Coordination of distributed energy resources for a reliable, resilient, and affordable decarbonized grid

by

Vineet Jagadeesan Nair

B.S. Mechanical Engineering and Economics, UC Berkeley, 2018
MPhil Energy Technologies, University of Cambridge, 2019
S.M. Computational Science and Engineering, MIