(\pi) \geq h_k(\pi^*) \Rightarrow q^\dagger \geq q^*$ because g_v is monotonically increasing and h_k is monotonically decreasing. Thus, $q^\dagger \geq q^* \geq q' \Rightarrow q^{(i+1)} = \min(q^\dagger, q' + \delta) \Rightarrow q^{(i+1)} \geq q'$. By the same logic, $q' \geq q^* \Rightarrow q^{(i+1)} \leq q'$. Moreover, $q' \leq q^* < p_v^s \Rightarrow q' = q^{(i)}$; therefore, $q^{(i+1)} \geq q' \Rightarrow q^{(i+1)} \geq q^{(i)}$. Finally, because $q' = \min(p_v^s, q^{(i)}) \leq q^{(i)}$, and by showing that $q' \geq q^* \Leftrightarrow q^{(i+1)} \leq q'$, it follows that $q' \geq q^* \Rightarrow q^{(i+1)} \leq q^{(i)}$. \square

Lemma 2 (Entry to the oscillatory mode) *If the system is not in the oscillatory mode at iteration i , then $\exists l > 0$ such that if the algorithm does not terminate at iteration $s < i + l$, it will be in the oscillatory mode at $i + l$.*

Proof: First, consider the case $q^* < p_v^s$. We will show the case when $q^{(i)} < q^*$, since the other case is analogous. By Lemma 1, $q^{(i+1)}$ moves towards the equilibrium and $\delta^{(i+1)}$ is not reduced when moving in the same direction. Thus, for some $j > i$, $q^{(j)} \geq q^*$ (with $q^{(j-1)} \leq q^*$); hence, by Lemma 1, $q^{(j+1)} \leq q^{(j)}$ and we enter the oscillatory mode at $j + 1 = i + l$.

Second, we consider the case of $q^* = p_v^s$. Observe that when eventually $q^{(j)} \geq p_v^s$, it will be projected back to $q' = \min(q^{(j)}, p_v^s) = p_v^s$ to ensure feasibility for the π -agent. Then, since there is no intersection of marginal utility curves in the interior, it implies that $g_k(q') \geq g_v(q') = \pi^{(j)}$, and hence, the q -agent again requests $q^{(j+1)} \geq p_v^s$, that gets projected back to $q' = p_v^s$. Thus, since repeated values of q are received, it will enter the oscillatory mode (and eventually converge to p_v^s). \square

Lemma 3 (Boundedness of distance from the equilibrium) *Assume $q^* < p_v^s$ and suppose the system is in the oscillatory mode at iteration i . Then, $|q' - q^*| < \gamma^{-1} \delta^{(i)}$.*

Proof: Let $q'^{(i-1)}$ denote the offer from the π -agent at the previous iteration. We first prove the case when $q' \geq q^*$ by contradiction. To this end, assume $q' - q^* \geq \gamma^{-1} \delta^{(i)}$. This implies $q^{(i)} - q^* \geq \gamma^{-1} \delta^{(i)}$ because $q^{(i)} \geq q'$ by $q' = \min(q^{(i)}, p_v^s)$. This in turn implies $q'^{(i-1)} \geq q^*$ by the step-limiting constraint (3.12c) at the previous iteration (observe $\delta^{(i-1)} \leq \gamma^{-1} \delta^{(i)}$). Then, $q'^{(i-1)} \geq q^* \Rightarrow q^{(i)} \leq q'^{(i-1)} \leq q^{(i-1)}$ by Lemma 1. For the system to be oscillating with $q^{(i)} \leq q^{(i-1)}$, either $q^{(i+1)} > q^{(i)}$ (which contradicts $q' \geq q^*$ by Lemma 1), or we have equality at $q^{(i-1)} = q^{(i)}$ or $q^{(i)} = q^{(i+1)}$ (which implies $q^* = q'$, by the unique fixed point, contradicting $q' - q^* \geq \gamma^{-1} \delta^{(i)}$).

We show the second case, $q' < q^*$, directly. We have $q^{(i+1)} > q'$ by Lemma 1 and $q^{(i)} = q'$ because $q^* \leq p_v^s$. Thus, the oscillating mode implies $q^{(i)} \leq q^{(i-1)}$. It holds that $q'^{(i-1)} \geq q^*$: if $q^{(i-1)} > p_v^s$, then $q'^{(i-1)} = p_v^s \geq q^*$. Alternatively, if $q^{(i-1)} \leq p_v^s$, then $q'^{(i-1)} = q^{(i-1)} \geq q^{(i)}$ and $q^{(i-1)} \geq q^*$ by Lemma 1. This implies $q^{(i)} \geq q^* - \delta^{(i-1)}$ by the step-limiting constraint at $i - 1$. Since $q' = q^{(i)}$ and $\delta^{(i-1)} \leq \gamma^{-1} \delta^{(i)}$, this proves the lemma. \square

Lemma 4 (Arbitrarily small δ) *For any tolerance $\varepsilon > 0$, there exists K indicating the number of finite iterations, such that $\delta^{(K)} \leq \varepsilon$.*

Proof: Let $m^{(i)}$ denote the cumulative number of times the system has been in the oscillatory mode at iteration i , with $m^{(1)} = 0$ and $m^{(i+1)} = m^{(i)} + o^{(i)}$. Thus $\delta^{(i)} = \delta^{(0)} \gamma^{m^{(i)}}$. Following Lemma 2, for any $m > 0$, if $m^{(i)} = m$, there exists $l > 0$ such that $m^{(i+l)} = m + 1$. Thus, we can make $m^{(i)}$ arbitrarily large with sufficient iterations, and therefore, $\delta^{(i)} = \delta^{(0)} \gamma^{m^{(i)}}$ can be made arbitrarily small. \square

Lemma 5 (Termination) *If the algorithm terminates at iteration i due to the stopping criterion, then $|q'^{(i)} - q^*| < \varepsilon$.*

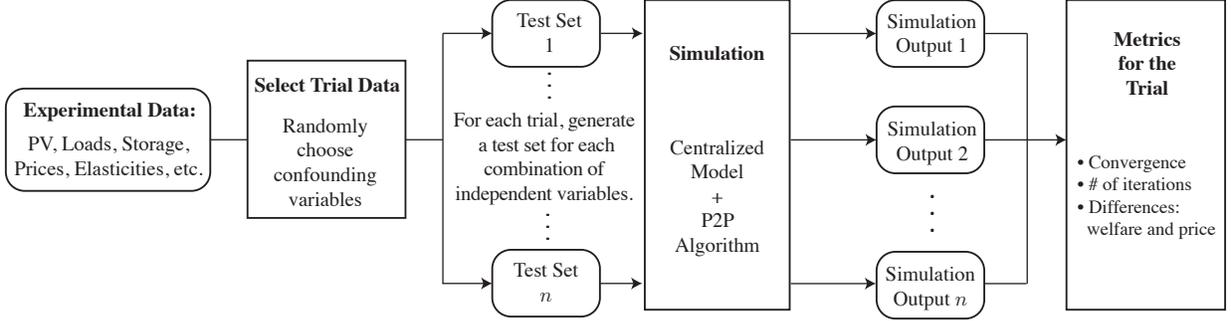


Figure 3.4: Trial procedure flowchart.

Proof: We will prove the case when $q^* < p_v^s$: First, consider the case when $q^{(i-1)} < q^*$ (and hence $q'^{(i-1)} = q^{(i-1)}$), then it follows from Lemma 1 that $q^{(i)} \geq q^{(i-1)}$. There are two cases, if (i) $q^{(i)} \geq q^*$, then $q^* < q'^{(i)} \leq q^{(i)}$ and it follows directly that $|q'^{(i)} - q^*| \leq |q'^{(i)} - q'^{(i-1)}| \leq \gamma\varepsilon < \varepsilon$, since the algorithm terminated at iteration i . Alternatively, if (ii) $q^{(i)} \leq q^*$, we have $q'^{(i)} = q^{(i)}$. From (3.15) and the intersection of g_k and g_v , it must be true that $q^{(i)} = q^{(i-1)} + \delta^{(i-1)}$, and thus $|q^{(i)} - q^{(i-1)}| = \delta^{(i-1)} \leq \gamma\varepsilon$. Now, we have two cases: (a) if the system was oscillating, we have that $q'^{(i-2)} > q^*$ and so $q'^{(i-2)} > q'^{(i-1)}$ from Lemma 1. It is also true that $\delta^{(i-1)} = \gamma\delta^{(i-2)}$ and that $|q'^{(i-1)} - q'^{(i-2)}| \leq \delta^{(i-2)}$, therefore, $|q'^{(i)} - q^*| \leq |q'^{(i)} - q'^{(i-2)}| \leq (1 - \gamma)\delta^{(i-2)} = (1 - \gamma)\gamma^{-1}\delta^{(i-1)} \leq (1 - \gamma)\gamma^{-1}\gamma\varepsilon = (1 - \gamma)\varepsilon < \varepsilon$. Or (b), if the system was not oscillating, this implies that $\delta^{(i-1)} = \delta^{(i-2)}$, and from the same argument as before satisfying that $|q^{(i)} - q^{(i-1)}| \leq \gamma\varepsilon$. This is a contradiction, since that would also imply that $|q'^{(i-1)} - q'^{(i-2)}| \leq \delta^{(i-2)} = \delta^{(i-1)} \leq \gamma\varepsilon$, hence terminating before i . Second, for the case when $q^{(i-1)} > q^*$, the proof is equivalent to the first case, but considering the special instance that if $q^{(i-1)} > p_v^s$, then $q'^{(i-1)} = p_v^s$, but is still larger than q^* so the same idea holds by invoking Lemma 1.

Finally, we address the case when $q^* = p_v^s$. As described in Lemma 2, the algorithm will get stuck at $q' = p_v^s$ for more than two iterations, terminating the algorithm. Since $p_v^s = q^*$, then $|q'^{(i)} - q'^{(i-1)}| = |q'^{(i)} - q^*| = 0 < \varepsilon$. \square

Theorem 6 (Optimality of Algorithm 1) *For 2 agents with strictly concave utility functions, $T = 1$, and with sufficiently large max iterations M , Algorithm 1 returns a quantity within ε of the centralized optimum q^* .*

Proof: By Lemma 4, if we set $M \geq K$, then the algorithm will terminate due to the stopping criterion in at most K iterations, and by Lemma 5, the quantity is within ε of q^* . \square

In Fig. 3.3 we depict a case that converges to the centralized solution via the step-limiting constraint. This case otherwise diverges based on the classic result [52] without the step-limiting constraint. This theoretical analysis provides the foundation for extending the algorithm to multiple agents $|\mathcal{C}| > 2$ with finite time horizon $T > 1$. We explore the behavior of the algorithm numerically for such cases in the next section.

3.4 Computational Experiments and Simulations

To provide additional insight into the algorithm performance, we perform two simulation-based computational experiments following the methodology and nomenclature in [57]. The simulation flowchart for both experiments are summarized in Fig. 3.4. The first, examines how the two algorithm parameters γ , $\delta^{(0)}$ affect the rate of convergence. The second, tests convergence for the unproven cases for $|\mathcal{C}| > 2$ and $T > 1$, and studies the effect of battery energy and power capacity on convergence and explores welfare differences between the centralized and P2P approaches.⁴ In all experiments we use hourly load and PV profiles from Pecan Street [95], and constant price elasticity utility functions fit to the baseline load with elasticities random on $[-1.5, -0.5]$ as in section 3.2.

Effect of parameters γ and $\delta^{(0)}$ on convergence

In this experiment, we study the convergence rate for the 2-agent, single period case. We systematically vary $\gamma \in \{0.05, 0.1, \dots, 0.95\}$, $\delta^{(0)} \in \{0.1, 0.2, \dots, 2\}$ kWh as independent variables, generating 380 unique pairs of $(\gamma, \delta^{(0)})$. For each pair, we execute 100 trials with randomly generated confounding variables (the two load profiles, hour of the year, price elasticities, and solar power between zero and twice the load) and compute the iterations to convergence. We use a stopping tolerance $\varepsilon = 10^{-3}$ for all trials.

The results in Fig. 3.5 show that γ has a strong effect on the convergence rate and exhibits a minimum for $\gamma \in [0.3, 0.5]$ that is consistent across the different ranges of $\delta^{(0)}$, and that the algorithm converges in on the order of 10-20 iterations on average for γ in the middle range. We found that $\delta^{(0)}$ was not very significant in influencing the number of iterations except for causing an increase at especially small values, suggesting the parameter ought to be set to a relatively large value. A possible intuition behind the effect of γ is that especially small values shrink the box too quickly away from the equilibrium, while large values do not shrink rapidly enough.

Performance for unproven cases

To study the performance in the general (multi-agent) case, we vary the total battery capacity $\bar{S}_{\text{tot}} \in \{15, 25, 40, 80, 300\}$ kWh and the maximum rate of charge/discharge of the battery $\bar{P}_i^b \in \{1, 2, 4, 8\}$ kW as independent variables, yielding 20 distinct pairs.

Similar to section 3.4, for each pair $(\bar{S}_{\text{tot}}, \bar{P}_i^b)$ we execute 60 trials ($60 \times 20 = 1200$ simulations), randomly selecting PV and load profiles, price elasticities, an hour of the year, $T \in \{1, 12, 24\}$ hours, and number of agents $N \in [2, 10]$. A battery capacity fraction is assigned uniformly to each agent (and then normalized) from the total battery capacity. The PV profiles are scaled so the total PV energy equals the total baseline load energy, and $(\gamma, \delta^{(0)}) = (0.5, 0.5)$.

⁴The experiment parameters, data files, and MATLAB code to reproduce the experiments can all be found at <https://github.com/Energy-MAC/TSG-P2P-Pricing>.

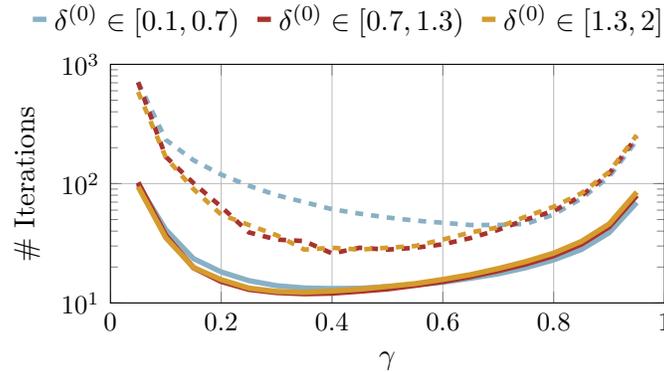


Figure 3.5: Effect of γ on the number of iterations to convergence. The solid lines show the mean over all trials where $\delta^{(0)}$ lies in the interval specified in the legend. The dashed lines show the maximum.

Convergence performance

All of the 1200 treatments converge to a solution. The average iterations required to convergence is 112.5, with a standard deviation of 257.3 and a median of 61. We observe that larger time horizons with more agents require more iterations for the algorithm to converge.

Effect of battery parameters

The effect of battery capacity on convergence is illustrated via boxplots in Fig. 3.6, depicting the distribution of the number of iterations for convergence against battery capacity (with outliers omitted). In general, a higher battery capacity requires more iterations to converge. The intuition being that with higher battery availability, the flexibility for each agent to adapt to successive trades increases, thus requiring more iterations. This highlights the importance of storage in a P2P setting and the effect on the implementation of energy trading algorithms. In contrast, the maximum charge/discharge rate of the battery does not significantly affect the number of iterations. This is expected, because given the demand profiles, a maximum rate of 1 kW is usually enough to achieve a trade.

Welfare comparison

In order to compare the total welfare of all agents for the centralized and the iterative P2P algorithm, we classify the trials by grouping the time horizon. The statistics of welfare difference percentages ΔW_p and absolute welfare differences ΔW are presented in Table 3.1. We note that most of the entries for ΔW_p are lower than 0.1%, i.e., in the range of numerical tolerance used for MATLAB based optimizers. These results indicate that in most cases the centralized welfare is close to that of the proposed algorithm. However, there exist cases

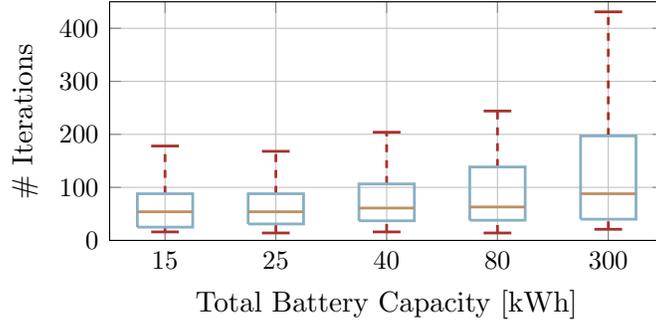


Figure 3.6: Number of iterations to converge with varying battery capacity.

Table 3.1: Welfare difference statistics for different time horizons.

	T	1	12	24
	#Simulations	300	420	476
ΔW_p	Mean [%]	0.023	0.001	0.072
	Std [%]	0.079	0.004	0.706
	Max [%]	0.558	0.034	7.758
ΔW	Mean [\$]	0.004	0.002	0.319
	Std [\$]	0.015	0.007	3.338
	Max [\$]	0.086	0.057	36.717

Table 3.2: Welfare difference statistics for the special instance considered in Section 3.4.

	W_{no} [\$]	W_{centr} [\$]	W_{P2P} [\$]	ΔW [\$]	ΔW_p [%]
ag-1	11.923	19.804	14.292	5.512	27.833
ag-2	6.617	16.785	9.079	7.706	45.910
ag-3	2.784	2.933	3.516	-0.583	-19.877
ag-4	202.124	202.906	202.920	-0.014	-0.007
ag-5	164.184	229.159	203.345	25.814	11.265
π -ag	1.633	1.711	3.429	-1.718	-100.409
Total	389.265	473.298	436.581	36.717	7.758

when $T > 1$, for which although the algorithm converges, the welfare is significantly different from the centralized solution.

Special instance

In this section we explore one instance where there is a considerable mismatch ($\Delta W = \$36.71$) between the welfare values obtained from the two approaches. This occurs for $T = 24$, $N = 6$,

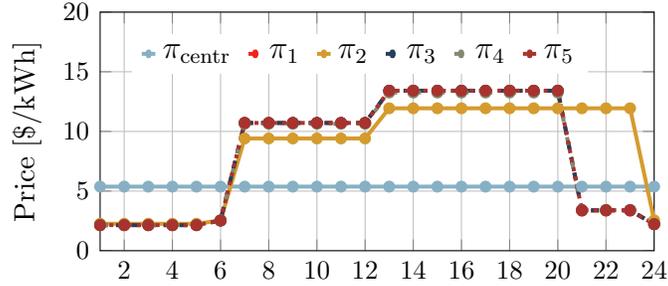


Figure 3.7: Centralized and P2P algorithm price profiles for the special instance of considerable difference in welfare.

and low total battery capacity of $\bar{S}_{\text{tot}} = 15$ kWh. The key difference is that the prices for the agents in the algorithm are significantly different than those obtained in the centralized solution, as observed in Fig. 3.7.

This simulation converges in 59 iterations, when agent-1 (ag-1) exits the algorithm. However, at iteration 32, agent-2 exits based on its stopping criteria, while the remaining agents continue trading, before exiting at iterations 59, 58, 58, and 55 respectively with similar price profiles, as indicated in Fig. 3.7. The consumption profiles and hence the individual welfare of each agent are thus significantly different from the centralized solution. Table 3.2 summarizes the total welfare (consumption + trading) of each agent using the centralized and P2P algorithm. The welfare for the no-trading case W_{no} , is also presented for comparison. A closer inspection reveals that while that agent-3 and the π -agent are better off in the P2P case, agents 1, 2, and 5 are well placed in the centralized case. Furthermore, for this particular simulation, exiting earlier is not optimal for agent-2, although the price is lower than the other q -agents.

Summary

The simulation results in this section highlight the main contributions of this work:

- (1) The P2P algorithm achieves similar welfare results as the centralized approach in most of the cases (Table 3.1), with the caveat that an early exit by some agents may introduce sub-optimality, in which case different agents end up as winners and losers relative to the social optimum (Table 3.2), but all agents are better off than no trading.
- (2) More flexibility for the agents via larger storage or longer time horizons increases the number of iterations (Fig. 3.6).
- (3) The expected number of iterations is minimized by setting the shrinking rate of the step-size γ around 0.4 (Fig. 3.5).
- (4) Real-time prices in a zero marginal-cost system arise from the marginal utility of consumption under scarcity.

3.5 Conclusions

In this paper, we address the question of optimal pricing and mechanisms for achieving optimal dispatch in microgrids with scarce, zero marginal-cost energy resources. We contribute a novel analysis of a centralized economic dispatch with welfare maximization that uses a Lagrangian dual decomposition to state the equilibrium optimal price conditions and show a previously unstated result that although optimal prices can induce unique and optimal consumption profiles and generator output, they do not yield unique or power-balanced battery dispatch decisions except in particular circumstances. Next, we propose a P2P algorithm where agents keep utility functions private and iteratively interact by exchanging price and quantity offers to arrive at mutually agreeable and weakly Pareto-optimal trades. We theoretically prove this outcome converges to the social optimum within a specified tolerance for the 2-agent case, and show via numerical experiments that the P2P algorithm converges in the multi-agent case, but we do not derive specific bounds. Although we find that the P2P algorithm obtains total welfare on average within 0.1% of the centralized solution for a wide range of parameters, significant differences in welfare and allocation can arise for longer time horizons and larger numbers of agents. We also find from simulations that the number of iterations for the P2P algorithm to converge increases with the total storage capacity, and that the P2P algorithm shrinking parameter γ impacts the number of iterations, and should be set to the neighborhood of 0.3-0.5 to minimize iterations in contexts similar to our simulations, while the initial maximum step-size $\delta^{(0)}$ is not significant.

The proposed P2P algorithm was designed to resemble an *informal* decentralized trading process where prices arise from the value of electricity consumption under scarcity. We envision it is feasible to implement such an interaction in practice via a software platform that defines the rules and aids in the iteration, or even with informal negotiation between neighbors in a community. However, there are several limitations that need to be addressed for this approach to be useful in practice. First, we do not study the impact of strategic gaming between agents, which could be significant in small markets, nor the equity of outcomes. Conducting this analysis likely requires removing the assumption that π -agents offer prices equal to their dual variables and considering their profit maximizing strategy, given expectations of q -agents' demand curves. Second, our analysis only considers strictly concave utility functions. This is a common assumption, but may not capture the discrete nature of decisions to use particular loads at small time-scales. We expect it will be difficult for researchers to derive useful theoretical insights with non-concave utility functions, but the construction of realistic utility functions and consumption decision models for use in simulation would be of tremendous value to this and related work. Third, we do not include network constraints or validate power flow. While it is relatively straightforward to validate whether a particular negotiated dispatch is feasible given a network model, the impact of binding constraints on pricing and negotiation is non-trivial and warrants further study. Fourth, a system for matching agents into smaller negotiation pools based on expected outcomes may be necessary to handle large numbers of agents, e.g., hundreds. Here, a challenge is to design suitable exit strategies for satisfied agents without compromising the inviolability of

agreements, and should also account for network constraints in creating market power (see [54]). This introduces significant complexity, where methods to certify optimality or bound the outcome are important theoretical directions for future work. Lastly, we suggest the inclusion of uncertainty via scenarios in a stochastic programming framework to deal with uncertainty in solar forecasts and load estimation. The inclusion of power flow feasibility and network validation in the P2P algorithm, and extensions to a broader class of DERs are next steps in this research.

3.A Appendix: Non-ideal battery modeling

In this appendix, we extend the analysis in Section 3.2 to a more realistic model of the battery and show that the results hold when battery inefficiencies and asymmetric charge and discharge power constraints are accounted for. We emphasize that we assume a linear model and the results may not hold for nonlinear models. However, agents may use decision models of varying complexity in practice, so understanding the implications of simplified models remains highly relevant.

The extended model replaces the net battery discharge $p_{i,t}^b$ with its positive discharge and charge components $p_{i,t}^{b,+}$ and $p_{i,t}^{b,-}$. We allow different discharge and charge power constraints $P_{i,t}^{b,+} > 0$ and $P_{i,t}^{b,-} > 0$ and assume that power is converted to and from the stored energy with charge efficiency $\sigma_i^- \in (0, 1]$ and discharge efficiency $\sigma_i^+ \in (0, 1]$, and that a battery self-discharges at a rate $(1 - \theta_i)$ proportional to the state-of-charge, with $\theta_i \in [0, 1)$. Thus, with the extended model, the problem formulation can be stated as:

$$\min_{\mathbf{p}, \mathbf{d}, \mathbf{s}} - \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{C}} U_{n,t}(d_{n,t}) \quad (3.16a)$$

$$\text{s.t. } \pi_t : \sum_{n \in \mathcal{C}} d_{n,t} = \sum_{i \in \mathcal{B}} (p_{i,t}^{b,+} - p_{i,t}^{b,-}) + \sum_{g \in \mathcal{G}} p_{g,t}^s, \forall t \in \mathcal{T} \quad (3.16b)$$

$$\lambda_{g,t}^s : 0 \leq p_{g,t}^s \leq \bar{P}_{g,t}^s, \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (3.16c)$$

$$\lambda_{n,t}^{d,-} : -d_{n,t} \leq 0, \forall n \in \mathcal{C}, \forall t \in \mathcal{T} \quad (3.16d)$$

$$\lambda_{i,t}^{b,+} : 0 \leq p_{i,t}^{b,+} \leq \bar{P}_{i,t}^{b,+}, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.16e)$$

$$\lambda_{i,t}^{b,-} : 0 \leq p_{i,t}^{b,-} \leq \bar{P}_{i,t}^{b,-}, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.16f)$$

$$p_{i,t}^{b,+} p_{i,t}^{b,-} = 0, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.16g)$$

$$\lambda_{i,t}^c : 0 \leq s_{i,t} \leq \bar{S}_{i,t}, \forall i \in \mathcal{B}, \forall t \in \mathcal{T} \quad (3.16h)$$

$$s_{i,t} = \theta s_{i,t-1} + \sigma_i^- p_{i,t}^{b,-} - (\sigma_i^+)^{-1} p_{i,t}^{b,+} \Delta T, \forall i \in \mathcal{B}, \forall t \in \mathcal{T}. \quad (3.16i)$$

In the following, we show a technical detail needed to extend **Remark 1**, state the modified forms of **Remark 2** and **Remark 4**, and show a modified example with discussion for **Remark 5**. **Remark 3** does not depend on the battery model and is unaffected.

For **Remark 1**, the analogous versions of (3.3)-(3.7) follow mechanically in the same way as the main text in Section 3.2. However, the complementarity constraint (3.16g) implies that $W_i(\pi)$ is now nonlinear, and thus strong duality does not necessarily hold, so an additional technical step is needed. We define a relaxed problem $W'_i(\pi)$ that drops the complementarity constraint to become a linear program. Under the assumption that $\pi_t \geq 0, \forall t$, which holds in particular for π^* by **Remark 2**, we show below that the problems have equal optimal objectives $W'_i(\pi) = W_i(\pi)$. Therefore, we can substitute $W'_i(\pi)$ for $W_i(\pi)$ into (3.7), and again rely on strong duality in $W'_i(\pi)$ to establish π^* as the optimal price.

To show that $W'_i(\pi) = W_i(\pi)$ if $\pi_t \geq 0, \forall t$, we take $(p_i'^{b,+}, p_i'^{b,-})$ to be any any optimal solution to $W'_i(\pi)$, and compute a particular corresponding solution $(p_i^{*,b,+}, p_i^{*,b,-})$. We verify below that $(p_i^{*,b,+}, p_i^{*,b,-})$ is both feasible for the original problem and optimal for the relaxed problem, and therefore it is optimal for the original problem and each problem's optimal objectives are equal. The alternative solution is defined specifically to be the solution that has equal net charge/discharge as $(p_i'^{b,+}, p_i'^{b,-})$ while satisfying the complementarity constraint:

$$p_{i,t}^{*,b,+} = \sigma_i^+ \max((\sigma_i^+)^{-1} p_{i,t}'^{b,+} - \sigma_i^- p_{i,t}'^{b,-}, 0) \quad (3.17a)$$

$$p_{i,t}^{*,b,-} = (\sigma_i^-)^{-1} \max(\sigma_i^- p_{i,t}'^{b,-} - (\sigma_i^+)^{-1} p_{i,t}'^{b,+}, 0). \quad (3.17b)$$

By construction, $(\sigma_i^-)^{-1} p_{i,t}^{*,b,+} - \sigma_i^- p_{i,t}^{*,b,-} = (\sigma_i^-)^{-1} p_{i,t}'^{b,+} - \sigma_i^- p_{i,t}'^{b,-}$, and thus the stored energy trajectory is identical and (3.16h)-(3.16i) are satisfied. The power constraints (3.16e)-(3.16f) are satisfied because $0 \leq p_{i,t}^{*,b,+} \leq p_{i,t}'^{b,+}$ and $0 \leq p_{i,t}^{*,b,-} \leq p_{i,t}'^{b,-}$. The complementarity constraint (3.16g) can be checked by inspection of the different cases of the max functions. Finally, $p_{i,t}^{*,b,+} - p_{i,t}^{*,b,-} \geq p_{i,t}'^{b,+} - p_{i,t}'^{b,-}, \forall t$, so if $\pi_t \geq 0$, then $-\sum_t \pi_t (p_{i,t}^{*,b,+} - p_{i,t}^{*,b,-}) \leq -\sum_t \pi_t (p_{i,t}'^{b,+} - p_{i,t}'^{b,-})$ and $(p^{*,b,+}, p^{*,b,-})$ is optimal for the relaxed problem. \square

As an aside, note that if $\sigma_i^+ \sigma_i^- < 1$ and $\pi_t > 0, \forall t$, then the relaxation is exact. This is because $(p^{*,b,+}, p^{*,b,-}) \neq (p'^{b,+}, p'^{b,-})$ implies $\sum_t \pi_t (p_{i,t}^{*,b,+} - p_{i,t}^{*,b,-}) > \sum_t \pi_t (p_{i,t}'^{b,+} - p_{i,t}'^{b,-})$, which is a contradiction, therefore $(p_i^{*,b,+}, p_i^{*,b,-}) = (p_i'^{b,+}, p_i'^{b,-})$.

The analogous form of **Remark 2** follows from the stationarity conditions of (3.16), and is

$$\pi_t^* = \partial U_{n,t}(d_{n,t}^*) / \partial d_{n,t} + \lambda_{n,t}^{*,d,-}, \forall n \in \mathcal{C} \quad (3.18a)$$

$$= \lambda_{i,t}^{*,b,+} - \Delta T (\sigma_i^+)^{-1} \sum_{\tau \geq t} \theta^{\tau-t} \lambda_{i,\tau}^{*,c}, \forall i \in \mathcal{B} \quad (3.18b)$$

$$= -\lambda_{i,t}^{*,b,-} - \Delta T \sigma_i^- \sum_{\tau \geq t} \theta^{\tau-t} \lambda_{i,\tau}^{*,c}, \forall i \in \mathcal{B} \quad (3.18c)$$

$$= \lambda_{g,t}^{*,s}, \forall g \in \mathcal{G}. \quad (3.18d)$$

Remark 4 again follows from **Remark 2**, and takes the form

$$\theta \pi_{t+1}^* - \pi_t^* = \theta \lambda_{i,t+1}^{*,b,+} - \lambda_{i,t}^{*,b,+} + \Delta T (\sigma_i^+)^{-1} \lambda_{i,t}^{*,c} \quad (3.19a)$$

$$= \lambda_{i,t}^{*,b,-} - \theta \lambda_{i,t+1}^{*,b,-} + \Delta T \sigma_i^- \lambda_{i,t}^{*,c}. \quad (3.19b)$$

These dynamics imply that the equilibrium optimal price, and equivalently the marginal value of consumption at optimum, will evolve depending on the battery inefficiencies and which battery constraints are active at the optimum.

For **Remark 5**, consider again the example from the main text with $T = 5$, $\Delta T = 1$, $\bar{P}_{i,t}^{b,+} \equiv 3$, $\bar{S}_{i,t} = 10$, $s_{0,i} = 5$. Take $\sigma_i^+ = 0.95$, $\sigma_i^- = 0.9$, $\theta_i = 0.98$, $\bar{P}_{i,t}^{b,-} \equiv 2$, and a modified price $\pi^* = [1, 1.0204, 2, 3, 1.2680]$. Analogous to the example in the main text, solutions $p^b = [-0.9587, -0.9587, 3, 3, 0]$, $p^b = [0, -1.8983, 3, 3, 0]$, and $p^b = [-1.5863, -1.5863, 3, 3, 1]$ are all optimal solutions with net benefit -13.063 .

In this example, π_2 was constructed by noting that if an optimal dispatch has the battery charging and unconstrained in power and stored energy at both $t = 1$ and $t = 2$, then by (3.19a), $\pi_{t+1} = \theta^{-1}\pi_t$; i.e., the equilibrium price trajectory must have these dynamics when the storage is charging between successive periods and is not constrained. It then follows by the same principle as described in the main text that charge power can be feasibly shifted from one period to another without affecting the cost, thus the storage dispatch is not unique. Here π_5 was chosen as $\pi_5 = (\sigma^+\sigma^-)^{-1}\theta^{-4}$ to show equilibrium conditions where a higher future price exactly compensates for the lost energy from charging. Additional discussion about the impact of storage inefficiencies on optimal pricing can be found in [46].

3.B Appendix: Constant price-elasticity utility functions

Here we describe the procedure used for developing sample utility functions from data, assuming constant price-elasticities, which is one of two common simple assumptions in pricing theory (the other being a linear demand curve / quadratic utility function). We emphasize that this was chosen for example purposes only, and that all of the analysis only assumes that utility functions are strictly concave, and would apply to logarithmic or quadratic utility functions as well.

Let the marginal utility of consumption be denoted by $g(d) = \partial U(d)/\partial d$. As described in the main text, the equilibrium price is equal to the marginal utility of consumption, i.e., $\pi^* = g(d^*)$. Then the demand function of price is given by h (as the inverse of g), where $d = h(\pi) := g^{-1}(\pi)$. The price-elasticity, which we denote as $r(\pi)$ is defined as the ratio of the percentage change in quantity to the percentage change in price, and in general depends on the price

$$r(\pi) = \frac{dh(\pi)}{d\pi} \frac{\pi}{h(\pi)}. \quad (3.20)$$

A constant price-elasticity implies $r(\pi) \equiv \hat{r}$. The general family of demand functions with this property has the form $h(\pi) = a\pi^{\hat{r}}$ for some constant a . This can be fit to an empirical price and consumption pair (π_0, d_0) by setting $h(\pi_0) = d_0$ and obtaining $a = d_0 \pi_0^{-\hat{r}}$.

Inverting this to marginal utility and integrating to utility, one obtains

$$g(d) = \pi_0 \left(\frac{d}{d_0} \right)^{\frac{1}{\hat{r}}}, \quad (3.21a)$$

$$U(d) = \frac{\hat{r} \pi_0 d^{\frac{1}{\hat{r}}+1}}{(\hat{r} + 1) d_0^{\frac{1}{\hat{r}}}} + c. \quad (3.21b)$$

However, general downward-sloping demand curves imply $r < 0$, thus $\lim_{d \rightarrow 0^+} g(d) = \infty$ and $\lim_{d \rightarrow 0^+} U(d) = -\infty$ which can be problematic for optimization solvers and is also an unrealistic extreme in practice. Thus, we modify the function to have “quasi-constant” price elasticity, by shifting the marginal utility curve to the left by a small $\delta > 0$ and compensating the exponent for the shift so that $r(\pi_0) = \hat{r}$. We also choose a c such that $U(0) = 0$. The resulting marginal utility and utility functions are

$$g(d) = \pi_0 \left(\frac{d + \delta}{d_0 + \delta} \right)^{\frac{1}{r'}}, \quad (3.22a)$$

$$U(d) = \frac{r' \pi_0 \left((d + \delta)^{\frac{1}{r'}+1} - \delta^{\frac{1}{r'}+1} \right)}{(r' + 1)(d_0 + \delta)^{\frac{1}{r'}}}, \quad (3.22b)$$

$$r' = \hat{r} \left(1 + \frac{\delta}{d_0} \right)^{-1}. \quad (3.22c)$$

Chapter 4

Decentralized optimal power dispatch and secondary control with agent-based ADMM

Abstract

In this chapter, we extend and modify the peer-to-peer (P2P) negotiation system introduced in Chapter 3 to include electricity network constraints and handle more general agent utility functions that are concave, but not necessarily *strongly* concave. To do this, we use the Alternating Direction Method of Multipliers (ADMM) to state a decentralized optimal power flow problem that can be cast as a mechanism for optimal P2P energy exchange. Unlike existing work on P2P systems that use the power flow formulations that are valid only for radial networks, the approach we propose applies to general networks, and uses an unconventional iterative linearization technique to generate locally optimal solutions to the full AC power flow problem. Additionally, we develop a novel approach for economically optimal secondary control that solves the optimal power flow problem in real-time using agents in the loop, contributing to the growing body of work on online decentralized optimal secondary control in power systems with increasing converter interfaced generation.

4.1 Introduction

Motivation and contributions

Chapter 3 studied pricing and peer-to-peer energy trading in 100% renewable power systems with storage. It contributed a characterization of the optimal price that shows battery

storage dispatch is not necessarily unique, and a novel peer-to-peer negotiation algorithm that was shown to converge to near the social optimum when agents have strongly concave utility functions. However, the analysis left several important areas to future work that we address in this chapter. The first is the behavior of the algorithm with concave, but not strongly concave, utility functions. The second is a model of the power network and associated line loading and voltage drop constraints. The third is an algorithm for “many-to-many” matching of agents. The previous chapter also highlighted the importance of a need for analysis of strategic price-manipulation behavior and the effects of forecast uncertainty; these are not addressed here.

Following the publication of Chapter 3, we observed through numerical simulation that the P2P algorithm proposed in Chapter 3 in fact does not necessarily converge in the one-to-many case when agents have piecewise-linear concave utility functions (i.e. non-smooth and not strongly concave). To address this we reviewed potential modifications to the algorithm against well-known decentralized optimization techniques using the optimal price and quantity analysis of Section 3.2. Using the understanding of average power quantities as the primal variable and price as a dual variable, we considered the classic dual ascent, primal-dual subgradient, and proximal point methods.

It can be shown that the dual ascent algorithm, which can be viewed as a centralized price adjustment process that increments the price up if there is a net shortage and down if there is a net surplus, converges to the optimal price and also gives the optimal social welfare by strong duality. However, in general for non-strictly convex functions, and in particular for batteries as highlighted in Section 3.2, for a given dual variable, the corresponding optimal primal variable is not necessarily unique. Moreover, it does not necessarily satisfy the equality constraint associated with the dual variable, which in this case is power balance, or “clearing the market”. The dual ascent algorithm can still be made to converge in this case by a diminishing step-size is used; that is, the price increment is scaled smaller and smaller at each iteration k by an amount that converges as $\frac{1}{k}$ to zero. However the problem persists that even if the optimal price is obtained, it does not necessarily induce agents with storage to inject power quantities that net out to a feasible dispatch. For this reason, we considered the dual-ascent algorithm insufficient to yield an optimal dispatch; however, a modified approach that includes shrinking step-limiting constraints on the primal quantities as in the P2P algorithm of Chapter 3 could prove fruitful. Note that the centralized price adjustment process in the dual-ascent algorithm can be viewed analogously to the π -agent in the P2P bidding algorithm, as both return a price in response to a vector of quantities. In the P2P algorithm, the π -agent uses local resources to satisfy the imbalance of the aggregated quantities and sets a price equal to the marginal opportunity cost of satisfying the imbalance. In dual ascent, the price is simply incremented up or down by the imbalance amount until convergence.

The primal-dual subgradient algorithm addresses the problem of the primal problem not necessarily converging to a feasible solution by updating the primal quantities according to a subgradient step rule. The challenge here is that the quantity given by the gradient step is not necessarily feasible, and thus the subgradient does not necessarily exist at the

next iteration. This can potentially be addressed with projected primal-dual subgradient algorithms; however, the development of an appropriate decentralized projection method requires further research.

In contrast, proximal point methods, and in particular the Alternating Direction Method of Multipliers (ADMM), addresses both of these concerns and has robust convergence properties [89, 14]. This method essentially uses a price update similar to dual-ascent, but includes a quadratic regularization penalty on the primal problem that penalizes the distance of the primal variable from a reference point. It can be proven that this reference point converges to the optimum. This algorithm is very popular for decentralized optimization for multi-agent systems and is likewise popular in the P2P and decentralized optimal power flow literature. In general, ADMM is not quite fully decentralized, as it includes centralized steps for updating price and computing a reference point, although we will show how these steps can be carried out with only neighbor-to-neighbor communication in power networks in certain formulations. Despite not being fully P2P, we find it satisfies many of the motivations for pursuing P2P negotiation highlighted in Chapter 3: it is scalable, does not require any centralized knowledge of private utility functions or DER characteristics, lends itself to a market implementation for 100% renewable systems with storage, and gives flexibility for prosumers to participate in the bidding as little or much as they wish, with the caveat that they may still face price volatility, which is always a possibility in real-time markets to the extent that pricing is not centrally controlled. Therefore, we use ADMM to develop an agent-based bidding system for coordinating DER dispatch and give network-aware formulations for DC power flow and both linear and non-linear / non-convex AC power flow. Although this is developed for the context of 100% renewable microgrids, it has applications to general decentralized optimal power flow.

This type of economic bidding is categorized of as “tertiary control” in power systems control, which is generally conducted as a forward market in liberalized electricity systems, but at the scale of large power plants. Primary and secondary control are responsible for the stability and rebalancing of the system in real-time, respectively. We find that not only does the solution of ADMM yield injection setpoints for the secondary control of DERs, but that an analogous ADMM algorithm can be executed online in real time as a novel form of secondary control, enabling an economic re-balancing. We derive this for the DC power flow model, and contributes to the growing body of work on feedback control systems that exploit the physics of the grid to solve optimal power flow problems in real time.

In summary, the contributions of the paper are:

- (1) An algorithm for a forward bidding system in agent-based smart grids based on ADMM with robust convergence properties that is scalable, privacy preserving, and plug-and-play. Although the use of ADMM in this context to achieve these features is not novel, we contribute a novel variation that yields a feasible power dispatch at every iteration. Moreover, the algorithm gives a potentially powerful approach for finding local optima of the general non-linear and non-convex AC optimal power flow problem.

- (2) An online ADMM algorithm that exploits the physics of the electricity network to yield a novel economic secondary control system. This applies to both agent-based microgrids with controllable converter-interfaced DERs and more generally to power systems with droop-controlled generation and storage.

Literature review

In this section, we provide a brief overview of ADMM, review the relevant literature on applications of ADMM to P2P energy and decentralized optimal power flow as well as on online algorithms for optimal control. We note some of the limitations of similar approaches and differentiate our approach throughout, and conclude with a summary of the gaps and the contributions to address them herein.

The Alternating Direction Method of Multipliers (ADMM) is a popular iterative algorithm for decentralized optimization and has attracted substantial attention for general decentralized optimization, and specifically for optimal power flow problems and control of distributed energy resources [82, 56, 80]. Like all iterative algorithms, ADMM starts with a candidate solution to the problem and applies a procedure repeatedly until it converges to a solution. Our contribution is twofold: first to present the application of the algorithm to agent-based control of distributed energy resources in more detail than can typically be written in a short journal article, including to show how the algorithm can be applied to power flow models of varying complexity; and second to elaborate how the algorithm could be implemented in a real-time market framework, including a discussion of how the algorithm suggests a payment structure, market timing, and warm-starting in a receding horizon framework.

The contemporary canonical reference for ADMM is [14], which at the time of this thesis has more than 17,000 citations on Google Scholar. Of particular relevance in [14] is the application of the model to “consensus” and “sharing” problems. ADMM applies very generally to optimization problems with some separable structure, but it is especially powerful for multi-agent systems where the objective is to minimize the sum of individual agents’ private cost functions subject to constraints that couple some of the agents’ variables. This applies directly to optimal power flow problems, where the cost functions are generator costs, or agents’ utility functions in our context, and the coupling constraints are the power flow equations. Consensus ADMM creates *local* targets and a *global* reference of the same coupling variable, e.g. power injection or voltage, and uses an iteration that converges to consensus, meaning the local targets and global reference are equal in the limit. The iteration involves a local update step, where each agent updates their local target based on a price and a reference point, a global update step, where a new reference point is computed by projecting the local target onto a feasible set, and a dual variable or price update, which increments a dual variable that can be interpreted as a price based on the difference between the local target and global reference. The local updates can be solved in parallel, and are typically subproblems with much smaller dimension than the full optimization problem, which makes the approach powerful for large scale systems. The global update often, but not always,

consist of a simple average of all the local targets that correspond to it. Thus, each iteration can often be solved relatively quickly, especially in parallel computing environments typically found for multi-agent systems. The algorithm has proven convergence properties [14], and converges asymptotically on order $\mathcal{O}(\frac{1}{M})$, where M is the number of iterations [73].

Many authors point to the simplicity of the algorithm and its robust convergence as its primary strengths, noting that it often performs well in practice for non-convex problems [14, 109]. Its primary weakness is slow convergence (in terms of number of iterations) to optimal solutions with high accuracy, but many authors also note that it can yield approximately optimal solutions that are “good enough” in on the order of tens of iterations (e.g. [14, 80]). Another potentially significant limitation is that the convergence guarantees for the classical algorithm assume synchronization, which can be challenging to enforce in large scale systems. Although many authors state that in practice asynchronous ADMM typically converges, theoretical convergence guarantees are an active area of research.

As noted in the previous section and in Chapter 3, ADMM is a popular approach for coordinating dispatch in P2P microgrids and more generally for decentralized optimal power flow (OPF). Of the many references focusing on P2P systems, [82] and [70] are relatively early ones in the space (2017 and 2016 respectively) and stand out as especially general. Reference [96] is contemporary in 2016 and uses ADMM for DER optimization, which although not specifically aimed at P2P systems, employs the same methodology. Many of the other references can be understood as variations of the same approach. References [63, 81, 83, 69] pay particular attention to how the centralized update of ADMM can be solved with bilateral or neighbor-to-neighbor communication, with [81] also providing an excellent computation analysis. References [83, 117] make additional assumptions on cost functions being quadratic or strongly convex that allow them to derive analytic update equations and a provably convergent accelerated iteration, respectively. Other works, e.g. [1, 45], apply ADMM more specifically to energy sharing communities in grid-connected microgrids, where the presence of the external grid is fundamental to the feasibility of the approach. Our work can be differentiated from these approaches in two primary ways.

First, these approaches all use either the DC power flow model or the DistFlow model (see [7]) and its second-order cone (SOCP) relaxation (see [35]). The DistFlow model neglects voltage phase angle and is only valid for radial networks. In meshed networks, an additional non-convex constraint called the cycle constraint is needed to state the more general BranchFlow model [35]. Whereas the SOCP relaxation is provably exact for some cases of radial networks [38], and widely believed to typically be exact in practice even without theoretical guarantees, the same does not hold for meshed networks. Our approach instead uses a dynamic linearization based on current and voltage phasors in rectangular coordinates that is updated every iteration of the algorithm. This is not guaranteed to converge, but if it does, it will converge to a local optimum satisfying the non-linear AC optimal power flow problem. We developed this approach independently, but have since observed that it is very similar to an algorithm developed for transmission systems with quadratic cost generators in the 2015 paper [71]. We have not yet carried out numerical experiments to validate its convergence in practice in agent-based microgrids, but [71] observes excellent convergence

performance and recovery from infeasible iterations. Although this iterative linearization is highlighted here, we also give a solution to the DC power flow problem and show how any static power flow linearization can be used instead of the iterative linearization. Both of these are guaranteed to converge, but except in the trivial case, the solution will only solve the linearized power flow equations and not the non-linear ones.

The second differentiating characteristic of this work with respect to the P2P literature is to use a variation of consensus ADMM that produces a solution to AC optimal power flow at every iteration. This is made possible by stating the coupling variables in ADMM as voltage and current, which are related by linear constraints. We also show how this gives an analytic global update that is related to edge-Laplacian of the graph representing the power network, which establishes a theoretical connection to an area of graph theory studying “edge agreement” dynamics [128] and may prove useful in developing convergence bounds for ADMM specifically in power networks (e.g. [73] employs a graph-theoretic proof using the edge-Laplacian).

Work on ADMM for decentralized optimal power flow in transmission networks focuses on decomposing the problem by “areas” instead of nodes and is given a relatively detailed treatment in the comprehensive review of decentralized and distributed optimization in power networks [80]. These works predate the P2P literature by some 5-10 years. References [56, 34] are foundational and demonstrate the scalability and performance of the approach. The review [80] also highlights sensitivity of the convergence rate to the penalty parameter, and points to the importance of this design parameter, which we do not address here. As noted above, [71] uses a very similar approach to ours where the power flow constraints are given in terms of voltage and current phasors, and a successive convex approximation is used at every iteration. However, they do not derive analytic solutions to the global update step, leaving it instead as an optimization problem. Our work gives an analytic solution with the edge-Laplacian, which involves a pre-computed matrix inversion and has the potential to significantly reduce the computation time of each iteration of the algorithm, thus contributing to the broader decentralized optimal power flow literature.

Our work also contributes the literature on online algorithms solving optimal power flow in real-time using feedback control and the closely related study of economically optimal primary and secondary frequency control. The review [80] includes a survey of this active area of research, with notable works include [30, 4, 21, 37, 25] and related publications by their authors. The approach taken in these papers is to solve the optimization problems analytically using the Karush-Kuhn Tucker (KKT) conditions, which yields an equation for the optimal operating point. Then, the control system is reverse-engineered so that this optimal point is the unique stable equilibrium of the closed-loop system. Typically, this yields a control rule that, once parameterized, can be executed with local feedback only. One of the important areas for future work highlighted by the authors of [80] is that the analysis restricted to quadratic cost functions. In this paper, we use ADMM to develop an approach that is effective for arbitrary convex cost functions or concave utility functions.

4.2 Agent models, notation, and basic graph-theoretic concepts

In this section we introduce the agent models, the specific notation that we use to represent the agents and power network, and state some basic graph theoretic concepts that we use. Most of the notation is introduced throughout the following subsections, but we introduce a couple standard choices here that are used throughout.

The symbol I represents the identity matrix and $\mathbb{1}$ the vector of all ones in dimensions conforming to the context. The superscript x^T denotes the transpose of x . $\|x\|_p$ denotes the p -norm of x . We use subscripts to index variables that are elements of a set and superscripts to name them. The braces $\{\}$ denote a set either consisting explicitly of the elements inside, or implicitly over all variables if an index is used. That is, $\{x\}$ refers to a set consisting solely of x , but $\{x_n\}$ refers to the set of all x indexed by n , which should be clear from the context. The functions $\text{Re}(x)$ and $\text{Im}(x)$ denote the real and imaginary parts of a complex number x , and $\text{conj}(x)$ denotes the complex conjugate of x .

Agent models

In this section, we introduce models for agents associated with distributed energy resources (DERs) and electric loads, and their respective utility functions and constraints. We first define an abstract generic agent, and then show concrete instances for specific DERs and loads. The concrete agent models themselves are not the focus of this chapter. Rather, the purpose of defining the models here are 1) to make the agent concept more concrete for the reader, and 2) to illustrate how agents could be implemented in practice in a programming language with abstract classes and inheritance. The models themselves follow closely from those used in Chapters 2 and 3, with the exception of loads and an extension to both grid-connected and islanded microgrids. Here, a piecewise linear utility function is used for loads, whereas in Chapter 2 both quadratic approximations and a discrete choice / mixed integer model were used, and in Chapter 3 a strictly concave load utility function was used. This is done to relax the strictly concave assumption, but to avoid using non-convex integer models. These choices are discussed further below.

As in the previous chapter, the context is motivated by the scenario of 100% renewable systems composed of microgrids and interconnected microgrids on the kW to MW scale. We model a microgrid composed of three types of local resources – solar photovoltaics, batteries, and loads – plus a fourth type of resource that is an interconnection to the broader “macrogrid”. We associate an *agent* with each individual resource.¹ An individual resource

¹Note that this differs from the previous chapter, which associated an agent with a person or entity deriving value from the resource. There are several reasons for this change. One is to be able to state simpler models for each agent that only consider one resource type, each of which has distinct characteristics. Moreover, in many cases, the physical resource will be connected to a (communication) networked device, such as a power converter, smart meter, or controller. Assigning one agent per resource simplifies communication in these cases, as it necessitates only a one-to-one link, and even allows a potential implementation for the

can be an aggregate of resources of the same type that act as a single unit; i.e. an array of solar panels, a bank of batteries, or a group of loads.

As in the previous chapter we assume agents choose quantity to maximize their net value over a look-ahead time horizon, subject to the constraints of their resource. This choice is formulated as an optimization problem called the *agent's decision model* and the net value is the objective. We assume that the agents are rational, or more accurately that they are a computer program that solves the decision model deterministically. For load agents, the value is a function that models benefits derived from the activities and services enabled by electricity consumption. For interconnection agents, the value is a net price paid/received for power consumed/supplied to the macrogrid. This price is exogenous. Solar and battery agents do not have intrinsic value or costs in this model. Once trading is introduced, they have exchange value, and naturally the energy stored in batteries has a future exchange value, but this arises entirely from the value of energy to the loads and/or the macrogrid. Load, solar, and interconnection agents make or receive forecasts of demand, irradiance, and grid price, respectively, over the time horizons, and this information is used to parameterize their decision model. In general, these forecasts are assumed to be based on external information, such as weather forecasts or user input on their preferences, and potentially on past observations as well.

Also as in the previous chapter, the quantity is *average* net power over the time interval, but in this case it is the net power consumption instead of supplied, which is simply a sign change. The *average* power is equivalent to energy when scaled by the length of the time interval, so we use the terms interchangeably. This notably neglects the *instantaneous* peak power over the interval, which is very relevant for physical constraints on resources and power lines in the grid. An expanding time horizon that starts with shorter intervals (i.e. 15 minutes or less) closer to the present, mitigating this modeling error somewhat, but addressing instantaneous flows directly remains an area for further development.

Before stating the agent models mathematically, it is important to reflect on the perspective taken here, which is fundamentally utilitarian, treats electricity as a commodity with private value, and neglects shared value to the extent that it is not captured by the exchange of electricity for payments. There is a vast landscape of critical questions that need to be addressed around the relationships between agents and the entities they are assumed to act on behalf of, the overarching governance of the system itself, and its complex interactions with society. Foremost among these are the extent to which the agents are able to represent individual preferences, the equity of outcomes and distribution of energy access, and understanding how social norms around shared infrastructure both affect and are affected by the physical infrastructure itself. This work is not meant to argue that a utilitarian approach is ethical, moral, or leads to sustainable infrastructure that serves the purpose of society. Nevertheless, it does show how such an approach provides a mechanism

agent to run as a program on the microprocessor of the device itself. From a perspective of abstraction, this also allows additional types of resources and corresponding agents to be added to the framework, for example fuel cells, gas generators, wind turbines, etc., without changing the fundamental architecture.

for efficiently managing a very complex infrastructure system that is being asked to rapidly evolve to become more resilient, provide more access to the global population, incorporate more renewable and distributed renewable energy, and support the increasing electrification of end-use energy. We advocate that the critical infrastructure study should be carried out in conversation with the practical demands of the electricity grid.

General agent model

The generic decision model for agent n can be stated mathematically as

$$V_n(x_n, w_n, \phi_n) = - \min_{p, z_n} - U_n(z_n, x_n, w_n, \phi_n) \quad (4.1a)$$

$$\text{s.t. } (p, z_n) \in \mathcal{C}_n(x_n, w_n, \phi_n) \quad (4.1b)$$

V_n is a “value function” that reflects the maximum value an agent can derive, given a time-varying state variable x_n , an exogenous, time-varying forecast w_n , and time-invariant (static) parameters of the system ϕ_n . Examples of the state x_n are the battery state of charge for a battery agent, the cumulative energy consumed for certain types of loads, or the internal state of a controllable load such as a thermostat. Solar agents and interconnection agents do not have a state. Examples of exogenous forecasts w_n are the solar irradiance, consumption or value of loads, or the macrogrid electricity price. In this formulation, the forecasts are assumed to be deterministic for simplicity, but it can also be formulated stochastically, in which case the value function can be the expected value (see Chapter 2). Examples of ϕ_n are battery capacity (both in terms of stored energy and charge and discharge rates), battery efficiencies, solar panel capacity and orientation, etc. charge and discharge capacity the static parameters of the system. The subscripts n are used to signify that these variables are parameters specifically for the n 'th agent's value function; another agent m can have parameters drawn from a different space with difference dimensions.

The variables z_n and p are the agents' decision variables. These are defined so that p is the power *withdrawal* from system and z_n are the set of remaining decision variables, which may also include the power. The decision variables are separated in this way so that p can be isolated and analyzed as the variable that couples agents together through the microgrid (this analysis is the focus of the remaining sections of the chapter following the agent models). The variable z_n is given the subscript n to signify that it is drawn from a space particular to agent n , whereas p does not have a subscript because it is drawn from the same space for all agents. z_n and p are coupled by the constraint set \mathcal{C}_n , which represent the union of the physical resource constraints with any additional constraints necessary to define auxiliary decision variables in z_n . \mathcal{C}_n can depend on the state, forecasts, and static parameters; for example, the feasible charge or discharge power of a battery over the time horizon depends on its initial state of charge, or the maximum power that a solar array can generate depends on the irradiance. p should be understood as a vector of power over the time horizon. The other variables are more general sets, but their elements should be understood implicitly as vectors or scalar values from the context. We do not explicitly denote the temporal

dimension except where it is particularly relevant to prevent the notation from being overly complex.

U_n is the agent's utility function, which is their intrinsic value derived from power consumption, and can also depend on the state, forecast, and static parameters. It is important to note that U_n is written in a way that does not explicitly depend on p , but it still depends on p implicitly through the constraint $(p, z_n) \in \mathcal{C}_n$. For example, we can (and will in some cases) trivially define p_n as a variable in z_n , and include the constraint $p = p_n$ in \mathcal{C}_n so that utility function depends on power. p is the vector over the time horizon of average power consumption, U_n is the intrinsic utility of power consumption, and \mathcal{C}_n are the physical resource constraints. Both U_n and \mathcal{C}_n can depend on the state and the exogenous parameters, but do not necessarily in all cases. The two negative signs in (4.1a) are used to convert the maximization problem to an equivalent minimization problem, adopting the standard convention in convex optimization theory.

It is also critical to stress that U_n is the intrinsic value of power, absent a mechanism for exchange within the microgrid. Therefore, the battery and solar agents have 0 valued utility functions because they function only to generate and store the resource and do not have intrinsic value.² If fuel-based generators were included, they would have a (negative) utility function in the form of the fuel costs and associated costs for emissions.³ Only the load agents and the interconnection agents have non-zero valued utility functions, which arise from the value of electricity use and the cost/revenue of importing/exporting electricity from the macrogrid.

In Section 4.3, we will analyze the exchange of electricity with a modified decision model assuming that average power can be bought or sold for a price over a network. To define this, let π be the vector of prices over time and $\bar{P} > 0$ be the maximum power transfer between the agent and the network, which is generally the rating of their connection to the network. Then, the *agent's decision model with exchange* can be written as

$$V_n^\pi(\pi, \bar{P}, x_n, w_n, \phi_n) = - \min_{p, z_n} - U_n(z_n, x_n, w_n, \phi_n) + \pi^T p \quad (4.2a)$$

$$\text{s.t. } (p, z_n) \in \mathcal{C}_n(x_n, w_n, \phi_n) \quad (4.2b)$$

$$|p| \leq \bar{P} \quad (4.2c)$$

where V_n^π is the “net” value function given price π .

We will also use an *augmented decision model with exchange* that includes a penalty on the distance of the chosen power consumption from a power withdrawal reference point \hat{p} .

$$V_n^{\pi, \rho}(\hat{p}, \pi, \bar{P}, x_n, w_n, \phi_n) = - \min_{p, z_n} - U_n(z_n, x_n, w_n, \phi_n) + \pi^T p + \frac{\rho}{2} \|p - \hat{p}\|_2^2 \quad (4.3a)$$

²It is likely that people derive an intrinsic value from having solar or battery systems. This is effectively ignored here, but technically it can be included as constant value without changing the model statement. Even if we posit that generating solar power as an altruistic act having some value, there is a strong argument that the value is derived from avoided emissions or the electricity uses that it enables; it is hard to argue that generating solar power that is never used has intrinsic value.

³Batteries could also easily be modeled to have a degradation cost that is a function of their power throughput, but this is neglected for simplicity here.

$$\text{s.t. } (p, z_n) \in \mathcal{C}_n(x_n, w_n, \phi_n) \quad (4.3b)$$

$$|p| \leq \bar{P} \quad (4.3c)$$

The relevant quantity here for bidding systems will be the optimal power, not the value itself. We denote this decision as a function \mathcal{P}_n :

$$\mathcal{P}_n(\hat{p}, \pi, \rho, \bar{P}, x_n, w_n, \phi_n) = \arg \min_{p|z_n} -U_n(z_n, x_n, w_n, \phi_n) + \pi^\top p + \frac{\rho}{2} \|p - \hat{p}\|_2^2 \quad (4.4a)$$

$$\begin{aligned} \text{s.t. } (p, z_n) &\in \mathcal{C}_n(x_n, w_n, \phi_n) \\ |p| &\leq \bar{P} \end{aligned} \quad (4.4b)$$

where the notation $p|z_n$ means that the optimal value of p is returned by arg min but that z_n is also a decision variable. Note that if U_n is convex (but not necessarily strictly convex) in z_n and $\mathcal{C}_n(x_n, w_n, \phi_n)$ is a convex set, then (4.4) has a unique solution and \mathcal{P} is indeed a function.

The agent model is stated in terms of real power p only. It could be augmented to include reactive power, which would in practice typically appear in the constraints on power injection, which are generally more correctly stated in terms of apparent power but this is not done for simplicity.

The specifics of this abstract model for each agent type are given in the following subsections. To simplify notation, the subscript n is dropped for the remainder of the section; all variables and functions are implicitly associated to an agent n .

Solar agent model

The solar agent's model is relatively simple. It is stateless and has zero intrinsic value. The exogenous parameter is a forecast of maximum power point solar production, \bar{P}^s , and the constraint set is simply that the power produced (i.e. the negative of that consumed) must be less than this value at each time. This is summarized as

$$x := \{\} \quad (4.5a)$$

$$w := \{\bar{P}^s\} \quad (4.5b)$$

$$\phi := \{\} \quad (4.5c)$$

$$z := \{\} \quad (4.5d)$$

$$U(z, x, w, \phi) \equiv 0 \quad (4.5e)$$

$$\mathcal{C}(x, w, \phi) := \{(p, z) : -\bar{P}^s \leq p \leq 0\} \quad (4.5f)$$

Note that z is superfluous in this case because the utility function is constant (specifically equal to zero) and the private constraints are simple.

Battery agent model

Like the solar agent, the battery agent has zero intrinsic value. However, it does have a state, which is its stored energy, and it has a more complex constraint set that includes intertemporal constraints. In order to capture charging and discharging efficiencies to the first order, the model from Section 3.A is used, neglecting self-discharge and changing the sign convention of the notation. This model introduces auxiliary decision variables e , p^+ and p^- to denote the stored energy, charge power, and discharge power, which are all vectors over time. The charge and discharge efficiencies $\sigma^+ \in (0, 1]$ and $\sigma^- \in (0, 1]$ are constant and fixed. The maximum charge and discharge powers, \bar{P} and \underline{P} , and the maximum and minimum stored energy \bar{E} and \underline{E} can either be time varying or time variant; however, if they are time varying, or if it is possible for initial stored energy, denoted e_0 , to be outside of $[\underline{E}, \bar{E}]$, then special attention must be paid to potential infeasibility cases that are not addressed in this chapter. Let ΔT_t denote the length of the time step at time t . This allows for the time step to be non-uniform; setting the time periods be longer when they are farther into the future is a useful technique for capturing longer time horizons. The timing information is assumed to be global and not part of any of the agents' parameters.

The decision model for a battery agent is defined by

$$x := \{e_o\} \quad (4.6a)$$

$$w := \{\} \quad (4.6b)$$

$$\phi := \{\underline{E}, \bar{E}, \underline{P}, \bar{P}, \sigma^-, \sigma^+\} \quad (4.6c)$$

$$z = \{p^+, p^-, e\} \quad (4.6d)$$

$$U(p, x, w, \phi) \equiv 0 \quad (4.6e)$$

$$\mathcal{C}(x, w, \phi) := \{(p, z) : p = p^+ - p^-, z \in \mathcal{C}^z\} \quad (4.6f)$$

with the private constraints \mathcal{C}^z defined as the intersection of the following convex constraints:

$$\mathcal{C}^z := \mathcal{C}^c(e_0) \cap \mathcal{C}^e(\underline{E}, \bar{E}) \cap \mathcal{C}^{p^-}(\underline{P}) \cap \mathcal{C}^{p^+}(\bar{P}) \quad (4.7a)$$

$$\mathcal{C}^c(e_0) := \{z : e_t = e_{t-1} + (\sigma^+ p_t^+ - \frac{1}{\sigma^-} p_t^-) \Delta T_t = 0 \quad \forall t\} \quad (4.7b)$$

$$\mathcal{C}^e(\underline{E}, \bar{E}) := \{z : \underline{E} \leq e \leq \bar{E}\} \quad (4.7c)$$

$$\mathcal{C}^{p^-}(\underline{P}) := \{z : 0 \leq p^- \leq \underline{P}\} \quad (4.7d)$$

$$\mathcal{C}^{p^+}(\bar{P}) := \{z : 0 \leq p^+ \leq \bar{P}\} \quad (4.7e)$$

There is an additional non-convex technical constraint, that $p^{+\top} p^- = 0$, enforcing that a battery cannot simultaneously charge and discharge. Following the analysis in Section 3.A, this constraint can be relaxed and then recovered when prices are positive, so it is relaxed here to maintain convexity. As shown in the previous chapter, negative prices should not arise in equilibrium as solar and battery generation can be freely curtailed; however negative

prices might arise during bidding iterations, and thus this condition should be checked in practice, and additional modifications may be necessary.

Load agent model

The load agent model plays the important role of defining the utility functions that define the value of energy and drive the optimal allocation. It also attempts to model human preferences, which are notoriously difficult, even impossible by many philosophies, to quantify. The approach taken here is to use a relatively simple model that captures some salient features of electricity use that can be represented mathematically to varying degrees of complexity. This section first compares two distinct approaches to load modeling in an optimization framework, then states a flexible model that can be used to approximate utility functions.

The most common approach in load modelling is to model different types of loads by having particular constraints, and potentially also having some marginal utility. This typically includes some must-run or inflexible loads that do not respond to price and some loads that have inherent energy storage that gives them flexibility as to when they consume as long as some total consumption is met. At the household level, examples of the former are lighting, cooking, or watching television; examples of the latter are an electric vehicle that must charge completely before it is needed in the morning, or a water heater or refrigerator that must keep its temperature in some allowable range. This approach is logical from the perspective of the legacy grid that meets inflexible demand, essentially leveraging some flexibility in the loads to reduce supply-side costs, but it has the potential to fall short in 100% renewable systems where electricity supply faces variable constraints. Particularly, during times of extreme shortage, modeling loads with hard constraints can lead to infeasible scenarios if there is simply not enough energy to satisfy the constraints, and leads to the question of which critical loads to satisfy under extreme scarcity. This could lead to a convergence failure, or cause users' agents to bid for astronomically high energy prices to satisfy the hard constraints.

Rather than taking this approach of hard constraints, all loads are considered to be flexible and to have a particular value. Critical loads should be given a very high value commensurate with a maximum willingness to pay that reflects their high priority.

The total load for an agent is composed of multiple individual loads. The l 'th load has a marginal value of energy, or maximum price that the user is willing to pay, denoted ν_l , and the load agent receives (or makes) a forecast of the load \bar{P}_l . In the simplest case from the user's perspective, their entire load is represented as a single load, and they are required to enter one value ν_1 reflecting their willingness to pay for electricity. This could be recommended to the user based on their expected consumption and monthly budget for electricity consumption, and adjusted up or down by them.⁴ Advanced users could optionally

⁴For a residential user, a recommendation system could estimate their monthly consumption based on households with similar characteristics, calculate by simulation a typical monthly cost and associated reliability for high, medium, and low willingness to pay values, and have the user select an option.

define different willingness-to-pay for different load circuits, or differentiate their willingness-to-pay by time of day.

$$x := \{\} \quad (4.8a)$$

$$w := \{\bar{P}_l\} \quad (4.8b)$$

$$\phi := \{\nu_l\} \quad (4.8c)$$

$$z := \{p_l\} \quad (4.8d)$$

$$U(z, x, w, \phi) := \sum_l \nu_l^T p_l \quad (4.8e)$$

$$\mathcal{C}(x, w, \phi) := \{(p, \{p_l\}) : p = \sum_l p_l, 0 \leq p_l \leq \bar{P}_l, \forall l\} \quad (4.8f)$$

This model is limited, and most notably does not include shiftable loads or loads with energy storage. There is ample research modelling these loads that can be drawn on to expand this model; the main issue that arises is that many of these models either require integer variables or use convex formulations that are not necessarily feasible, which introduces additional complexity.

Interconnection agent model

Lastly, we give a simple interconnection agent that can represent the point of common coupling of a microgrid with an external grid. This allows the methods used here to apply to both islanded and grid connected systems. π^g denotes the forecasted grid price for electricity, and $\underline{P} \leq 0$ and $\bar{P} \geq 0$ denote constraints on the power transfer. The private variable p^g is defined to represent the power exported to the external grid. Using a similar technique for splitting the battery into charge and discharge components, this model could be extended to have asymmetric import and export prices. The interconnection agent model is thus:

$$x := \{\} \quad (4.9a)$$

$$w := \{\pi^g\} \quad (4.9b)$$

$$\phi := \{\underline{P}, \bar{P}\} \quad (4.9c)$$

$$z := \{p^g\} \quad (4.9d)$$

$$U(z, x, w, \phi) := \pi^{gT} p^g \quad (4.9e)$$

$$\mathcal{C}(x, w, \phi) := \{(p, z) : p = p^g, \underline{P} \leq p^g \leq \bar{P}\} \quad (4.9f)$$

Power network model and graph theoretic concepts

This section introduces graph theoretic notation and some non-restrictive assumptions to impose a structure that is useful for describing and analyzing the electrical network.

Let the network be represented as a graph. A graph consists of a set of nodes \mathcal{N} and a set of edges \mathcal{L} . We define the number of nodes and lines in the graph to be $N := |\mathcal{N}|$ and $L := |\mathcal{L}|$, respectively. Electrically, the nodes correspond to “buses” and the edges correspond to “lines”, i.e. power lines, or conductors. We will use the terms interchangeably. Some texts also use the term “vertex” in place of a node. Each line $l \in \mathcal{L}$ connects an unordered pair of nodes $(i, j) \in \mathcal{N}$. We will denote this $l : i \leftrightarrow j$.

The above describes an *undirected* representation of the network. We will also use a directed representation of the same network, where each l is given a unique arbitrary direction. For a line l , the direction is defined by *ordering* the pair of nodes (i, j) . This directed line is denoted $l : i \rightarrow j$. In this ordering, the first node i is called the “sending” node, and node j is called the “receiving” node. We use \rightarrow and \leftrightarrow to distinguish whether we are referring to the directed or undirected representation of the line connecting i and j . Both directed and undirected representations of the graph will prove useful mathematically for representing the network constraints. We use the term “the directed graph” and “the undirected graph” to refer to the directed and undirected representations of the same network.

We assume that the undirected graph is connected, which means there is a path of undirected lines connecting every node in the graph to every other node. This is equivalent to the directed graph being weakly connected. This is a trivial assumption: if it did not hold, the graph could be represented as two or more distinct networks, i.e. distinct grids, that are all completely isolated from each other and could be analyzed separately. We also assume that each agent n is associated with a physical resource (DER) that is located at a unique node i in the network. This is also trivial to impose from a physical perspective: it is simply that each DER cannot be in two places at once. We assume for simplicity that the network is single phase, or is the single phase equivalent of a balanced three phase network. Extensions to unbalanced three phase networks is an area for future work, where intuitively the essence of the main results should still hold, but with different equations accounting for the phase imbalance.

Let \mathcal{N}_i denote the open undirected neighborhood of a node i , i.e. all the nodes connected by a single undirected line to i . Technically speaking, this is different from the closed neighborhood, which also includes node i itself, but whenever only the term “neighborhood” is used here, it is meant to refer to the open neighborhood. Let \mathcal{L}_i be called the incident lines of i , defined as the set of all undirected lines that connect i to another node. Let \mathcal{N}_i^s be called the sending neighborhood of i , which is defined as all the nodes j in the neighborhood of i for which i is the sending node of the directed line connecting them. Let the associated lines, that is all lines such that i is a sending node of l be called the sending lines of i , denoted \mathcal{L}_i^s . Let the receiving neighborhood \mathcal{N}_i^r and the receiving lines \mathcal{L}_i^r of node i be defined analogously. Let N_i, N_i^s, N_i^r and L_i, L_i^s, L_i^r denote the cardinality of these sets, that is, the number of nodes or lines in each. It is clear that these subsets form a partition of the neighborhood and incident lines of i , that is $\mathcal{N}_i^s \cup \mathcal{N}_i^r = \mathcal{N}_i$ and $\mathcal{N}_i^s \cap \mathcal{N}_i^r = \emptyset$, and $\mathcal{L}_i^s \cup \mathcal{L}_i^r = \mathcal{L}_i$ and $\mathcal{L}_i^s \cap \mathcal{L}_i^r = \emptyset$.

These neighborhoods and incident line sets have associated matrices called adjacency and incidence matrices. Let J^s and J^r be matrices in $\{0, 1\}^{N \times N}$ that are defined as $J_{ij}^s = 1$ iff

there is a directed line connecting $i \rightarrow j$ and $J_{ij}^r = 1$ iff there is a directed line connecting $j \rightarrow i$. Thus the *undirected* adjacency matrix is $J^u = J^s + J^r$ and the *directed* adjacency matrix is $J^d = J^s - J^r$. Similarly, let H^s and H^r be matrices in $\{0, 1\}^{N \times L}$ with $H_{il}^s = 1$ iff i is the sending node of line l and $H_{jl}^r = 1$ iff j is the receiving node of line l . Thus $H^u = H^s + H^r$ and $H^d = H^s - H^r$ are the undirected and directed incidence matrices, respectively. Let $J_i^{(0)\text{T}}$ and $H_i^{(0)\text{T}}$ denote the i 'th row of the respective $J^{(0)}$ and $H^{(0)}$ matrix with $(\)$ being the appropriate superscript $s, r, u,$ or d . It can be seen by definition that the neighborhood, sending neighborhood, and receiving neighborhood of i are equal to the set of indices of the nonzero elements of J_i^u, J_i^s and J_i^r respectively. Likewise the incident, sending, and receiving lines of i are equal to the indices of the nonzero elements of H_i^u, H_i^s and H_i^r .

Let the neighborhood of each node be ordered arbitrarily and let $\mathcal{N}_i(m)$ denote the m 'th neighbor of i . Let e_j denote the j 'th standard basis vector in \mathbb{R}^N and let E_i be a matrix in $\{0, 1\}^{N \times N_i}$ with $e_{\mathcal{N}_i(m)}$ as the m 'th column. Thus $E_i^{\text{T}} J^{(0)}$ selects columns of the respective adjacency matrix corresponding to nodes that are in the neighborhood of i . Likewise $E_i^{\text{T}} H^{(0)}$ selects the columns of the respective incidence matrix corresponding to lines incident to i .

The edge Laplacian of the directed graph will be relevant in the AC power flow problem. For a reference on the edge Laplacian and its properties, see [128]. The edge Laplacian can be defined in terms of the incidence matrix as $H^{d\text{T}} H^d$ and is an $L \times L$ matrix. It is closely related to the well-studied nodal Laplacian which is an $N \times N$ matrix equal to $H^d H^{d\text{T}}$. The nodal Laplacian can also be stated as the degree matrix minus J^u , where the degree matrix is a diagonal matrix with N_i in the i 'th diagonal entry. This is stated by the identity $H^d H^{d\text{T}} = \text{diag}(J^u \mathbf{1}) - J^u$.

Some important properties of the edge and nodal Laplacians can be shown from well-known results in linear algebra. All real symmetric matrices in the form of XX^{T} or $X^{\text{T}}X$ are positive semidefinite, meaning that all of its eigenvalues are greater than or equal to zero, and that the nonzero eigenvalues of XX^{T} and $X^{\text{T}}X$ are identical. This can be derived directly from the singular value decomposition of X . This implies a well-known fact in graph theory that the nodal and edge Laplacians share non-zero eigenvalues, and that all these eigenvalues are non-negative. It is also clear from the definition of the nodal Laplacian in terms of the degree and undirected adjacency matrices that the nodal Laplacian times the unit vector is always equal to the zero vector for any graph, thus zero is always an eigenvalue of the nodal Laplacian. If the graph is connected, which we assume, then the zero eigenvalue is not repeated. Thus, for any connected graph, zero is an eigenvalue with multiplicity $L - N + 1$ in the edge Laplacian. $L - N + 1$ is also equal to the number of cycles in the undirected network, and thus each cycle is associated with a zero eigenvalue in the edge Laplacian. The smallest non-zero eigenvalue in the edge and nodal Laplacian is called the algebraic connectivity, which decreases with the number of vertices for random graphs.

Let the electrical impedance of each line $l : i \leftrightarrow j$ be denoted $Z_l := Z_{ij} = Z_{ji}$, which is represented by a phasor described by a complex number. The corresponding admittance phasor is $Y_l := \frac{1}{Z_l} := Y_{ij} = Y_{ji}$. Let Z be a diagonal $L \times L$ matrix with the impedance of line l on the l 'th diagonal. The nodal admittance matrix Y is a weighted nodal laplacian, given

by $Y = H^d Z^{-1} H^{dT}$. This is not to be confused with the diagonal matrix of line admittances Z^{-1} . Let $B := \text{Im}(Y)$ be the susceptance matrix used in DC power flow models.

For the admittance and impedances to be properly defined as multiplicative inverses of each other, they must be nonzero. Let $|Y|^\infty \gg 0$ and $|Z|^\infty \gg 0$ be real numbers denoting admittances and impedances with very large magnitudes and arbitrary phase, and $0 < |Y|^0 \ll 1$ and $0 < |Z|^0 \ll 1$ be their very small counterparts so that $|Y|^\infty = \frac{1}{|Z|^0}$. This will be used as a placeholder value to represent lines with approximately 0 impedance.

Note that the graphical representation of a network is not unique. For example, a line could be “cut in half” and replaced by two lines in series, each with half of the impedance, and represent the same electrical network. We will say that a graph is an electrically equivalent representation of another if they both describe Kirchoff’s Voltage Law and Kirchoff’s Current Law holds at the nodes of interest of the same electrical network in the limit of $|Y|^\infty \rightarrow \infty$ and $|Z|^\infty \rightarrow \infty$.

We do not assume that the network is radial, and allow it to be meshed. This means the associated graph is not necessarily a tree, and that it may contain cycles.

The next set of assumptions are not trivial in the way that the assumption of the graph being connected is, but they can be made without loss of generality. Without loss of generality means here that if they are not satisfied for a given graph, there is an electrically equivalent graph that does satisfy the assumption. Thus, throughout the chapter, everything is written in terms of the electrically equivalent graph satisfying the assumptions. These assumptions and the procedures to construct electrically equivalent graphs are described in the next two paragraphs.

First, assume that there are no parallel lines; i.e. there is at most one $l : i \leftrightarrow j$. It follows that if the direction of l is defined as and if $l : i \rightarrow j$, then there is no $l : j \rightarrow i$. If it is not true, the equivalent graph satisfying these assumption is easily produced by adding the admittances of all parallel lines, and replacing these parallel lines with a single line with its admittance being this sum. This is just the basic procedure of reducing parallel lines taught in introductory circuit classes.

Next, assume that at most one resource is located at a node i , and that if a node i has a resource located it, then it has exactly one neighbor j . The equivalent graph satisfying these two conditions can be obtained by introducing additional nodes and lines. Specifically, suppose there are more than one DERs located at a node j . For each m of these nodes, add an additional node i to the graph and a line $l : i \rightarrow j$ connecting i to j with admittance $Y_l := |Y|^\infty$, and locate m at the new node i . Thus, each of these new nodes i has no more than one resource, and it has exactly one neighbor j . In the case that a single resource is located at node j , but that j has more than one neighbor, the same procedure of adding a node i and moving the resource to node i results in node i having exactly one neighbor j . Applying the procedure across all such nodes produces an electrically equivalent graph satisfying the assumption. The introduction of lines with very large admittance $|Y|^\infty$ is actually a quite faithful physical representation of the series elements connecting the DERs to the network, because in practice, there will be some conductor with very small impedance

that makes the connection. The main issue to be aware of is that numerically, including a very large number can result in ill-conditioned numerical operations; thus care needs to be used in implementation to make sure that it is appropriately divided out before executing numerical operations. This will be addressed in the derivations that follow.

This assumption allows us to partition \mathcal{N} into the set of nodes that have exactly one resource connected to them and those that have no resources connected to them. We call these the agent nodes, denoted \mathcal{N}^A , and the internal nodes, \mathcal{N}^I , respectively. These sets will be useful in defining the ADMM update equations.

We also assume that there all loads or resources have an agent associated with them. This implies that there is no withdrawal or injection into the network at the internal nodes \mathcal{N}^I . This assumption is more or less without loss of generality. Relaxing it simply means that the homogeneous linear flow conservation constraints at these nodes would be replaced with affine constraints. In the section on online control we will relax this and have uncontrollable and controllable nodes.

4.3 Economic power dispatch with ADMM

In this section we first state the general ADMM algorithm and the specific consensus-type ADMM algorithm commonly used for multi-agent systems. We next show how ADMM can be applied to solve the economic power dispatch problem with two network-aware models of the power flow constraints. The first is a lossless line model neglecting voltage drop that is equivalent to the commonly used DC power flow model. With this model, we show that the algorithm is guaranteed to converge to the optimum and requires only synchronized neighbor-to-neighbor communication. The second is the general non-linear and non-convex AC power flow model. In this case we state the solution in non-convex form, and show how an iterative linearization can be used efficiently to solve the non-convex problem by solving convex subproblems at each iteration. In comparison to the lossless line model neglecting voltage drop, the iterations are not carried out with only neighbor-to-neighbor communication, but require information from all nodes mapped through an inverse of the diagonally-modified edge Laplacian of the network. Finally, we conclude with a high-level discussion of the practical implementation of a system based on ADMM, including potential market and transaction structures.

General ADMM and consensus ADMM for multi-agent systems

Optimization problems of the form for ADMM can be written as:

$$\min_{u,v} f(u) + g(v) \tag{4.10a}$$

$$s.t. \quad Au + Bv = c \tag{4.10b}$$

$$u \in \mathcal{C}^u \tag{4.10c}$$

$$v \in \mathcal{C}^v \quad (4.10d)$$

We will call u the local variables and v the global variables for reasons that will be apparent when we state the consensus ADMM. A and B are matrices defining an affine constraint with constant c , and f and g are the cost functions of the local and global variables.

The algorithm uses the augmented Lagrangian of the problem with scalar parameter $\rho > 0$ and a dual variable y associated with the constraint $Au + Bv = c$:

$$\mathcal{L}_\rho(u, v, y) := f(u) + g(v) + y^T(Au + Bv - c) + \frac{\rho}{2} \|Au + Bv - c\|_2^2 \quad (4.11)$$

The algorithm can be stated concisely in three lines, with k representing the iteration count:

$$u^{k+1} = \arg \min_u \mathcal{L}_\rho(u, v^k, y^k) : u \in \mathcal{C}^u \quad (4.12a)$$

$$v^{k+1} = \arg \min_v \mathcal{L}_\rho(u^{k+1}, v, y^k) : v \in \mathcal{C}^v \quad (4.12b)$$

$$y^{k+1} = y^k + \rho(Au^{k+1} + Bv^{k+1} - c) \quad (4.12c)$$

Essentially, it performs alternating updates of the local and global variables u and v , and then updates the dual variable y , effectively integrating the constraint violation. In this statement, it is assumed that the local and global updates are feasible. The state of the iteration is stored jointly in v^k and y^k . Initial values of $y^0 = 0$ and $v^0 = 0$ can be used, or the iteration can be “warm-started” at some other initial point.

It can be shown that if f and g are closed and convex functions, if \mathcal{C}^u and \mathcal{C}^v are closed and convex sets, and with the additional technical condition that the unaugmented Lagrangian has a saddle point, then as $k \rightarrow \infty$, $Au^k + Bv^k \rightarrow c$ and $f(u^k) + g(v^k)$ converges to the optimal cost [14]. With additional technical assumptions, then u^k and v^k also converge to points that solve (4.10) [33]. Infeasibility can be detected using an approach described in [6]. The convergence analysis can be proven from first principles [14], or more generally because the algorithm is a specific case of proximal methods [33, 89].

If u , v , and c are vectors of complex numbers, f and g must be real valued functions with complex arguments. The iteration is equivalent if the transpose operator is generalized to the conjugate transpose, the term $y^T(Au + Bv - c)$ in (4.11) is replaced with $\text{Re}(y^T(Au + Bv - c))$, and the norm is appropriately defined using the conjugate transpose. In this case y is also a complex number. This follows directly from stating the problem in terms of the real and imaginary components of u and v , and assigning the real part of y to the constraint on the real part of $Au + Bv - c = 0$, and likewise for the imaginary part.

ADMM is especially powerful for multi-agent systems in the case where the the cost function to be optimized is a sum of individual cost functions, and the agents’ constraints can be separated into individual private constraints and coupling constraints that connect the agents. The individual cost functions can be private as well. In this case, one can generally define the global variable v to correspond to the variables that are involved in the

coupling constraints and the set of remaining private decision variables as the local variable u . With this definition, the local update can be executed in parallel for each agent with respect to the price y and the global variable v from the previous iteration.

To define the consensus ADMM for multi-agent systems, let n indicate the n 'th of N agents, u_n be their local decision variables, and f_n be their local cost function. The total cost f is the sum of all f_n . Partition the local variables u_n into private variables, denoted z_n , and “local targets” of the variables involved in the coupling, denoted v'_n . Let each agents' private constraint be denoted \mathcal{C}_n^u , which depends on both z_n and v'_n . Define the “global reference” of the local targets so they correspond one-to-one with v'_n , and denote these v_n . Let all the v_n be stacked in a vector v . The coupling constraints that connect the agents can now be defined as only a constraint on v , which is denoted \mathcal{C}^v . The linear constraint relating the local to global variables is now simply that $v_n = v'_n$. The objective of the global variables is zero, $g(z) \equiv 0$. The optimization problem (4.10) can be stated as:

$$\min_{\{v'_n, z_n, v_n\}} \sum_n f_n(v'_n, z_n) \quad (4.13a)$$

$$s.t. \ v'_n - v_n = 0 \quad \forall n \quad (4.13b)$$

$$(v'_n, z_n) \in \mathcal{C}_n^u \quad \forall n \quad (4.13c)$$

$$v \in \mathcal{C}^v \quad (4.13d)$$

In this special form of ADMM, the iteration drives the global copies and the local variables to consensus, i.e., satisfying $v_n - v'_n = 0$, while enforcing the coupling constraint in the global update at each iteration. Essentially, in the local update, each agent acts like its target variables v'_n are free, but the “price” y_n and the 2-norm penalty term drive their decisions to an equilibrium with $v_n^k \rightarrow v_n^k$ as $k \rightarrow \infty$. To generalize the model a little bit, note that we can arbitrarily weight each of the constraints $v_n - v'_n = 0$ by restating it as $W_n(v_n - v'_n)$, where W_n is any diagonal matrix with nonzero weights on the diagonal, because W_n is non-singular and therefore $W_n(v_n - v'_n) = 0$ if and only if $v_n - v'_n = 0$. The ADMM iteration equations for the weighted consensus algorithm are written as:

$$v_n'^{k+1} = \arg \min_{v'_n | z_n} f_n(v'_n, z_n) + y_n^{kT} v_n + \frac{\rho}{2} \|W_n(v'_n - v_n^k)\|_2^2 \quad \forall n \quad (4.14a)$$

$$s.t. \ (v'_n, z_n) \in \mathcal{C}_n^u$$

$$v^{k+1} = \arg \min_v \sum_n -y_n^{kT} v_n + \frac{\rho}{2} \|W_n(v_n^{k+1} - v_n)\|_2^2 \quad (4.14b)$$

$$s.t. \ v \in \mathcal{C}^v$$

$$y_n^{k+1} = y_n^k + \rho W_n(v_n'^{k+1} - v_n^{k+1}) \quad (4.14c)$$

The key feature here is that the local update for u can be updated in parallel and that the private constraints and cost functions need not be shared with any other agents or with a

central operator, thus preserving privacy. From a complexity perspective, the iteration also decomposes the problem into simpler subproblems that can be easier to solve or analyze. There is also some discretion as to whether the coupling constraints are placed in the global constraint set \mathcal{C}^v or in the local constraint set \mathcal{C}_n^u . For example, assume that a coupling constraint depends only on a variable that is shared by n and m . Naturally, this constraint could be in the global constraint set \mathcal{C}^v . However, one can also define a duplicate copy of the variable, and distribute one copy to agent n and another to agent m , distribute the constraint to one or both of them, and add the constraint that the two copies must be equal to \mathcal{C}^v . This yields an equivalent solution. This can be useful to do in a context where the global update needs to be relatively simple, or if the global constraint set can be decomposed into the intersection of constraints involving only a few agents each, which will be the case in the power flow problem.

If we further assume that the global coupling constraint \mathcal{C}^v is a linear equality constraint, i.e., that it can be written as $Gv = 0$ for some matrix G , then the global update can be solved analytically and takes the form of a linear mapping. In fact, this assumption can always be met by using the approach of distributing multiple copies of the same variable to each agent that is described in the previous paragraph. To obtain the solution in this case, let μ^{k+1} be the dual vector associated with the constraint $Gv = 0$, and let $G := [G_1, \dots, G_n]$ so that $Gv = \sum_n G_n v_n$. Then the Karush-Kuhn-Tucker conditions for (4.14) are:

$$v_n^{k+1} = v_n'^{k+1} + \rho^{-1} W_n^{-2} (y_n^k - G_n^T \mu^{k+1}) \quad \forall n \quad (4.15a)$$

$$\sum_n G_n v_n^{k+1} = 0 \quad (4.15b)$$

By multiplying both sides of (4.15a) by G_n , summing over n , and then applying (4.15b) to eliminate the v_n^{k+1} variables, we obtain a solution for μ^{k+1} in terms of intermediate variables \tilde{G} , ξ^k , and η^k :

$$\tilde{G} := \sum_n G_n W_n^{-2} G_n^T \quad (4.16a)$$

$$\xi^k := \sum_n G_n W_n^{-2} v_n'^{k+1} \quad (4.16b)$$

$$\eta^k := \tilde{G}^\dagger \xi^k \quad (4.16c)$$

$$\mu^{k+1} = \rho \eta^k + \tilde{G}^\dagger \sum_n G_n y_n^k \quad (4.16d)$$

where $()^\dagger$ denotes the pseudo-inverse of a matrix in the case that \tilde{G} is not invertible.⁵ The additional vectors ξ^k and η^k are introduced to illustrate the components of the update. ξ is defined as the “error” vector for the local targets v_n' with respect to the global constraints

⁵In the application to optimal power flow in this chapter, the matrix will be invertible, and the solution for μ is unique, but the more general statement is given here.

G and weighted by W^2 , and η is defined by applying \tilde{G}^\dagger to the error, which can be thought of as a vector that projects v' onto the feasible space of the coupling constraints. Applying this to (4.14), we get equations for the weighted consensus iteration for the case when the global update is linear:

$$v_n'^{k+1} = \arg \min_{v_n'|z_n} f_n(v_n', z_n) + y_n^{k\text{T}} v_n' + \frac{\rho}{2} \|W(v_n' - v_n^k)\|_2^2 \quad (4.17a)$$

$$\text{s.t. } (v_n', z_n) \in \mathcal{C}_n^u$$

$$\xi^k := \sum_n G_n W_n^{-2} v_n'^{k+1} \quad (4.17b)$$

$$\eta^k := \tilde{G}^\dagger \xi^k \quad (4.17c)$$

$$\mu^{k+1} = \rho \eta^k + \tilde{G}^\dagger \sum_n G_n y_n^k \quad (4.17d)$$

$$v_n^{k+1} = v_n'^{k+1} - G_n^\text{T} \eta^k + \rho^{-1} W_n^{-2} (y_n^k - G_n^\text{T} \sum_m G_m y_m^k) \quad (4.17e)$$

$$y_n^{k+1} = y_n^k + \rho W_n (v_n'^{k+1} - v_n^{k+1}) \quad (4.17f)$$

It is often the case that the iteration is not weighted and W is the identity matrix I . In this case, (4.15a) implies that $y^k + \rho(v_n'^{k+1} - v_n^{k+1}) = G^\text{T} \mu^{k+1}$, which in turn implies $y^k = G_n^\text{T} \mu^{k+1}$ and allows (4.17f) and (4.17d) to be simplified further. If we also assume that \tilde{G} is full column rank, then $\tilde{G}^\dagger \tilde{G} = I$ and (4.17e) can also be simplified. Thus the unweighted iteration with $W = I$ can be stated as:

$$v_n'^{k+1} = \arg \min_{v_n'|z_n} f_n(v_n', z_n) + y_n^{k\text{T}} v_n' + \frac{\rho}{2} \|v_n' - v_n^k\|_2^2 \quad (4.18a)$$

$$\text{s.t. } (v_n', z_n) \in \mathcal{C}_n^u$$

$$\xi^k := \sum_n G_n v_n'^{k+1} \quad (4.18b)$$

$$\eta^k := \tilde{G}^\dagger \xi^k \quad (4.18c)$$

$$v_n^{k+1} = v_n'^{k+1} - G_n^\text{T} \eta^k \quad (4.18d)$$

$$\mu^{k+1} = \mu^k + \rho \eta^k \quad (4.18e)$$

$$y_n^{k+1} = G_n^\text{T} \mu^{k+1} \quad (4.18f)$$

Note that in this case the algorithm state is described entirely by μ , which can be seen as a kind of integration of the effective error term. The remaining variables ξ , η , v_n , and y can all be defined in terms of μ , and thus are not strictly necessary to define the algorithm. However, ξ , η , v_n , and y , have useful interpretations as the constraint error, the effective constraint error, the projection of the local target variables onto the global constraint set

(and also the reference point for each local update), and a price signal, respectively, so they are defined explicitly. In the case of the weighted algorithm with $W \neq I$, then the state is most compactly described by the price y instead of μ .

Note also that the updated value of each global coupling variable v_n^{k+1} only depends on η^k , does not depend on any of the price variables, and can be computed in parallel for each n once η is obtained. The constraint violation ξ^k can also be computed in parallel as a running sum. The only step that is not directly parallelizable for large scale systems is computing the effective error η^k .

Application of ADMM to optimal energy allocation

Case 1: Lossless lines ignoring voltage drop

Recall the general utility-based model for the generic agent decision from Section 4.2. To simplify notation, the utility functions U_n and the constraints \mathcal{C}_n are assumed to depend implicitly on x_n, w_n, ϕ_n from here on out, and are only written in terms of (p_n, z_n) .

We drop the temporal dimension from the notation to avoid clutter. All equations and inequalities can be interpreted as over all time steps. We can safely do this because the only intertemporal coupling that arises is within the private constraints of agents, specifically the batteries. Because of the intertemporal coupling, the equilibrium values in different time periods will depend on each other, but all of the algorithm steps except for the agent's local update are separable in time. Thus the algorithm can be analyzed for a single time period, keeping in mind that the intent is for it to be run over multiple time periods in a horizon simultaneously and synchronously.

Let p_{ij} denote the power flow from node i to j . By the assumption of lossless lines, $p_{ij} = -p_{ji}$. Assume that agent n is at node i . Let \bar{P}_l be the maximum power transfer on line l .

The agent-based, welfare-maximizing power dispatch problem with lossless lines ignoring voltage drop can be stated as:

$$p^* = \arg \min_{\{p_n\}} - \sum_n U_n(z_n) \quad (4.19a)$$

$$\text{s.t. } (p_n, z_n) \in \mathcal{C}_n, \quad \forall n \quad (4.19b)$$

$$p_{ij} + p_n = 0 \quad \forall i \in \mathcal{N}^A \quad (4.19c)$$

$$\sum_{j \in \mathcal{N}_i} p_{ij} = 0 \quad \forall i \in \mathcal{N}^I \quad (4.19d)$$

$$p_{ij} \leq \bar{P}_l \quad \forall i \in \mathcal{N}^A, \forall j \in \mathcal{N}_i, l : i \leftrightarrow j \quad (4.19e)$$

Note that this model includes line loading constraints to model congestion, but does not include voltage constraints. Mathematically, the power flow constraints are equivalent to the DC power flow model: with a change of variable $p_{ij} := B_{ij}\theta_{ij}$, where B_{ij} is the susceptance on line $l : i \leftrightarrow j$ and θ_{ij} is the voltage phase angle difference, we get the DC power flow

model directly. This model is a linear approximation of the full AC power flow model. It is generally considered to be a good approximation only for high voltage transmission – not microgrids – but it allows line flow constraints and network congestion to be incorporated into the framework to a first order. It is used here as a stepping stone to full linearized power flow incorporating voltage drop and also as a computationally less expensive approach.

We will show that by using consensus ADMM, this problem can be solved with synchronized peer-to-peer communication. That is, the “global” updates in each iteration depend only on information from neighboring nodes in the network.

To set up the consensus ADMM iteration, we define the coupling variables as the power transfer between nodes. For every node i define the local targets as p'_{ij} for all $j \in \mathcal{N}_i$ and the global reference as p_{ij} . As a result, there are two distinct variables p_{ij} and p_{ji} describing the same power transfer across line $l : i \leftrightarrow j$. We associated p_{ij} with node i and p_{ji} with node j . To be “consistent”, each must be equal to the negative of the other. We define the global constraints to enforce this consistency, that is, that $p_{ij} = -p_{ji}$ for all lines l . By the nature of consensus ADMM, the global variables are consistent at every iteration. The local targets are not necessarily consistent at every iteration, but converge to consistency because they converge to their global counterparts.

To define these consistency constraints using a matrix G , let each set of power transfers $\{p_{ij}\}$ from node i to each of its neighbors be represented as a vector p_i . This is not to be confused with the power injection at node i . Let the neighbors be ordered arbitrarily so that the m 'th neighbor of i , denoted $\mathcal{N}_i(m)$, is node j , and likewise the m 'th element of p_i is p_{ij} . The same applies to the corresponding global variables p'_i . To represent the constraint set $p_{ij} + p_{ji} = 0, \forall l : i \leftrightarrow j$, G must be constructed with a row for each line, and the same number in each of the two columns corresponding to p_{ij} and p_{ji} . This is achieved by:

$$G := [G_1 \dots G_i] \tag{4.20a}$$

$$G_i := [G]_{lm} = \begin{cases} 1 & \text{if } l : i \leftrightarrow \mathcal{N}_i(m) \\ 0 & \text{otherwise} \end{cases} \tag{4.20b}$$

$$= \text{diag}(H_i^u) H^{uT} E_i \tag{4.20c}$$

G_i can either be defined element-wise or in terms of the undirected incidence matrix H^u and the matrix E_i which selects lines in the incidence matrix that are connected to i . It can be seen by carrying out the matrix multiplication that $G_i G_i^T$ is a diagonal matrix with a 1 on the l 'th diagonal iff $l \in \mathcal{L}_i$, where \mathcal{L}_i is the set of all lines incident to i . When summing G_i over i , each line l will pick up a 1 for each of its sending and receiving nodes, so

$$\tilde{G} := \sum_i G_i G_i^T = 2I \tag{4.21a}$$

$$\tilde{G}^\dagger = \tilde{G}^{-1} = \frac{1}{2}I \tag{4.21b}$$

With this in hand, we can write the global variable and price updates for ADMM at iteration k , assuming the local variables p'_i have been obtained at all nodes i . Let ξ , η , and μ

denote the constraint violation, averaged error, and integral error, and π the price on power transfer. By substituting these variables into (4.18) it can be seen that the global updates simplify to updates for each line $l : i \leftrightarrow j$ and their associated nodes i and j . These can be written, for all $l : i \leftrightarrow j$, as:

$$\xi_l^k = p'_{ij}{}^{k+1} + p'_{ji}{}^{k+1} \quad (4.22a)$$

$$\eta_l^k = \frac{1}{2}\xi_l^k = \frac{1}{2}(p'_{ij}{}^{k+1} + p'_{ji}{}^{k+1}) \quad (4.22b)$$

$$\mu_l^{k+1} = \mu_k + \rho\eta_l^k = \mu_k + \frac{\rho}{2}(p'_{ij}{}^{k+1} + p'_{ji}{}^{k+1}) \quad (4.22c)$$

$$p_{ij}{}^{k+1} = p_{ij}{}^{k+1} - \eta_l^k = \frac{1}{2}(p_{ij}{}^{k+1} - p_{ji}{}^{k+1}) \quad (4.22d)$$

$$p_{ji}{}^{k+1} = p_{ji}{}^{k+1} - \eta_l^k = \frac{1}{2}(p_{ji}{}^{k+1} - p_{ij}{}^{k+1}) \quad (4.22e)$$

$$\pi_{ij}{}^{k+1} = \mu_l^k \quad (4.22f)$$

$$\pi_{ji}{}^{k+1} = \mu_l^k \quad (4.22g)$$

Thus, the “global” update can actually be decomposed as an update for each line, which only involves information exchange with its sending and receiving nodes. This means that information exchange can be entirely peer-to-peer and does not need to be centralized. The key nuance to this is that the iteration must still be synchronized across all nodes for the convergence guarantee of ADMM, although asynchronous ADMM shows promising results in practice.

The local updates can be stated separately for nodes with agents and those without. For all internal nodes $i \in \mathcal{N}^I$, their local problem does not include decision variables and utility functions for any DERs because they don't have any. We will use their local problem to satisfy power transfer constraints $|p_{ij}| \leq \bar{P}_l$. Thus the local update at an internal node $i \in \mathcal{N}^I$ can be written as:

$$\begin{aligned} \{p'_{ij}\}^{k+1} &= \arg \min_{\{p'_{ij}\}} \sum_{j \in \mathcal{N}_i} \pi_{ij}^k{}^T p'_{ij} + \frac{\rho}{2} \|p'_{ij} - p_{ij}^k\|_2^2 \\ &\text{s.t. } |p_{ij}| \leq \bar{P}_l \\ &\quad \sum_{j \in \mathcal{N}_i} p'_{ij} = 0 \end{aligned} \quad (4.23)$$

This problem returns a set of power transfers, one for each neighbor j , and can be thought of as a simple projection.

For the agent nodes $i \in \mathcal{N}^A$, the local problem includes the line flow to their single neighboring node as well as the private decision of the agent. Let the agent connected be n , and their neighboring node be j connected by line l . The local problem is:

$$p'_{ij}{}^{k+1} = \arg \min_{p'_{ij}|p_n, z_n} -U_n(z_n) + \pi_{ij}^k{}^T p'_{ij} + \frac{\rho}{2} \|p'_{ij} - p_{ij}^k\|_2^2 \quad (4.24a)$$

$$\begin{aligned} \text{s.t. } (p_n, z_n) &\in \mathcal{C}_n \\ p'_n + p'_{ij} &= 0 \\ |p'_{ij}| &\leq \bar{P}_l \end{aligned}$$

By substituting $-p'_n$ for p'_{ij} , this clearly equivalent to the agent's augmented decision model \mathcal{P} defined in (4.4). Thus the local update at agent nodes can be written compactly as

$$p'_{ij}{}^{k+1} = -\mathcal{P}(-p_{ij}{}^k, -\pi_{ij}{}^k, \rho, \bar{P}_l) \quad (4.25)$$

Now, we can state the full set of equations that define the ADMM update. We solve and substitute for ξ , η , and μ variables to keep the equations relatively compact. The update equations at each iteration k are:

$$p'_{ij}{}^{k+1} = -\mathcal{P}^\rho(-p_{ij}{}^k, -\pi_{ij}{}^k, \bar{P}_l) \quad \forall i \in \mathcal{N}^A \quad (4.26a)$$

$$\{p'_{ij}\}^{k+1} = \arg \min_{\{p'_{ij}\}} \sum_{j \in \mathcal{N}_i} \pi_{ij}{}^{kT} p'_{ij} + \frac{\rho}{2} \|p'_{ij} - p_{ij}{}^k\|_2^2 \quad \forall i \in \mathcal{N}^I \quad (4.26b)$$

$$\text{s.t. } |p_{ij}| \leq \bar{P}_l$$

$$\sum_{j \in \mathcal{N}_i} p'_{ij} = 0$$

$$p_{ij}{}^{k+1} = \frac{1}{2}(p'_{ij}{}^{k+1} - p'_{ji}{}^{k+1}) \quad \forall i \in \mathcal{N}, j \in \mathcal{N}_i \quad (4.26c)$$

$$\pi_{ji}{}^{k+1} = \pi_{ij}{}^k + \frac{\rho}{2}(p'_{ij}{}^{k+1} + p'_{ji}{}^{k+1}) \quad \forall i \in \mathcal{N}, j \in \mathcal{N}_i \quad (4.26d)$$

The local update involves solving a private decision at the agent nodes and a simple projection at the internal nodes, and the global update is decomposed into simple summations. Key features are:

- All of the “for all” equations can be solved in parallel in both the local and global update
- The global update for the reference power transfer p_{ij} at each node i depends only on its target power and its neighbor's target power transfer. The same holds for the price, with the additional dependency on the previous price.
- At each iteration $\pi_{ij} = \pi_{ji}$, so the prices are symmetric and only one price need be computed and stored for each line.
- At each iteration $p_{ij} = -p_{ji}$ and $|p_{ij}| \leq \bar{P}$. Thus the global reference power transfers are feasible with respect to line loading at every iteration.
- The only network information required can be embedded in a messaging framework for each agent to communicate with its neighboring node along with a clock to synchronize the iterations.

The iteration can proceed until a stopping criteria is met, which is generally defined in terms of two residuals [14], or it can continue indefinitely in a real time setting.

Case 2: Linear and non-linear AC power flow with voltage constraints

This section extends the previous DC power flow model to linearized and full non-linear AC power flow. The section begins with the global update step and its solution, which is constructed so that it applies to both linearized and non-linear AC power flow. Then, the respective local updates for each of the linearized and non-linear models and the associated ADMM iterations are given. It is important to note that because non-linear AC power flow is non-convex, the standard convergence guarantee of ADMM does not hold for this model.

We define the coupling variables to be the voltage phasors at each node and the current phasors on each line. The linear (and invertible) constraint relating voltage to current is represented as a constraint in the global update, meaning that at every iteration of ADMM, the global reference voltage and current phasors satisfy Ohm's Law and map to a set of power flows that satisfy the AC power flow equations. The non-linear equation relating voltage and current to power, which is necessary because the resource agent's decision models are defined in terms of power, are included in the local update problem for each agent node. This equation is non-convex. Additionally, the local problem for all nodes includes the constraint on the minimum voltage magnitude, which is non-convex. We propose linearizing these constraint at each iteration and show that this linearization can be computed with trivial effort. As a consequence, if the iteration converges to a solution, then both the local targets and global references satisfy the full AC power flow equations and are a local minimum. The linearization of the voltage magnitude constraint is a convex restriction, so the voltage constraints are satisfied at every iteration by both the local targets and global reference variables. Likewise the line loading constraints, which we define in terms of current, are also satisfied at every iteration by both the local targets and global references.

We point out that including voltage drop and voltage constraints in the problem substantially changes the complexity of the global update. Whereas in the DC power flow model, the global update can be carried out with only neighbor-to-neighbor communication, the inclusion of the voltage magnitude necessitates that the global update use information from all nodes. It will be shown that the global update for voltage magnitude uses a linear mapping that involves the inverse of a modified edge Laplacian of the network. It is possible to side-step this by defining the coupling variables in a different way, as is done in [82], but this sacrifices having power flow solutions satisfying Ohm's law at every iteration and the features of the linearization described above. In radial networks where the DistFlow model is valid and its SOCP power flow relaxation is valid, this alternate approach may be advantageous; however, it does not apply directly to meshed networks.

Global variable and price update

To derive the model, we first analyze and solve the global update, which follows the form of the consensus update used previously.

To solve the global update step, let the coupling variables be the complex nodal voltages and the current transferred across each line. Let V_i be the voltage at node n , and I_{ij} be the current on line $l : i \leftrightarrow j$ flowing from i to j . These variables are *phasors* represented by complex numbers in rectangular coordinates. As in the previous section, where the coupling variables was the power flow on each line, we define two coupling variables for the current, I_{ij} and I_{ji} . By assuming the π model for power lines with no shunt admittance, $I_{ij} + I_{ji} = 0$ by conservation of flow. Including Ohm's Law relating voltage and current by the impedance Z_{ij} , the constraints that must be satisfied by the global update are:

$$V_i - V_j - Z_l I_{ij} = 0 \quad \forall l : i \rightarrow j \quad (4.27a)$$

$$I_{ij} + I_{ji} = 0 \quad \forall l : i \leftrightarrow j \quad (4.27b)$$

Note that (4.27a) defined with respect to the *directed* line $l : i \rightarrow j$ and that a constraint for $V_j - V_i - Z_l I_{ji} = 0$ for the same line is not explicitly defined because would be redundant and is satisfied implicitly by the intersection of constraints (4.27). The constraints hold for both the real and imaginary parts of each variable.

As in the prior section, this can be written in terms of a matrix G with a particular structure, with m denoting the m 'th neighbor of i :

$$G := [G_1 \quad \dots \quad G_N] \quad (4.28a)$$

$$G_i := \begin{bmatrix} G_i^V \\ G_i^I \end{bmatrix} := \begin{bmatrix} G_i^{VV} & G_i^{VI} \\ 0 & G_i^{II} \end{bmatrix} \quad (4.28b)$$

$$G_i^{VV} := [G^{VV}]_{lm} = \begin{cases} 1 & \text{if } l : n \rightarrow \mathcal{N}_i(m) \\ -1 & \text{if } l : \mathcal{N}_i(m) \rightarrow i \\ 0 & \text{otherwise} \end{cases} \quad (4.28c)$$

$$G_i^{VI} := [G^{VI}]_{lm} = \begin{cases} -Z_l & \text{if } l : i \rightarrow \mathcal{N}_i(m) \\ 0 & \text{otherwise} \end{cases} \quad (4.28d)$$

$$G_i^{II} := [G^V]_{lm} = \begin{cases} 1 & \text{if } l : n \leftrightarrow \mathcal{N}_i(m) \\ 0 & \text{otherwise} \end{cases} \quad (4.28e)$$

G_i^V , the upper block of G_i , enforces constraint (4.27a). G_i^I , the lower block, enforces (4.27b). Note that G_i^{II} is identical to the matrix used in the DC power flow case, defined by (4.20b), because the constraints have an identical structure representing flow conservation. Also as in the prior section, these constraints can be written in terms of standard basis vectors and the components of the incidence matrix H , with the addition of the diagonal matrix of line impedances Z . Let $|Z|^2 := Z \times \text{conj}(Z)$ denote the impedance magnitude squared.

$$\begin{bmatrix} G_i^V \\ G_i^I \end{bmatrix} = \begin{bmatrix} H_i^{dT} & -\text{diag}(H_i^s) Z H^{uT} E_i \\ 0 & \text{diag}(H_i^u) Z H^{uT} E_i \end{bmatrix} \quad (4.29)$$

It follows that

$$\tilde{G} := \sum_i G_i G_i^T = \begin{bmatrix} H^{d^T} H^d + |Z|^2 & -|Z|^2 \\ -|Z|^2 & 2|Z|^2 \end{bmatrix} \quad (4.30)$$

The ADMM iterations require computing the inverse (or pseudo-inverse) of \tilde{G} . This can be done analytically using the Schur complement if it is well-defined and invertible, in which case \tilde{G} is also invertible. For a block matrix $X := \begin{bmatrix} A & B \\ C & D \end{bmatrix}$, with D invertible and each block having conforming dimensions, the Schur complement of D in X is written as X/D and defined as $A - BD^{-1}C$. If the Schur complement is invertible, then there is a formula for the inverse of the block matrix involving the Schur complement's inverse. In the following paragraphs, we give the Schur complement, show that it is invertible, and qualitatively discuss the meaning of the inverse and conditions necessary for its stable numerical computation. After this is established, we will give the formula for \tilde{G}^{-1} and the following global update equations for ADMM.

Let M be the Schur complement of the lower-right block $2|Z|^2$ of \tilde{G} :

$$M := \tilde{G}/(2|Z|^2) = H^{d^T} H^d + \frac{1}{2}|Z|^2 \quad (4.31)$$

The formula for G^{-1} depends on the M^{-1} , the inverse of the Schur complement. It is straightforward to show that M is invertible. This is because it is the sum of a symmetric positive semi-definite matrix (the directed edge Laplacian $H^{d^T} H^d$) and a symmetric positive definite matrix $\frac{1}{2}|Z|^2$. By Weyl's inequalities, the smallest eigenvalue of the sum of two symmetric matrices is bounded below by the sum of the smallest eigenvalues of summands. In this case, the smallest eigenvalue of $H^{d^T} H^d$ is bounded below by 0: it is 0 if the network is meshed (contains any cycles), or is the (positive) algebraic connectivity of the network if it is radial (contains no cycles / is a tree). The smallest eigenvalue of $\frac{1}{2}|Z|^2$ is $\min_l \frac{1}{2}|Z_l|^2$, which is clearly given by the smallest line impedance magnitude, which is strictly positive. Thus, the smallest eigenvalue of M is bounded below by the sum of a nonnegative and a strictly positive number, which is strictly positive, so it is positive definite and invertible.

To compute M^{-1} numerically, it is important to consider also how close the smallest eigenvalue of M is to zero. If it is very close to zero, then the matrix is approximately singular, and numerical algorithms can fail to produce accurate inverses. This is especially important in the case of meshed networks because the lower bound on the smallest eigenvalue is given by the smallest line impedance, and we used a procedure to add nodes and lines with very small impedance $|Z|^0$. Moreover, there is no guarantee against having a small line impedance in an arbitrary network, so potential issues with small impedance lines are important to address. Without proof and using only an intuitive argument, we conjecture that as long as there are no cycles in the network where all lines in the cycle have approximately zero impedance, then the matrix M is well-conditioned; i.e. the smallest eigenvalue is not very close to zero. To see this, recall from the preliminaries that each 0 eigenvalue in the edge Laplacian corresponds

to an undirected cycle in the network. Each of these has an eigenvector with 1 or -1 as its l 'th element for every line l in the cycle, where the sign is given by the direction of the line as it appears in the undirected cycle. Call this eigenvector λ for some cycle. If all lines in the cycle have the same small impedance magnitude $|Z|^0$, then it is trivial to see that v is also an eigenvector of M with eigenvalue $\frac{1}{2}|Z|^0$. Thus, we have shown that if all lines in a cycle have very small impedance, then M will be ill-conditioned for computing the inverse. It is much more complex to describe the smallest eigenvalue when all lines in a cycle have different impedances, because in this case v is no longer an eigenvector of M , but some numerical examples show roughly that the size of each eigenvalue associated with a cycle scales with the largest impedance in the cycle. Therefore, as long as all lines in the cycle are not close to zero, M^{-1} can be computed numerically. This is why we require the ‘‘no short-circuit’’ assumption.

With the Schur complement M and its inverse M^{-1} , the inverse of \tilde{G} can be computed directly as:

$$\tilde{G}^{-1} = \begin{bmatrix} M^{-1} & \frac{1}{2}M^{-1} \\ \frac{1}{2}M^{-1} & \frac{1}{2}Z^{-2} + \frac{1}{4}|M|^{-1} \end{bmatrix} \quad (4.32)$$

Before stating the complete analytic solution to the global update in terms of M^{-1} , we solve for the error ξ and effective error η as used in (4.18):

$$\xi := \begin{bmatrix} \xi^V \\ \xi^I \end{bmatrix} := \begin{bmatrix} \sum_i G_i^V [V_i & I_i]^T \\ \sum_i G_i^I [V_i & I_i]^T \end{bmatrix} = \begin{bmatrix} \sum_i (H^d)_i^T V_i - \text{diag}(H_i^s) Z (H^u)^T E_i I_i \\ \sum_i \text{diag}(H_i^u) (H^u)^T E_i I_i \end{bmatrix} \quad (4.33)$$

The l 'th elements are simply:

$$\xi_l^V = V_i' - V_j' - Z_{ij} I_{ij}' \quad (4.34a)$$

$$\xi_l^I = I_{ij}' + I_{ji}' \quad (4.34b)$$

The effective error can now be computed using M^{-1} and the formula (4.32):

$$\eta := \begin{bmatrix} \eta^V \\ \eta^I \end{bmatrix} := \eta \tilde{G}^{-1} \xi = \begin{bmatrix} M^{-1} (\xi^V + \frac{1}{2}\xi^I) \\ \frac{1}{2} (M^{-1}(\xi^V + \frac{1}{2}\xi^I) + |Z|^{-2}\xi^I) \end{bmatrix} \quad (4.35)$$

With η and ξ computed, the rest of the global update follows mechanically by converting G_i^T into a summation operation over the lines l . Written altogether, the global update equations are:

$$\xi_l^{V^k} = V_i'^{k+1} - V_j'^{k+1} - Z_{ij} I_{ij}'^{k+1} \quad (4.36a)$$

$$\xi_l^{I^k} = I_{ij}'^{k+1} + I_{ji}'^{k+1} \quad (4.36b)$$

$$\eta_l^{V^k} = M^{-1} \left(\xi^V + \frac{1}{2}\xi^I \right) \quad (4.36c)$$

$$\eta_l^{I^k} = \frac{1}{2} \left(M^{-1}(\xi^V + \frac{1}{2}\xi^I) + |Z|^{-2}\xi^I \right) \quad (4.36d)$$

$$V_i^{k+1} = V_i'^{k+1} - \left(\sum_{l:\mathcal{L}_i^s} \eta_l^{V^k} - \sum_{l:\mathcal{L}_i^r} \eta_l^{V^k} \right) \quad (4.36e)$$

$$I_{ij}'^{k+1} = I_{ij}'^k - \eta_{l:i \leftrightarrow j}^I + \sum_{l:\mathcal{L}_i^s} Z_l \eta_l^V \quad (4.36f)$$

$$\mu_l^{V^{k+1}} = \mu_l^{V^k} + \rho \eta_l^{V^k} \quad (4.36g)$$

$$\mu_l^{I^{k+1}} = \mu_l^{I^k} + \rho \eta_l^{I^k} \quad (4.36h)$$

$$\pi_i^{V^{k+1}} = \sum_{l:\mathcal{L}_i^s} \mu_l^{V^{k+1}} - \sum_{l:\mathcal{L}_i^r} \mu_l^{V^{k+1}} \quad (4.36i)$$

$$\pi_{ij}^{I^{k+1}} = \mu_l^{I^{k+1}} - \begin{cases} \mu_l^{V^{k+1}} & \text{if } l : i \rightarrow j \\ 0 & \text{if } l : j \rightarrow i \end{cases} \quad (4.36j)$$

The mapping M^{-1} , which describes how the error on one line relates to the “effective error” on another, is the central part of the complexity of the global update. Unlike G , \tilde{G} , and M , which are sparse $L \times L$ matrices, M^{-1} is a dense $L \times L$ matrix. Entries close to the diagonal represent lines that are close together in the network, and thus their values are relatively large and tend to get smaller away from the diagonal, but do not go to zero. It is not present in the updates for the lossless power flow model of the previous section because that model neglects voltage magnitude, which is the nodal “state” from a network flow perspective, and thus the error on all lines are independent and the analogous mapping is simply the identity scaled by a constant. In power flow models that incorporate voltage drop into the problem, the error on all lines has an affect on the other. A potential extension to this formulation could use a simpler approximation for M^{-1} , for example a low-rank approximation following from its eigendecomposition (spectral decomposition), or a sparse approximation. With a sparse approximation, the approximate effective error η would not depend on error from all other lines, only the most significant ones. The close relationship between M to the edge Laplacian, and the fact that the nonzero eigenvalues of the edge Laplacian are identical to those of the nodal Laplacian, point to an avenue for studying spectral decomposition techniques.

Local update

This section gives the local update equations. As with the previous model of lossless lines, the local update can be partitioned into update equations for agent nodes and for internal nodes. In this model, the local update includes the nonlinear power flow equations and constraints on voltage magnitude. These equations are non-convex and do not have an analytic inverse, and are typically solved iteratively using a procedure such as the Newton-Raphson or Gauss-Siedel methods over the entire network. Alternatively, linearizations or

convex relaxations of the constraints are used to solve the problem with specific algorithms for linear and convex problems. There are no known algorithms to solve the optimal power flow problem that are guaranteed to converge in general. In some cases, such as for radial networks, convex relaxations can be structured so that they can be guaranteed to be exact.

Here, we propose an approach that performs a linearization at each local update. This is similar to [71]. We show that this linearization can be performed efficiently by the decomposition structure of ADMM and yields a local update optimization problem that is convex with dimension on the order of the number of lines incident to each node. Before stating the non-linear and linearized subproblems, we introduce some notation and set up the model.

Let S_{ij} be the phasor representing the complex power flow on line $l : i \rightarrow j$. It is given from the voltage and current as:

$$S_{ij} = V_i \times \text{conj}(I_{ij}) \quad (4.37)$$

To write the equation for real power flow, let the vectors u_i^V and u_{ij}^I represent the real and imaginary parts of V_i and I_{ij} as their coordinates in \mathbb{R}^2 . The real power flow, p_{ij} , which is the real part of the complex number S_{ij} , is:

$$p_{ij} = u_{ij}^{I\text{T}} u_i^V \quad (4.38a)$$

$$u_i^V := \begin{bmatrix} \text{Re}(V_i) \\ \text{Im}(V_i) \end{bmatrix} \quad (4.38b)$$

$$u_{ij}^I := \begin{bmatrix} \text{Re}(I_{ij}) \\ \text{Im}(I_{ij}) \end{bmatrix} \quad (4.38c)$$

where Re and Im denote the real and imaginary parts of a complex number.

In this model, we will define the line-loading and injection constraints in terms of the magnitude of current instead of apparent power. This is done mainly for convenience, but it does have a physical justification in that the circuit protection that enforces line limit constraints, especially on lower voltage systems, is typically rated in terms of current. In any case, it is not a significant assumption because constraints on apparent power can be defined with an analogous formulation, and in the case that the voltages are close to nominal values, the constraint in terms of apparent power and current are approximately equivalent. The result is that there is no need to define reactive power explicitly in the model. There are still implicit reactive power flows, but its interaction with the constraints are fully described by the currents and voltages.

At an internal node i , the non-convex local update is:

$$\begin{aligned} (u_i^{V'}, \{u_{ij}^{I'}\})^{k+1} = & \arg \min_{u_n^{V'}, \{u_{ij}^{I'}\} | p_n, z_n} \left[\text{Re}(\pi_n^{V^k}) \quad \text{Im}(\pi_n^{V^k}) \right] u_n^{V'} + \frac{\rho}{2} \|u_n^{V'} - u_n^{V^k}\|_2^2 \\ & + \sum_{m \in \mathcal{N}_n} \left[\text{Re}(\pi_{ij}^{I^k}) \quad \text{Im}(\pi_{ij}^{I^k}) \right] u_{ij}^I + \frac{\rho}{2} \|u_{ij}^{I'} - u_{ij}^{I^k}\|_2^2 \\ \text{s.t. } & \sum_{j \in \mathcal{N}_i} u_{ij}^{I'} = 0 \end{aligned} \quad (4.39a)$$

$$\|u_{ij}^I\|_2^2 \leq \overline{I_{ij}}, \quad \forall m \in \mathcal{N}_n \quad (4.39b)$$

$$\|u_i^{V'}\|_2^2 \leq |\overline{V_n}|^2 \quad (4.39c)$$

$$- \|u_i^{V'}\|_2^2 \leq -|\underline{V_n}|^2 \quad (4.39d)$$

The only non-convex constraint is the lower bound on voltage magnitude in rectangular coordinates (4.39d). Note that this problem is separable with respect to the voltage and current variables.

At an agent node n , recall that by construction an agent has only one line l connecting to a neighbor j . Thus, we can define the power withdrawal by agent n as $p_n = -p_{nj} = -u_{ij}^{I\top} u_n^V$ in terms of the voltage and current flow to neighbor j . The non-convex local update is:

$$(u_n^{V'}, u_{nj}^{I'})^{k+1} = \arg \min_{u_n^{V'}, u_{nj}^{I'} | p_n, z_n} -U_n(z_n) \quad (4.40a)$$

$$\begin{aligned} & + \left[\operatorname{Re}(\pi_n^{V^k}) \quad \operatorname{Im}(\pi_n^{V^k}) \right] u_n^{V'} + \frac{\rho}{2} \|u_n^{V'} - u_n^{V^k}\|_2^2 \\ & + \left[\operatorname{Re}(\pi_{nj}^{I^k}) \quad \operatorname{Im}(\pi_{nj}^{I^k}) \right] u_{nj}^{I'} + \frac{\rho}{2} \|u_{nj}^{I'} - u_{nj}^{I^k}\|_2^2 \end{aligned}$$

$$\text{s.t. } (p_n, z_n) \in \mathcal{C}_n \quad (4.40b)$$

$$p_n + u_{nj}^{I\top} u_n^{V'} = 0 \quad (4.40c)$$

$$\|u_{nj}^I\|_2^2 \leq \overline{I_{nj}} \quad (4.40d)$$

$$\|u_n^{V'}\|_2^2 \leq |\overline{V_n}|^2 \quad (4.40e)$$

$$- \|u_n^{V'}\|_2^2 \leq -|\underline{V_n}|^2 \quad (4.40f)$$

Here, (4.40c) and (4.40f) are both non-convex constraints.

These problems could be left in non-convex forms and solved with an appropriate non-convex solver. Even if they are solved exactly, there is still no guarantee that the ADMM iteration will converge, as the convergence result is only for convex problems. It is also likely that solving the non-convex problem exactly at each stage is not necessary because it will be overwritten at the next iterate. Thus, we propose instead to linearize the non-convex constraints using a first-order Taylor approximation around the the global reference of the coupling variables. The intuition behind this is that because the ADMM iterations include the quadratic penalty on the distance from the global copy, that if we take the global copy as the linearization point, the solution will stay near the linearization point, making the linear approximation a good approximation. The constraints are re-linearized every iteration, which it turns out is not computationally intensive. Because of the re-linearization, it follows that if the algorithm converges to a fixed point, then it satisfies the full AC power flow equations, and is at least a local optimum of the non-linear and non-convex optimal AC power flow problem with voltage and line loading constraints.

To compute the linearization, we use the gradient of $u_{ij}^{I\top} u_i^{V'}$ at the point (u_i^V, u_{ij}^I) . The gradient is simply $\begin{bmatrix} u_{ij}^{I\top} & u_i^{V\top} \end{bmatrix}$. Thus, the linear Taylor approximation for p_{ij} with respect

to $(u_i^{V'}, u_{ij}^{I'})$ around the point (u_i^V, u_{ij}^I) is:

$$\begin{aligned} p'_{ij} &\approx u_{ij}^{I\text{T}}(u_i^{V'} - u_i^V) + u_i^{V\text{T}}(u_{ij}^{I'} - u_{ij}^I) + u_{ij}^{I\text{T}}u_i^V \\ &= u_{ij}^{I\text{T}}u_i^{V'} + u_i^{V\text{T}}u_{ij}^{I'} - u_{ij}^{I\text{T}}u_i^V \end{aligned} \quad (4.41)$$

Similarly, the linear approximation of voltage magnitude squared is

$$\begin{aligned} \|u_i^{V'}\|_2^2 &\approx 2u_i^{V\text{T}}(u_i^{V'} - u_i^V) + \|u_i^V\|_2^2 \\ &= 2u_i^{V\text{T}}u_i^{V'} - \|u_i^V\|_2^2 \end{aligned} \quad (4.42)$$

Both of these linearizations can be computed trivially at any operating point because the gradient is given directly by the global references. We only linearize the non-convex constraints. The other non-linear convex quadratic inequality constraints are left in their original form so as not to introduce additional unnecessary approximations. However, it may be advantageous from a computational perspective to linearize them as well so that the problem has only linear constraints.

The linearized convex local update at internal nodes $i \in \mathcal{N}^{\mathcal{I}}$ is:

$$\begin{aligned} (u_i^{V'}, \{u_{ij}^{I'}\})^{k+1} &= \arg \min_{u_n^{V'}, \{u_{ij}^{I'}\}|p_n, z_n} \left[\text{Re}(\pi_n^{V^k}) \quad \text{Im}(\pi_n^{V^k}) \right] u_n^{V'} + \frac{\rho}{2} \|u_n^{V'} - u_n^{V^k}\|_2^2 \\ &\quad + \sum_{m \in \mathcal{N}_n} \left[\text{Re}(\pi_{ij}^{I^k}) \quad \text{Im}(\pi_{ij}^{I^k}) \right] u_{ij}^{I'} + \frac{\rho}{2} \|u_{ij}^{I'} - u_{ij}^{I^k}\|_2^2 \\ \text{s.t.} \quad &\sum_{j \in \mathcal{N}_i} u_{ij}^{I'} = 0 \end{aligned} \quad (4.43a)$$

$$\|u_{ij}^{I'}\|_2^2 \leq \overline{I}_{ij}, \quad \forall m \in \mathcal{N}_n \quad (4.43b)$$

$$\|u_i^{V'}\|_2^2 \leq |\overline{V}_n|^2 \quad (4.43c)$$

$$-2u_i^{V^k\text{T}}u_i^{V'} \leq \|u_i^{V^k}\|_2^2 - |\underline{V}_n|^2 \quad (4.43d)$$

The linearized convex local update at agent nodes $n \in \mathcal{N}^{\mathcal{A}}$ is:

$$(u_n^{V'}, u_{nj}^{I'})^{k+1} = \arg \min_{u_n^{V'}, u_{nj}^{I'}|p_n, z_n} -U_n(z_n) \quad (4.44a)$$

$$\begin{aligned} &+ \left[\text{Re}(\pi_n^{V^k}) \quad \text{Im}(\pi_n^{V^k}) \right] u_n^{V'} + \frac{\rho}{2} \|u_n^{V'} - u_n^{V^k}\|_2^2 \\ &+ \left[\text{Re}(\pi_{nj}^{I^k}) \quad \text{Im}(\pi_{nj}^{I^k}) \right] u_{nj}^{I'} + \frac{\rho}{2} \|u_{nj}^{I'} - u_{nj}^{I^k}\|_2^2 \end{aligned}$$

$$\text{s.t.} \quad (p_n, z_n) \in \mathcal{C}_n \quad (4.44b)$$

$$p_n + u_{nj}^{I\text{T}}u_n^{V'} + u_n^{V\text{T}}u_{nj}^{I'} - u_{nj}^{I\text{T}}u_n^V = 0 \quad (4.44c)$$

$$\|u_{nj}^{I'}\|_2^2 \leq \overline{I}_{nj} \quad (4.44d)$$

$$\|u_n^{V'}\|_2^2 \leq |\overline{V}_n|^2 \quad (4.44e)$$

$$- \|u_n^{V'}\|_2^2 \leq -|\underline{V}_n|^2 \quad (4.44f)$$

Together with the global update, these local updates define the ADMM iteration for non-convex non-linear AC power flow. The key features are:

- As for general consensus ADMM, the local updates can be computed in parallel without sharing private information.
- The voltage phasors given by the global reference update completely define an AC power flow solution at every iteration.
- If the iterations converge, then the local target power injections give a solution to the AC power flow problem.
- The global update cannot be solved entirely in parallel. It is first necessary to map the error on each line through a dense matrix.

As an alternative to this approach, a static linearization or a convex relaxation that is not updated every iteration could be used to approximate AC power flow. These would be guaranteed to converge, but not necessarily to a solution of the AC power flow problem.

Discussion

The preceding section shows how the ADMM algorithm can be used to find the optimal power dispatch through agent-based decisions with coordination. At each iteration of the algorithm, agents receive a price signal and a reference quantity, and then update their quantity. As the process repeats, the prices and quantities converge asymptotically to an equilibrium. This can be thought of as a bidding system and leads to potential new markets for electricity incorporating DERs. Here we share some perspective on important considerations for implementing a system based on this algorithm in real electricity networks: monetary transactions associated with the bidding, synchronization of bidding, and bidding over horizons with multiple time periods.

The variable π that results from the iteration gives an efficient per unit price for electricity, which could be used as the basis for new electricity markets. In the lossless line formulation ignoring voltage drop, the price for power injection is defined at each agent node. This can be used as a price for agents participating in the market. At internal nodes, the prices are defined for power transfer on each line. While these are not directly presented to agents, they can be used to analyze congestion. In the full AC power flow formulation, the price is not directly defined in terms of power, but in terms of voltage at each node and current on each line. These prices also have two dimensions as they are phasors, which is a mathematical convention that is not intuitive to general market participants without a power systems engineering background. However, these complex prices can be translated to prices on real

and reactive power for any given voltage and current phasors, so this is not a blocking issue for structuring transactions. On the other hand, it also opens up a new avenue to consider pricing directly in terms of these voltage and current phasors, and to explore whether using phasor measurement units (PMUs) for electricity metering presents any advantages for efficient market design.

ADMM relies on the quadratic penalty term with parameter ρ as part of the cost function. How does this fit into a material transaction? We offer two approaches for exploration. One is to leave it out of any real monetary transaction and require that agents include this penalty in their decision – essentially programming the rule directly into the agent. Enforcing this would require both a set of standards and specifications along with a system to validate that the software agents conform to these. As any movement to use this kind of system would require large-scale institutional coordination among utilities, industry, government, and consumer advocates, developing this enforcement infrastructure is possible to consider. A second more decentralized approach is to make the penalty a real bidding cost that agents pay *each time* they update their bid; i.e. at each iteration. Two advantages to this are that it does not require additional rules and standards, and that it generates surplus revenue for the system. On the other hand, it imposes additional costs on the participants. Clearly, either of these options demand much more study, but we point out a couple of observations here about the second option of including the cost.

First, although these payments would generate surplus revenue for the system, it should not be confused with a per-unit tax or charge, especially in that it does not affect the equilibrium price and quantity. Second, at each iteration of the bidding, the agent always has the option to pay zero penalty by not changing their quantity from the reference. Thus, any penalty they pay at each iteration for adjusting their bid is less than their payoff from the new bid. However, this does not say definitively how their cumulative penalties over all iterations compare to their payoff from the final price and quantity. If these costs are large, then there would be a large payment to the system that might be unfair or could cause individuals to opt not to participate. Third, the payment depends strongly on ρ , but this parameter also affects how long the bidding takes to converge, so it cannot be adjusted in isolation solely to control this payment amount. Additional mathematical analysis on the cumulative payment would help to clarify how large these costs are and how they depend on ρ . Especially in the case that they are large, mechanisms to re-distribute excess surplus will be important.

The classic ADMM algorithm is synchronous, meaning that all agents update their bids before each iteration is complete. In practice, with large numbers of agents, perfect synchrony is unrealistic. Even if the algorithm is structured to be nominally synchronous, there will inevitably be communication failures or computational delays that prevent agents from submitting their updates reliably on time, and it necessitates choosing a regular interval for the iterations that may need to be overly long to nominally allow all agents to compute an update. Designing the iteration to be partially asynchronous, where only a subset of agents update their bids in a period of time, or fully asynchronous and event-driven, where the global variables are updated in response to each agent's update will facilitate deployment in

real systems and may even improve the algorithm’s convergence time. When the algorithm is asynchronous, the convergence analysis becomes more complicated, but approaches proposed in [125, 131] and more general methods described in [123] are promising foundations for an asynchronous approach.

Lastly, we bring the multi-period time horizon back into focus. The iteration presented in this chapter did not have an explicit temporal dimension, rather we treated it implicitly as happening across all time periods in a horizon simultaneously. There are a few important details and opportunities about this. First, all of the bidding described in this section is looking ahead into the future, so the agents’ models are based on forecasts, which are made explicit in the definition of the agent models.⁶ Second, the ADMM global and price updates for each time period are independent; the time periods are only coupled through agents’ decision models, and specifically battery agents or load agents with time-shiftable loads. So a battery will choose their quantity for all time periods simultaneously, but the price and reference quantities are adjusted in parallel for each time period. In an asynchronous architecture, agents would not need to update quantities for all time periods simultaneously, which could reduce the information processing burden by only having agents’ publish updates where there are significant changes. Lastly, there is a strong justification for using time periods of different lengths with periods growing longer the further out they are in the time horizon. For example, the first 1 to 2 hours could be discretized into 5 to 15 minute periods, the next 4 to 8 hours in 30 minute to 1 hour periods, the next 12 to 36 hours in 2 to 4 hour periods, and so on until the horizon extends at least several days and up to a week. This gives granularity in the short run while approximating the long run and limiting the number of time periods in the decision models. Using this approach, the markets can be extended even to monthly or annual scales to capture long term demand forecasts and the seasonality of renewables and guide investment decisions. It is important to keep in mind that the quantities are averaged over each period and thus obscure fluctuations within periods; i.e. the average price over a day may be very different from the peak price. Therefore, future prices will not accurately capture the charging and discharging constraints of storage within the period, but it does give an approximation of the relative value of energy in one period vs another.

As time moves forward, the longer periods move up in the horizon and get discretized into smaller periods. For example, at 1 AM, the 2 hour period from 9 AM to 11 AM would become 4 new 30 minute periods, and at 7 AM, these 30 minute periods would be discretized into 15 minute periods. The exact parameterization of when to increase the resolution and to what granularity should be designed to balance the benefits of increasing resolution against the additional complexity of having more periods. The key idea is that as a future period approaches real time, the period reduces to the order of minutes so the

⁶Briefly on the effect of forecast uncertainty, we can say that the decomposition still applies with respect to maximizing expected value, and if each agents’ forecasts are independent, then each agent can separately optimize under their own uncertainty and collectively maximize total welfare. However, when forecasts are not independent, the decentralized approach either requires shared information about forecast covariance, or it will be suboptimal.

prices approximate real-time pricing. When a period is discretized, there is an opportunity to “warm-start” the bidding by taking whatever the price and reference quantities were for the longer period and using it as the initial point for each of the smaller periods that come from the discretization. This means that it may only take a few bidding iterations for the dispatch to converge to a new optimum with increased resolution. Similarly, as time moves forward and the system begins to consider a new period, such as several days ahead, the prices and reference quantities can be warm-started to something like the prior last period or some other simple heuristic like an autoregressive moving average that gives an estimate of the equilibrium.

When the future period moves to real time, the bidding closes and control is handed off from the tertiary layer to the secondary. The prices are fixed and the quantities become reference injection or withdrawal setpoints at each node for the secondary control system. It is not necessary that the bidding has converged exactly to an equilibrium; there will generally be an imbalance between supply and demand that is handled by the primary and secondary controls. There is an opportunity to keep the quadratic penalty term with parameter ρ in play as a penalty on the tracking error; i.e. penalizing each agent from deviating from their forward commitment. Conceptually, this is in line with the idea of explicitly charging the cost penalty discussed above. The next section focuses on the real time system rather than the forward market through a novel secondary control that continues to include a type of bidding that incorporates economic information while re-balancing the system and stabilizing voltage and frequency.

4.4 Economic secondary control with online ADMM

In this section we use ADMM to derive a secondary control system that solves the optimal power flow problem in real time. This is closely connected in spirit to foundational papers [30, 21, 37] and other works that have derived feedback controllers that use the system dynamics to solve an optimization problem in real time. However, these specific papers rely on quadratic or strongly convex cost functions to derive the feedback controller, which limits their applicability to more general DER participation in smart grids. The key difference with the approach here is not that the optimization problem seeks to maximize utility instead of cost, but that it applies to more general concave utility functions (or equivalently convex cost functions) that are not required to be quadratic or even differentiable. The result is a novel approach to secondary control that converges to the optimal power flow solution.

The technique used here is first to state the optimal power flow problem and transform it into a form for consensus ADMM. However, unlike in the previous section, the coupling constraints will be enforced in the global update, and these constraints are exactly the power flow equations. Instead of being solved computationally, these constraints are solved physically. Next, the solution to the global update is derived analytically in terms of the price, network parameters, and local target variables. Lastly, the feedback control system is reverse-engineered so that it converges in steady state to the solution to the global update problem.

Thus, the global update is solved by the physical system rather than computationally, and the ADMM iterations proceed in real time. The real time iteration consists of each agent solving for their local target, then control system setpoints and/or gains are updated, then the system runs for a period of time until it reaches steady state, then the price is updated, and the iteration repeats. In practice, these iterations could proceed on the order of tens of seconds to minutes.

This online algorithm complements the preceding section. The ADMM-based bidding described previously can run as a forward market; however, due to changing conditions, forecast inaccuracy, model error, or incomplete convergence, its solution will not in general satisfy the power flow equations exactly, and some feedback control will be necessary to maintain stability in real time. The traditional function of secondary control is to re-balance the system after disturbances to try and stay close to the solution of the optimal power flow problem. The novel online secondary control algorithm proposed here performs this re-balancing in a way that allows market participation to continue in real time, thus responding to the changing state of the system in an economically optimal way.

We apply this approach using only the DC power flow model for simplicity and to illustrate the concept. The extension to full power flow is left for future work.

Optimal control as consensus ADMM

The optimization problem we are trying to solve is simply maximizing the sum of all individual utility functions subject to the DC power flow constraints. Before proceeding, it is important to specify what is controllable in this context. We assume that all batteries and solar resources have a controllable power output, but that the loads do not. How the resources are controlled does not come into play directly until the next section, but we will assume they are controlled with active power vs. frequency droop. We assume the loads are not controllable mostly to establish that the approach does not require all loads to be continuously controllable – in practice most loads can only be turned on or off – but also because we assume that by this point the electricity user or their smart load agent has already adjusted their planned consumption as much as they are willing in response to the forward price and that additional real-time control would be a nuisance. We also assume the microgrid is islanded, so there is no interconnection agent.

Assume that the n 'th agent is at node i , and that they are numbered so that $n = i$. Let p denote the vector of power injections at all controllable nodes, and p^u be vector of power injections at all other nodes. Similarly let θ be the vector of voltage phase angles at each controllable node and θ^u the angle at all other nodes. Let the nodes be ordered so that all controllable nodes appear before uncontrollable nodes in the ordering. Let B be the nodal susceptance matrix, the imaginary part of the nodal admittance matrix. Let θ be the voltage phase angle at each node, and partitioned analogously to p into the angle at controllable and

uncontrollable nodes θ^c and θ^u . The DC power flow equations are:

$$\begin{bmatrix} p \\ p^u \end{bmatrix} = \begin{bmatrix} B^{cc} & B^{cu} \\ B^{uc} & B^{uu} \end{bmatrix} \begin{bmatrix} \theta \\ \theta^u \end{bmatrix} \quad (4.45)$$

The block form of B is given so that we may compute the reduced power flow equations using the Kron reduction. See [29] for a detailed discussion of the Kron reduction that includes these reduced power flow equations and a proof that it is well-defined for the nodal admittance matrix. The Kron reduction uses the reduced matrix B^{red} and the accompanying matrix B^{ac} to define the reduced DC power flow equations:

$$B^{\text{ac}} := -B^{cu}B^{uu-1} \quad (4.46a)$$

$$B^{\text{red}} := B^{cc} + B^{\text{ac}}B^{uc} \quad (4.46b)$$

$$p = B^{\text{red}}\theta - B^{\text{ac}}p^u \quad (4.46c)$$

These equations give the controllable power injection in terms of the angle at the controllable nodes and the power injection at uncontrollable nodes; the angle at uncontrollable nodes has already been solved for in this form and does not appear in the equations. To connect the power flow to the agent's consumption, let p' denote the target power consumption for each controllable resource, with $p' = -p$. Thus, we can state the optimization problem as

$$\min_{p', p, \theta} \sum_n -U_n(z_n) \quad (4.47a)$$

$$s.t. (p'_n, z_n) \in \mathcal{C}_n \quad (4.47b)$$

$$p' + p = 0 \quad (4.47c)$$

$$p = B^{\text{red}}\theta - B^{\text{ac}}p^u \quad (4.47d)$$

This is conveniently in the form for consensus ADMM, with p' as the local target, and p as the global coupling variable. Unlike in the previous section, the coupling constraint (i.e. the reduced DC power flow equations) is stated in terms of the global variables. This is because we are developing an online approach that will exploit the fact that the physics of the system solve the coupling constraint, whereas in the previous section it was advantageous for computational simplicity to state the coupling constraint in terms of the local update.

The ADMM iterations to solve (4.47) are:

$$p_n^{k+1'} = \arg \min_{p'_n | z_n} -U_n(z_n) + \pi p'_n + \frac{\rho}{2}(p'_n - p_n^k)^2 \quad (4.48a)$$

$$s.t. (p'_n, z_n) \in \mathcal{C}_n \quad (4.48b)$$

$$p^{k+1} = \arg \min_{p | \theta} -\pi^T p + \frac{\rho}{2} \|p' - p\|_2^2 \quad (4.48c)$$

$$s.t. p = B^{\text{red}}\theta - B^{\text{ac}}p^u \quad (4.48d)$$

$$\pi^{k+1} = \pi^k + \rho(p_n^{k+1'} - p^{k+1}) \quad (4.48e)$$

The global update can be solved analytically:

$$p^{k+1} = B^{\text{red}} \left(\frac{1}{\rho} \pi + p^{k+1} + B^{\text{ac}} p^u \right) - B^{\text{ac}} p^u \quad (4.49)$$

In the next section, we show how a droop control system with appropriate biases can be designed so that the power p converges to this value in steady state.

Before proceeding, we note one limitation of the model (4.47), which is that it does not include line loading constraints or model congestion, unlike the models of the previous section. If these inequality constraints are included, we can no longer solve the global update analytically.

Reverse engineered secondary control

In this section, we show how appropriate biases can be set so that the power injections converge to the solution of the global update in steady state, assuming the controllable nodes use active power / frequency droop. We first derive the steady-state power injections in the time domain. We then set the steady-state value equal to (4.49), which yields a formula for optimal biases.

The active power / frequency droop equations are:

$$\dot{\theta} = -K(p - \hat{p}) \quad (4.50a)$$

$$= -KB^{\text{red}}\theta + K(\hat{p} + B^{\text{ac}}p^u) \quad (4.50b)$$

where $K \succ 0$ is a positive-definite diagonal matrix of droop gains, \hat{p} is the power bias, and $\dot{\theta}$ is the angular frequency. It is critical that all droop gains are positive for the system to be stable. The angular frequency is implicitly defined as the shift relative to the nominal synchronous frequency, i.e. 50 Hz in most of the world or 60 Hz in North America, and the angles themselves are likewise in this reference frame. The second equation follows from substituting the DC power flow equations.

In order to write the solution in the time domain, it is necessary to first analyze the matrix KB^{red} . First, the system is exponentially stable if and only if KB^{red} is positive definite (i.e. iff all eigenvalues of KB^{red} are strictly greater than zero). We can see directly that this is not the case, because B^{red} is a weighted nodal Laplacian,⁷ and thus has 0 as an eigenvalue, with $\mathbf{1}$ as the associated eigenvector. Therefore 0 is also an eigenvalue of KB^{red} , with $\mathbf{1}$ as the associated eigenvector. However, it can be shown that KB^{red} is positive semi-definite, and that the eigenvalue 0 is not repeated. This is because, for two symmetric matrices X and Y , if X is positive definite, then XY has the same number of positive, negative, and

⁷It is proved in [29] that the Kron reduction of the admittance matrix is also a nodal Laplacian, and thus has all the same structural properties of a generic admittance matrix

zero eigenvalues as Y . This fact follows from the identity $\lambda(XY) = \lambda(\sqrt{XY}\sqrt{X})$, where $\lambda()$ denotes the function yielding the set of eigenvalues of a matrix, and Sylvester's Law of Inertia. Thus, because B^{red} has no negative eigenvalues, a 0 eigenvalue with multiplicity 1, and all the remaining eigenvalues positive, the eigenvalues of KB^{red} have the same properties. Given this, we can diagonalize KB^{red} as follows:

$$KB^{\text{red}} = YDY^{-1} \quad (4.51)$$

$$D := \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix} \quad (4.52)$$

$$Y := [y_1 \dots y_N] \quad (4.53)$$

$$y_1 = \mathbb{1} \quad (4.54)$$

where Λ is a diagonal matrix of strictly positive eigenvalues. Each y_n is an eigenvector of KB .

With this diagonalization, we can proceed to write the solutions for θ and p over time. Assume that \hat{p} and p^u are constant inputs over the time horizon. The solutions for θ and p over time, denoted $\theta(t)$ and $p(t)$ are:

$$\begin{aligned} \theta(t) &= \exp(-KB^{\text{red}}t)\theta(0) + \int_0^t \exp(-KB^{\text{red}}(t-\tau))K(\hat{p} + B^{\text{ac}}p^u)d\tau \\ &= Y \exp(-Dt)Y^{-1}\theta(0) + Y \left(\int_0^t \exp(-D(t-\tau))d\tau \right) Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) \end{aligned} \quad (4.55a)$$

$$\begin{aligned} p(t) &= B^{\text{red}}\theta(t) - B^{\text{ac}}p^u \\ &= B^{\text{red}}Y \exp(-Dt)Y^{-1}\theta(0) \\ &\quad + B^{\text{red}}Y \left(\int_0^t \exp(-D(t-\tau))d\tau \right) Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \end{aligned} \quad (4.55b)$$

Note also that $B^{\text{red}}Y = K^{-1}YD$ by the definition of the diagonalization of KB^{red} . By making this substitution, solving the integral, and expanding the matrix exponential, we get

$$\begin{aligned} p(t) &= K^{-1}YD \begin{bmatrix} 1 & 0 \\ 0 & \exp(-\Lambda t) \end{bmatrix} Y^{-1}\theta(0) \\ &\quad + K^{-1}YD \begin{bmatrix} t & 0 \\ 0 & (I - \exp(-\Lambda t))\Lambda^{-1} \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \\ &= K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \exp(-\Lambda t) \end{bmatrix} Y^{-1}\theta(0) \\ &\quad + K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix} \begin{bmatrix} t & 0 \\ 0 & (I - \exp(-\Lambda t))\Lambda^{-1} \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \\ &= K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \exp(-\Lambda t) \end{bmatrix} Y^{-1}\theta(0) \end{aligned}$$

$$+ K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & I - \exp(-\Lambda t) \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \quad (4.56)$$

Finally, the steady-state solution is obtained by taking the limit as $t \rightarrow \infty$. Because $\Lambda \succ 0$, all the exponential terms decay to 0. Let the steady state solution be denoted simply by p :

$$p = K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \quad (4.57)$$

Our objective is for p to solve the global update equations. So we set p equal to p^{k+1} using equation (4.49):

$$B^{\text{red}}(\rho^{-1}\pi^k + p'^{k+1} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u = K^{-1}Y \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \quad (4.58)$$

$$\implies KB^{\text{red}}(\rho^{-1}\pi^k + p'^{k+1} + B^{\text{ac}}p^u) = Y \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) \quad (4.59)$$

$$\implies YDY^{-1}(\rho^{-1}\pi^k + p'^{k+1} + B^{\text{ac}}p^u) = Y \begin{bmatrix} 0 & 0 \\ 0 & I \end{bmatrix} Y^{-1}K(\hat{p} + B^{\text{ac}}p^u) \quad (4.60)$$

This equation is satisfied by:

$$\begin{aligned} \hat{p} &= K^{-1}YDY^{-1}(\rho^{-1}\pi^k + p'^{k+1} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \\ &= B^{\text{red}}(\rho^{-1}\pi^k + p'^{k+1} + B^{\text{ac}}p^u) - B^{\text{ac}}p^u \end{aligned} \quad (4.61)$$

Thus, if the bias for droop control is set by this formula, the steady-state power injections solve the global update.

To understand the information exchange required to compute these biases, consider the formula for a particular node n :

$$\hat{p}_n = B_n^{\text{red}}(\rho^{-1}\pi^k + p'^{k+1} + B_n^{\text{ac}}p^u) - B_n^{\text{ac}}p^u \quad (4.62)$$

$$\begin{aligned} &= (\rho^{-1}\pi_n^k + p_n'^{k+1} + B_n^{\text{ac}}p^u) \sum_{j \in \mathcal{N}_n} \text{Im}(Y_{nj}) \\ &\quad + \sum_{j \in \mathcal{N}_n} \text{Im}(Y_{nj})(\rho^{-1}\pi_j^k + p_j'^{k+1} + B_n^{\text{ac}}p^u) - B_n^{\text{ac}}p^u \end{aligned} \quad (4.63)$$

where B_n^{red} and B_n^{ac} are the n 'th rows of B^{red} and B^{ac} . Therefore, the only external information required to set the bias is the price and target power at neighboring nodes in the reduced network B^{red} , the power injection at neighboring nodes in the accompanying network B^{ac} , and the line admittance between these neighbors. The update can be computed with only information from its neighbors, so it can be accomplished with peer-to-peer communication. Contrast this against conventional secondary control, where generator biases are adjusted

by the centrally computed Area Control Error (ACE). Once the biases are obtained, the feedback control is done locally.

Note that the power injections at the uncontrollable nodes p^u may not in fact be constant, in which case the steady-state power injections will not be exactly the desired solution to the global update equations.

It is also important to consider how long the system will take to converge to steady-state, and whether this is practical. This is controlled by the smallest non-zero eigenvalue of KB^{red} , which is bounded below by the smallest eigenvalue of K . Because K is diagonal, its entries k_n are the eigenvalues. A typical value for K for fossil-fuel generators is in the neighborhood of 0.05 in per unit generator capacity notation, meaning a 5% increase in frequency corresponds to shedding the entire unit's capacity. This decays completely on the order of 1 minute. Converter interfaced generation is capable of responding much faster, so convergence on the order of seconds or tens of seconds is possible.

All of the above describes one step in the online ADMM iteration to “compute” the new global reference by physically steering the system to a new state. Once the system converges to this state, the price is updated as:

$$\pi_n^{k+1} = \pi_n^k + \rho(p_n'^{k+1} - p_n^{k+1}) \quad (4.64)$$

and the iteration repeats. The next iteration proceeds by updating the local targets p_n' given the state of the DERs and their associated utility functions, computing the biases \hat{p}_n using formula (4.61), waiting for the system to converge to a new p_n , and then updating the price by (4.64).

Discussion

The online ADMM algorithm is a novel approach to secondary control that rebalances the system in an economically optimal way. The forward bidding described in the previous section, also referred to as the tertiary layer, will generally not yield power injection setpoints that are perfectly balanced because of forecast error and model uncertainty. The power converters are built to automatically correct the imbalance between supply and demand by using droop control. Droop control adjusts the frequency of each DER proportionally to how much the DER's power injection deviates from its bias, which we denote \hat{p} . This stabilizes frequency and balances the supply and demand, or the injections and withdrawals, on the network. However, in doing so, the power injection of each generator deviates from its bias and the frequency deviates from the nominal value of 50 Hz or 60 Hz. The secondary control recomputes the biases iteratively, which we call rebalancing, until eventually the bias and the injections are equal and the frequency is restored to zero. Classic secondary control drives the system to a point that minimizes the Area Control Error (ACE), which is a composite of the deviation between the target injection from the tertiary layer and the frequency deviation. By using the online ADMM algorithm, the rebalancing instead moves the system towards the economically optimal operating point.

The economic information is captured by the agents' utility functions. For batteries, the utility functions capture the future value of stored energy, which is given by the prices in the forward markets of the tertiary layer. To achieve this, the battery agents must be simultaneously participating in the secondary control system and the tertiary bidding system, and there is information exchange between the two. The current state of charge of the battery is used in the agent's decision model in the tertiary system, and the future prices from the tertiary bidding system are used in the agent's decision model in the secondary system. In a synchronized implementation, it is possible that the secondary control action to adjust the real time biases of the DERs system could be synchronized with the bidding; i.e. both the real time biases and the forward bids could be updated something like every 15 or 30 seconds. As in the previous section, an asynchronous implementation is likely to be more flexible and robust, but further theoretical work to characterize the convergence of asynchronous approaches is needed.

The concept of warm starting discussed in Section 4.3 is useful to further clarify the interface between the tertiary and secondary layer. Suppose the smallest period in tertiary bidding layer is 5 minutes, starting at 9:00 AM. At 9:00 AM, the current state of the bidding layer, which are the prices and reference withdrawals, become the (negative of) the initial price and global reference in the online ADMM secondary control. This is a form of warm starting, as these initial values should be close to the economic equilibrium. Suppose the power injections take 15 seconds to converge to steady state. This means 20 ADMM iterations can take place before 9:05 AM when it starts over with a new price and global reference. Simulation experiments are needed to validate whether this is sufficient to converge to an economic equilibrium.

Lastly, we note that more study of the frequency deviations associated with this approach is necessary. While an increasing number of loads do not require strict regulation to 50 Hz or 60 Hz, many still do. Moreover, many protection systems, including automatic disconnection of DERs, are triggered by frequency deviations. The approach described here stabilizes frequency and does converge to an economic equilibrium that also has zero frequency deviation, but it remains to be studied how much frequency deviates before it converges and whether this is acceptable.

Conclusion

This thesis shows how energy management systems focusing on the value of electricity can improve access to electricity in low-carbon and zero-carbon microgrids. Chapter 1 conducts a high-level economic analysis of the costs of decentralized solar plus battery systems without exploring adjusting consumption in response to scarcity. The following chapters focus on managing electricity consumption and developing systems that improve the value of electricity as a constrained resource. Altogether, this work shows the efficacy of decentralized systems for increasing energy access while proposing ways to improve the technical operational systems for allocating electricity. It concludes here with three critical questions for researchers to shape how we approach designing accessible electricity systems.

Question 1: *What are the human interfaces to enable active prosumer participation in energy management systems?*

A critique of Chapters 2-4 is that they assume to varying degrees that individuals think about their electricity consumption rationally, with detailed knowledge of how much energy their appliances consume, and have quantified notions of value from electricity use, which is clearly not the case in reality. Although each of the chapters is careful to show that the systems themselves can still function to provide reliable electricity without these strong assumptions, the argument that they improve access through better value and more efficient use than the status quo requires a different level of individual engagement than exists today. Technology will be fundamental in this transition, specifically educational systems to make technical *knowledge* of electricity more accessible to people, such as to understand units of electricity and how much their activities and appliances consume, and also systems to collect and learn user's preferences so that agents can participate in transactional systems on their behalf. The fields of human-computer interaction, behavioral economics, and psychology can contribute here, and the emerging "smart-home" industry is an arena where real-life experience can validate design principles in practice. We should not, however, expect principles from households in one geography to necessarily hold in another, nor to necessarily hold in commercial and industrial settings.

Question 2: *What are the equity and justice implications of decentralization?*

This thesis argues for grid decentralization on its engineering and economic efficiency merits, but decentralization can have both positive and negative impacts on equity and justice. In the positive direction, decentralized systems provide an alternative means for those who are excluded from the status quo of electricity systems to gain access. In rural

regions throughout the global south, decentralized solar home systems and microgrids have brought unprecedented electricity access to millions of people. Decentralized systems also allow communities and individuals to own and control their electricity supplies, allowing them to use more renewable energy, lower their costs of energy, and increase their resilience to extreme weather and disruptions in the electricity system. One can also view decentralized and locally generated electricity as a form of restorative justice against the extractive legacy of fossil fuels. On the other hand, increasing decentralization is likely to lead to more variable and unequal access to electricity by potentially undermining societies' commitment to providing equal access to centralized grid infrastructure. Whether this potential for more variable access is equitable or just is subjective, but it is difficult to argue that increasing inequality is consistent with widely accepted social ideals such as those in the Sustainable Development Goals. Beyond the question of the distribution of access, we need to ask seriously how fragmenting electricity grids into smaller units affects the cohesion of society and debates over the role of electricity grids and infrastructure more generally as a common resource regulated by economic principles versus a public good.

Question 3: *How would the management systems described here change pathways to universal access?*

This is partly answered by question two, as equity and justice are a critical part of pathways to universal access, but there are other technical elements to this as well. International institutions rely on large scale electricity planning models on the time scales of years to decades to build and compare scenarios that achieve universal access, typically based on cost-minimization. Static models typically advise which regions are more cost effectively served by extending centralized grids vs decentralized solutions, and those with a dynamic component simulate expanding centralized grids while simultaneously deploying decentralized solutions. The maps of where to build when and the corresponding investment strategy advise electrification policy. If we incorporate the various operational regimes described in Chapters 2-4 into planning models of the kind addressed by Chapter 1, it might show, hypothetically, that a pathway that includes investing widely in decentralized solutions and incrementally expanding and interconnecting them with smaller scale networks may provide more access more quickly and efficiently.

References

- [1] Muhammad Raisul Alam, Marc St-Hilaire, and Thomas Kunz. “Peer-to-peer energy trading among smart homes”. In: *Applied energy* 238 (2019), pp. 1434–1443.
- [2] Greg Albright, Jake Edie, and Said Al-Hallaj. *A Comparison of Lead Acid to Lithium-ion in Stationary Storage Applications*. Tech. rep. AllCell Technologies LLC, 2012. URL: <http://www.batterypoweronline.com/main/wp-content/uploads/2012/07/Lead-acid-white-paper.pdf>.
- [3] Peter Alstone, Dimitry Gershenson, and Daniel M. Kammen. “Decentralized energy systems for clean electricity access”. In: *Nature Climate Change* 5.4 (2015), pp. 305–314. ISSN: 1758-678X. DOI: 10.1038/nclimate2512. URL: <http://www.nature.com/doifinder/10.1038/nclimate2512>.
- [4] Daniel 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 (2015), pp. 3583–3593.
- [5] M Imran Azim, Wayes Tushar, and Tapan K Saha. “Cooperative negawatt P2P energy trading for low-voltage distribution networks”. In: *Applied Energy* 299 (2021), p. 117300.
- [6] Goran Banjac et al. “Infeasibility detection in the alternating direction method of multipliers for convex optimization”. In: *Journal of Optimization Theory and Applications* 183.2 (2019), pp. 490–519.
- [7] MEME Baran and Felix F Wu. “Optimal sizing of capacitors placed on a radial distribution system”. In: *IEEE Transactions on power Delivery* 4.1 (1989), pp. 735–743.
- [8] Saule Baurzhan and Glenn P. Jenkins. “Off-grid solar PV: Is it an affordable or appropriate solution for rural electrification in Sub-Saharan African countries?” In: *Renewable and Sustainable Energy Reviews* 60 (2016), pp. 1405–1418. ISSN: 13640321. DOI: 10.1016/j.rser.2016.03.016. URL: <http://www.sciencedirect.com/science/article/pii/S1364032116002513>.
- [9] E. M. L. Beale and J. J. H. Forrest. “Global optimization using special ordered sets”. In: *Mathematical Programming* 10.1 (Dec. 1976), pp. 52–69. ISSN: 0025-5610. DOI: 10.1007/BF01580653.

- [10] Paul Bertheau et al. “Visualizing National Electrification Scenarios for Sub-Saharan African Countries”. In: *Energies* 10.11 (Nov. 2017), p. 1899. ISSN: 1996-1073. DOI: 10.3390/en10111899. URL: <http://www.mdpi.com/1996-1073/10/11/1899>.
- [11] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. 4th ed. Vol. 1. Athena Scientific, 2017. ISBN: 1886529264.
- [12] Mikul Bhatia and Nicolina Angelou. *Beyond Connections Energy Access Redefined, Technical Report 008/15*. Tech. rep. 2015, pp. 1–24. URL: <https://openknowledge.worldbank.org/handle/10986/24368>.
- [13] Roy Billinton and Ronald N. Allan. “Reliability of Electric Power Systems: An Overview”. In: *Handbook of Reliability Engineering*. Ed. by Hoang Pham. London: Springer-Verlag, 2003. Chap. 28, pp. 511–528. DOI: 10.1007/1-85233-841-5_{_}28. URL: http://link.springer.com/10.1007/1-85233-841-5_28%20http://www.springerlink.com/index/x3022415076173kj.pdf.
- [14] Stephen Boyd et al. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. In: *Foundations and Trends® in Machine learning* 3.1 (2011), pp. 1–122.
- [15] Christian Breyer et al. “On the role of solar photovoltaics in global energy transition scenarios”. In: *Progress in Photovoltaics: Research and Applications* 25.8 (2017), pp. 727–745. DOI: 10.1002/pip.2885.
- [16] Louis L. Bucciarelli. “Estimating loss-of-power probabilities of stand-alone photovoltaic solar energy systems”. In: *Solar Energy* 32.2 (1984), pp. 205–209. ISSN: 0038092X. DOI: 10.1016/S0038-092X(84)80037-7. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0038092X84800377>.
- [17] Fiona Burlig and Louis Preonas. “Out of the Darkness and Into the Light? Development Effects of Rural Electrification”. Berkeley, 2016. URL: <https://ei.haas.berkeley.edu/research/papers/WP268.pdf>.
- [18] C. Cader et al. “Global cost advantages of autonomous solar–battery–diesel systems compared to diesel-only systems”. In: *Energy for Sustainable Development* 31 (Apr. 2016), pp. 14–23. ISSN: 0973-0826. DOI: 10.1016/J.ESD.2015.12.007.
- [19] Catherina Cader, Philipp Blechinger, and Paul Bertheau. “Electrification Planning with Focus on Hybrid Mini-grids – A Comprehensive Modelling Approach for the Global South”. In: *Energy Procedia* 99 (Nov. 2016), pp. 269–276. ISSN: 1876-6102. DOI: 10.1016/J.EGYPRO.2016.10.116.
- [20] Debabrata Chattopadhyay, Morgan Bazilian, and Peter Lilienthal. “More power, less cost: Transitioning up the solar energy ladder from home systems to mini-grids”. In: *Electricity Journal* (2015). ISSN: 10406190. DOI: 10.1016/j.tej.2015.03.009.
- [21] Lijun Chen and Seungil You. “Reverse and forward engineering of frequency control in power networks”. In: *IEEE Transactions on Automatic Control* 62.9 (2016), pp. 4631–4638.

- [22] Ashish Cherukuri and Jorge Cortés. “Iterative bidding in electricity markets: rationality and robustness”. In: *IEEE Transactions on Network Science and Engineering* (2019).
- [23] Wesley J. Cole et al. “Utility-scale lithium-ion storage cost projections for use in capacity expansion models”. In: *2016 North American Power Symposium (NAPS)*. IEEE, Sept. 2016, pp. 1–6. DOI: 10.1109/NAPS.2016.7747866. URL: <http://ieeexplore.ieee.org/document/7747866/>.
- [24] *Copernicus Atmospheric Monitoring Radiation Service*. URL: <http://www.soda-pro.com/web-services/radiation/cams-radiation-service> (visited on 07/21/2019).
- [25] Emiliano Dall’Anese, Sairaj V Dhople, and Georgios B Giannakis. “Photovoltaic inverter controllers seeking AC optimal power flow solutions”. In: *IEEE Transactions on power systems* 31.4 (2015), pp. 2809–2823.
- [26] Uwe Deichmann et al. “The economics of renewable energy expansion in rural Sub-Saharan Africa”. In: *Energy Policy* 39.1 (2011), pp. 215–227. ISSN: 03014215. DOI: 10.1016/j.enpol.2010.09.034. URL: http://ac.els-cdn.com/S0301421510007202/1-s2.0-S0301421510007202-main.pdf?_tid=48f50c60-5c3d-11e7-95e0-0000aab0f6b&acdnat=1498680574_392c16ce7145d43a8782453e365fa191.
- [27] Boucar Diouf and Ramchandra Pode. “Potential of lithium-ion batteries in renewable energy”. In: *Renewable Energy* 76 (2015), pp. 375–380. ISSN: 09601481. DOI: 10.1016/j.renene.2014.11.058. URL: <http://www.sciencedirect.com/science/article/pii/S0960148114007885>.
- [28] K.C. Divya and Jacob Østergaard. “Battery energy storage technology for power systems—An overview”. In: *Electric Power Systems Research* 79.4 (2009), pp. 511–520. ISSN: 03787796. DOI: 10.1016/j.epsr.2008.09.017.
- [29] Florian Dorfler and Francesco Bullo. “Kron reduction of graphs with applications to electrical networks”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 60.1 (2012), pp. 150–163.
- [30] Florian Dörfler, John W Simpson-Porco, and Francesco Bullo. “Breaking the hierarchy: Distributed control and economic optimality in microgrids”. In: *IEEE Transactions on Control of Network Systems* 3.3 (2015), pp. 241–253.
- [31] John A. Duffie and William A. Beckman. *Solar Engineering of Thermal Processes, Fourth Edition*. Fourth. Hoboken, New Jersey: John Wiley & Sons, Inc., 2013. DOI: 10.1115/1.2930068.
- [32] Bruce Dunn, Haresh Kamath, and Jean-Marie Tarascon. “Electrical Energy Storage for the Grid: A Battery of Choices”. In: *Science* 334.6058 (2011), pp. 928–935. URL: <http://science.sciencemag.org/content/334/6058/928>.
- [33] Jonathan Eckstein and Dimitri P Bertsekas. “On the Douglas—Rachford splitting method and the proximal point algorithm for maximal monotone operators”. In: *Mathematical Programming* 55.1 (1992), pp. 293–318.

- [34] Tomaso Erseghe. “Distributed optimal power flow using ADMM”. In: *IEEE transactions on power systems* 29.5 (2014), pp. 2370–2380.
- [35] Masoud Farivar and Steven H Low. “Branch flow model: Relaxations and convexification—Part I”. In: *IEEE Transactions on Power Systems* 28.3 (2013), pp. 2554–2564.
- [36] Isa Ferrall. “Measuring electricity reliability of decentralized solar energy systems”. In: *C3E Women in Clean Energy Symposium*. Cambridge, MA., 2017.
- [37] Lingwen Gan and Steven H Low. “An online gradient algorithm for optimal power flow on radial networks”. In: *IEEE Journal on Selected Areas in Communications* 34.3 (2016), pp. 625–638.
- [38] Lingwen Gan et al. “Exact convex relaxation of optimal power flow in radial networks”. In: *IEEE Transactions on Automatic Control* 60.1 (2014), pp. 72–87.
- [39] J Duncan Glover, Mulukutla S Sarma, and Thomas Overbye. *Power system analysis & design, SI version*. Cengage Learning, 2012.
- [40] Lorenz Gollwitzer et al. “Rethinking the sustainability and institutional governance of electricity access and mini-grids: Electricity as a common pool resource”. In: *Energy Research & Social Science* 39 (2018), pp. 152–161.
- [41] Will Gorman, Stephen Jarvis, and Duncan Callaway. “Should I Stay Or Should I Go? The importance of electricity rate design for household defection from the power grid”. In: *Applied Energy* 262 (2020), p. 114494.
- [42] Michael Grant and Stephen Boyd. *CVX: Matlab Software for Disciplined Convex Programming, version 2.1*. <http://cvxr.com/cvx>. Mar. 2014.
- [43] Michael Grant and Stephen Boyd. “Graph implementations for nonsmooth convex programs”. In: *Recent Advances in Learning and Control*. Ed. by V Blondel, S Boyd, and H Kimura. Lecture Notes in Control and Information Sciences. Springer-Verlag Limited, 2008, pp. 95–110.
- [44] Jemma Green and Peter Newman. “Citizen utilities: The emerging power paradigm”. In: *Energy Policy* 105 (2017), pp. 283–293.
- [45] Mirna Grzanić et al. “Electricity cost-sharing in energy communities under dynamic pricing and uncertainty”. In: *IEEE Access* 9 (2021), pp. 30225–30241.
- [46] Gabriela Hug. “Integration of optimal storage operation into marginal cost curve representation”. In: *Energy Systems* 7.3 (2016), pp. 391–409.
- [47] Imelda, Matthias Fripp, and Michael J Roberts. *Variable Pricing and the Cost of Renewable Energy*. Tech. rep. National Bureau of Economic Research, 2018.
- [48] International Energy Agency. *World Energy Outlook 2016*. Tech. rep. 2016. URL: <https://www.iea.org/publications/freepublications/publication/WorldEnergyOutlook2016ExecutiveSummaryEnglish.pdf>.

- [49] International Energy Agency (IEA) and the World Bank. *Sustainable Energy for All 2017—Progress toward Sustainable Energy*. Tech. rep. Washington, DC: The World Bank, 2017. DOI: 10.1596/978-1-4648-1084-8. URL: http://gtf.esmap.org/data/files/download-documents/eegp17-01_gtf_full_report_for_web_0516.pdf.
- [50] International Renewable Energy Agency. *Solar PV in Africa: Costs and Markets*. Tech. rep. International Renewable Energy Agency, 2016. URL: https://www.irena.org/DocumentDownloads/Publications/IRENA_Solar_PV_Costs_Africa_2016.pdf.
- [51] International Renewable Energy Agency. *The Power to Change: Solar and Wind Cost Reduction Potential to 2025*. Tech. rep. International Renewable Energy Agency, 2016. URL: http://www.irena.org/DocumentDownloads/Publications/IRENA_Power_to_Change_2016.pdf.
- [52] Nicholas Kaldor. “A classificatory note on the determinateness of equilibrium”. In: *The review of economic studies* 1.2 (1934), pp. 122–136.
- [53] Francis Kemausuor et al. “Electrification planning using Network Planner tool: The case of Ghana”. In: *Energy for Sustainable Development* 19.1 (Apr. 2014), pp. 92–101. ISSN: 09730826. DOI: 10.1016/j.esd.2013.12.009. URL: <http://linkinghub.elsevier.com/retrieve/pii/S097308261300121X>.
- [54] Jip Kim and Yury Dvorkin. “A P2P-dominant distribution system architecture”. In: *IEEE Transactions on Power Systems* 35.4 (2019), pp. 2716–2725.
- [55] Mathias Koepke and Sebastian Groh. “Against the odds: The potential of swarm electrification for small island development states”. In: *Energy Procedia* 103 (2016), pp. 363–368.
- [56] Matt Kraning et al. “Dynamic network energy management via proximal message passing”. In: *Foundations and Trends® in Optimization* 1.2 (2013), pp. 73–126.
- [57] Jose Daniel Lara et al. “Computational experiment design for operations model simulation”. In: *Electric Power Systems Research* 189 (2020), p. 106680.
- [58] Jason Leadbetter and Lukas G. Swan. “Selection of battery technology to support grid-integrated renewable electricity”. In: *Journal of Power Sources* 216 (2012), pp. 376–386. ISSN: 03787753. DOI: 10.1016/j.jpowsour.2012.05.081. URL: <http://www.sciencedirect.com/science/article/pii/S0378775312009500>.
- [59] Jonathan T Lee and Duncan S Callaway. “The cost of reliability in decentralized solar power systems in sub-Saharan Africa”. In: *Nature Energy* 3.11 (2018), pp. 960–968.
- [60] Jonathan T Lee et al. “Non-Intrusive Load Management Under Forecast Uncertainty in Energy Constrained Microgrids”. In: *Electric Power Systems Research* 190 (2021), p. 106632.

- [61] Kenneth Lee et al. “Electrification for “Under Grid” households in Rural Kenya”. In: *Development Engineering* 1 (June 2016), pp. 26–35. DOI: 10.1016/j.deveng.2015.12.001. URL: <http://linkinghub.elsevier.com/retrieve/pii/S235272851530035X>.
- [62] Mitchell Lee, Daniel Soto, and Vijay Modi. “Cost versus reliability sizing strategy for isolated photovoltaic micro-grids in the developing world”. In: *Renewable Energy* 69 (Sept. 2014), pp. 16–24. ISSN: 09601481. DOI: 10.1016/j.renene.2014.03.019. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0960148114001633>.
- [63] Gijs van Leeuwen et al. “An integrated blockchain-based energy management platform with bilateral trading for microgrid communities”. In: *Applied Energy* 263 (2020), p. 114613.
- [64] Todd Levin and Valerie M. Thomas. “Can developing countries leapfrog the centralized electrification paradigm?” In: *Energy for Sustainable Development* 31 (2016), pp. 97–107. ISSN: 09730826. DOI: 10.1016/j.esd.2015.12.005.
- [65] Todd Levin and Valerie M. Thomas. “Least-cost network evaluation of centralized and decentralized contributions to global electrification”. In: *Energy Policy* 41 (Feb. 2012), pp. 286–302. ISSN: 03014215. DOI: 10.1016/j.enpol.2011.10.048. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0301421511008639>.
- [66] Chao Long et al. “Peer-to-peer energy sharing through a two-stage aggregated battery control in a community Microgrid”. In: *Applied energy* 226 (2018), pp. 261–276.
- [67] Chao Long et al. “Peer-to-peer energy trading in a community microgrid”. In: *2017 IEEE Power & Energy Society General Meeting*. IEEE. 2017, pp. 1–5.
- [68] Alexandra Lüth et al. “Local electricity market designs for peer-to-peer trading: The role of battery flexibility”. In: *Applied energy* 229 (2018), pp. 1233–1243.
- [69] Cheng Lyu, Youwei Jia, and Zhao Xu. “Fully decentralized peer-to-peer energy sharing framework for smart buildings with local battery system and aggregated electric vehicles”. In: *Applied Energy* 299 (2021), p. 117243.
- [70] Minyue Ma, Lingling Fan, and Zhixin Miao. “Consensus ADMM and Proximal ADMM for economic dispatch and AC OPF with SOCP relaxation”. In: *2016 North American power symposium (NAPS)*. IEEE. 2016, pp. 1–6.
- [71] Sindri Magnússon, Pradeep Chathuranga Weeraddana, and Carlo Fischione. “A distributed approach for the optimal power-flow problem based on ADMM and sequential convex approximations”. In: *IEEE Transactions on Control of Network Systems* 2.3 (2015), pp. 238–253.
- [72] Bodhisattwa P Majumder et al. “An efficient iterative double auction for energy trading in microgrids”. In: *2014 IEEE Symposium on Computational Intelligence Applications in Smart Grid (CIASG)*. IEEE. 2014, pp. 1–7.

- [73] Ali Makhdoumi and Asuman Ozdaglar. “Convergence rate of distributed ADMM over networks”. In: *IEEE Transactions on Automatic Control* 62.10 (2017), pp. 5082–5095.
- [74] Kostas Margellos and Shmuel Oren. “Capacity Controlled Demand Side Management: A Stochastic Pricing Analysis”. In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016), pp. 706–717. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2015.2406813.
- [75] Bill McKibben. *The Race to Solar Power Africa*. June 2017. URL: <http://www.newyorker.com/magazine/2017/06/26/the-race-to-solar-power-africa>.
- [76] Esther Mengelkamp et al. “Designing microgrid energy markets: A case study: The Brooklyn Microgrid”. In: *Applied Energy* 210 (2018), pp. 870–880.
- [77] Asami Miketa and Bruno Merven. *Southern African Power Pool: Planning and Prospects for Renewable Energy*. Tech. rep. International Renewable Energy Agency, 2013. URL: <http://irena.org/DocumentDownloads/Publications/SAPP.pdf>.
- [78] Asami Miketa and Bruno Merven. *West African Power Pool: Planning and Prospects for Renewable Energy*. Tech. rep. International Renewable Energy Agency, 2013. URL: <https://www.irena.org/documentdownloads/publications/wapp.pdf>.
- [79] Asami Miketa and Nawfal Saadi. *Africa Power Sector: Planning and Prospects for Renewable Energy*. Tech. rep. International Renewable Energy Agency, 2015. URL: http://www.irena.org/DocumentDownloads/Publications/IRENA_Africa_Power_Sector_synthesis_2015.pdf.
- [80] Daniel K Molzahn et al. “A survey of distributed optimization and control algorithms for electric power systems”. In: *IEEE Transactions on Smart Grid* 8.6 (2017), pp. 2941–2962.
- [81] Fabio Moret et al. “Negotiation algorithms for peer-to-peer electricity markets: Computational properties”. In: *2018 power systems computation conference (PSCC)*. IEEE, 2018, pp. 1–7.
- [82] Eric Münsing, Jonathan Mather, and Scott Moura. “Blockchains for decentralized optimization of energy resources in microgrid networks”. In: *2017 IEEE conference on control technology and applications (CCTA)*. IEEE, 2017, pp. 2164–2171.
- [83] Dinh Hoa Nguyen. “Optimal solution analysis and decentralized mechanisms for peer-to-peer energy markets”. In: *IEEE Transactions on Power Systems* 36.2 (2020), pp. 1470–1481.
- [84] R. Nyakudya et al. “A decision support tool for rural electrification grid design”. In: *2013 IEEE International Conference on Industrial Technology (ICIT)*. IEEE, Feb. 2013, pp. 1443–1449. ISBN: 978-1-4673-4569-9. DOI: 10.1109/ICIT.2013.6505884. URL: <http://ieeexplore.ieee.org/document/6505884/>.
- [85] Eunsung Oh and Sung-Yong Son. “Peer-to-peer energy transaction mechanisms considering fairness in smart energy communities”. In: *IEEE Access* 8 (2020), pp. 216055–216068.

- [86] Daniel E. Olivares et al. “Stochastic-Predictive Energy Management System for Isolated Microgrids”. In: *IEEE Transactions on Smart Grid* 6.6 (Nov. 2015), pp. 2681–2693. ISSN: 1949-3053. DOI: 10.1109/TSG.2015.2469631.
- [87] Elinor Ostrom et al. *Rules, games, and common-pool resources*. University of Michigan Press, 1994.
- [88] A. Papavasiliou, S. S. Oren, and R. P. O’Neill. “Reserve Requirements for Wind Power Integration: A Scenario-Based Stochastic Programming Framework”. In: *IEEE Transactions on Power Systems* 26.4 (Nov. 2011), pp. 2197–2206. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2011.2121095.
- [89] Neal Parikh and Stephen Boyd. “Proximal algorithms”. In: *Foundations and Trends in optimization* 1.3 (2014), pp. 127–239.
- [90] Alessandra Parisio and Luigi Glielmo. “Stochastic Model Predictive Control for economic/environmental operation management of microgrids”. In: *2013 European Control Conference (ECC)*. IEEE, July 2013, pp. 2014–2019. ISBN: 978-3-033-03962-9. DOI: 10.23919/ECC.2013.6669807.
- [91] Alessandra Parisio, Evangelos Rikos, and Luigi Glielmo. “A Model Predictive Control Approach to Microgrid Operation Optimization”. In: *IEEE Transactions on Control Systems Technology* 22.5 (Sept. 2014), pp. 1813–1827. ISSN: 1063-6536. DOI: 10.1109/TCST.2013.2295737.
- [92] Lily Parshall et al. “National electricity planning in settings with low pre-existing grid coverage: Development of a spatial model and case study of Kenya”. In: *Energy Policy* 37.6 (June 2009), pp. 2395–2410. ISSN: 03014215. DOI: 10.1016/j.enpol.2009.01.021. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0301421509000561>.
- [93] Amrit Paudel and Hoay Beng Gooi. “Pricing in peer-to-peer energy trading using distributed optimization approach”. In: *2019 IEEE Power & Energy Society General Meeting (PESGM)*. IEEE, 2019, pp. 1–5.
- [94] Amrit Paudel et al. “Peer-to-peer energy trading in a prosumer-based community microgrid: A game-theoretic model”. In: *IEEE Transactions on Industrial Electronics* 66.8 (2018), pp. 6087–6097.
- [95] Pecan Street, Inc. *Pecan Street Dataport*. 2020. URL: <https://www.pecanstreet.org/dataport/>.
- [96] Qiuyu Peng and Steven H Low. “Distributed optimal power flow algorithm for radial networks, I: Balanced single phase case”. In: *IEEE Transactions on Smart Grid* 9.1 (2016), pp. 111–121.
- [97] Morteza Rahimiyan and Luis Baringo. “Real-time energy management of a smart virtual power plant”. In: *IET Generation, Transmission & Distribution* 13.11 (June 2019), pp. 2015–2023. ISSN: 1751-8687. DOI: 10.1049/iet-gtd.2018.5637.

- [98] REN21. *Renewables 2017 Global Status Report*. Tech. rep. REN21, 2017. URL: http://www.ren21.net/wp-content/uploads/2017/06/170607_GSR_2017_Full_Report.pdf.
- [99] Jesse C Ribot and Nancy Lee Peluso. “A theory of access”. In: *Rural sociology* 68.2 (2003), pp. 153–181.
- [100] Joan Rocabert et al. “Control of power converters in AC microgrids”. In: *IEEE transactions on power electronics* 27.11 (2012), pp. 4734–4749.
- [101] P.A. Ruiz et al. “Uncertainty Management in the Unit Commitment Problem”. In: *IEEE Transactions on Power Systems* 24.2 (May 2009), pp. 642–651. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2008.2012180.
- [102] Tiago Sousa et al. “Peer-to-peer and community-based markets: A comprehensive review”. In: *Renewable and Sustainable Energy Reviews* 104 (2019), pp. 367–378.
- [103] Goran Strbac. “Demand side management: Benefits and challenges”. In: *Energy Policy* 36.12 (Dec. 2008), pp. 4419–4426. DOI: 10.1016/J.ENPOL.2008.09.030.
- [104] Goran Strbac et al. “Decarbonization of Electricity Systems in Europe: Market Design Challenges”. In: *IEEE Power and Energy Magazine* 19.1 (2021), pp. 53–63.
- [105] Wencong Su, Jianhui Wang, and Jaehyung Roh. “Stochastic Energy Scheduling in Microgrids With Intermittent Renewable Energy Resources”. In: *IEEE Transactions on Smart Grid* 5.4 (July 2014), pp. 1876–1883. ISSN: 1949-3053. DOI: 10.1109/TSG.2013.2280645.
- [106] Sándor Szabó et al. “Energy solutions in rural Africa: mapping electrification costs of distributed solar and diesel generation versus grid extension”. In: *Environmental Research Letters* 6.3 (July 2011), p. 034002. ISSN: 1748-9326. DOI: 10.1088/1748-9326/6/3/034002. URL: <http://stacks.iop.org/1748-9326/6/i=3/a=034002?key=crossref.89590e74fed5b2bdbff442f147cb9d62>.
- [107] Sándor Szabó et al. “Identification of advantageous electricity generation options in sub-Saharan Africa integrating existing resources”. In: *Nature Energy* 1.10 (2016). ISSN: 2058-7546. DOI: 10.1038/nenergy.2016.140. URL: <https://www.nature.com/articles/nenergy2016140.pdf%20http://www.nature.com/articles/nenergy2016140>.
- [108] Sándor Szabó et al. *Sustainable energy planning: Leapfrogging the energy poverty gap in Africa*. 2013. DOI: 10.1016/j.rser.2013.08.044.
- [109] Reza Takapoui et al. “A simple effective heuristic for embedded mixed-integer quadratic programming”. In: *International journal of control* 93.1 (2020), pp. 2–12.
- [110] S. Takriti, J.R. Birge, and E. Long. “A stochastic model for the unit commitment problem”. In: *IEEE Transactions on Power Systems* 11.3 (1996), pp. 1497–1508. ISSN: 08858950. DOI: 10.1109/59.535691.

- [111] Jay Taneja. “Measuring Electricity Reliability in Kenya”. 2017. URL: <http://blogs.umass.edu/jtaneja/files/2017/05/outages.pdf>.
- [112] The World Bank. *Access to electricity (% of population)*. 2014. URL: <https://data.worldbank.org/indicator/EG.ELC.ACCS.ZS>.
- [113] The World Bank. *Enterprise Surveys*. 2017. URL: <http://www.enterprisesurveys.org/data/exploretopics/infrastructure#sub-saharan-africa>.
- [114] Chris Trimble et al. *Financial Viability of Electricity Sectors in sub-Saharan Africa. Quasi-fiscal Deficits and Hidden Costs*. Tech. rep. 2016. URL: <https://data.worldbank.org/data-catalog/affordable-viable-power-for-africa>.
- [115] Wayes Tushar et al. “Peer-to-peer energy systems for connected communities: A review of recent advances and emerging challenges”. In: *Applied Energy* 282 (2021), p. 116131.
- [116] Wayes Tushar et al. “Peer-to-peer trading in electricity networks: An overview”. In: *IEEE Transactions on Smart Grid* 11.4 (2020), pp. 3185–3200.
- [117] Md Habib Ullah and Jae-Do Park. “Peer-to-Peer Energy Trading in Transactive Markets Considering Physical Network Constraints”. In: *IEEE Transactions on Smart Grid* (2021).
- [118] Khalid Umer et al. “A novel communication efficient peer-to-peer energy trading scheme for enhanced privacy in microgrids”. In: *Applied Energy* 296 (2021), p. 117075.
- [119] United Nations. *Transforming our world: the 2030 Agenda for Sustainable Development*. 2015. DOI: 10.1007/s13398-014-0173-7.2.
- [120] United States Energy Information Administration. *Annual Electric Power Industry Report (EIA-861 data file)*. Tech. rep. 2015. URL: <https://www.eia.gov/electricity/data/eia861/>.
- [121] United States of America National Aeronautics and Space Administration. *Surface meteorology and Solar Energy*. URL: <https://eosweb.larc.nasa.gov/cgi-bin/sse/sse.cgi>.
- [122] Mostafa Vahedipour-Dahraie et al. “Stochastic security and risk-constrained scheduling for an autonomous microgrid with demand response and renewable energy resources”. In: *IET Renewable Power Generation* 11.14 (Dec. 2017), pp. 1812–1821. ISSN: 1752-1416. DOI: 10.1049/iet-rpg.2017.0168.
- [123] Mengdi Wang and Dimitri P Bertsekas. “Incremental constraint projection-proximal methods for nonsmooth convex optimization”. In: *SIAM J. Optim.(to appear)* (2013).
- [124] Yang Wang et al. “Shadow price based co-ordination methods of microgrids and battery swapping stations”. In: *Applied Energy* 253 (2019), p. 113510.
- [125] Ermin Wei and Asuman Ozdaglar. “On the $o(1/k)$ convergence of asynchronous distributed alternating direction method of multipliers”. In: *2013 IEEE Global Conference on Signal and Information Processing*. IEEE. 2013, pp. 551–554.

- [126] Grace C. Wu et al. *Renewable Energy Zones for the Africa Clean Energy Corridor*. Tech. rep. International Renewable Energy Agency and Lawrence Berkeley National Laboratory, 2015. URL: http://www.irena.org/DocumentDownloads/Publications/IRENA_LBNL_Africa-REZ-CEC_2015.pdf.
- [127] Lei Wu, Mohammad Shahidehpour, and Tao Li. “Stochastic Security-Constrained Unit Commitment”. In: *IEEE Transactions on Power Systems* 22.2 (May 2007), pp. 800–811. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2007.894843.
- [128] Daniel Zelazo and Mehran Mesbahi. “Edge agreement: Graph-theoretic performance bounds and passivity analysis”. In: *IEEE Transactions on Automatic Control* 56.3 (2010), pp. 544–555.
- [129] Marianne Zeyringer et al. “Analyzing grid extension and stand-alone photovoltaic systems for the cost-effective electrification of Kenya”. In: *Energy for Sustainable Development* 25 (2015), pp. 75–86. ISSN: 09730826. DOI: 10.1016/j.esd.2015.01.003.
- [130] Min Zhang et al. “Energy Trading with Demand Response in a Community-based P2P Energy Market”. In: *2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2019, pp. 1–6.
- [131] Ruiliang Zhang and James Kwok. “Asynchronous distributed ADMM for consensus optimization”. In: *International conference on machine learning*. PMLR, 2014, pp. 1701–1709.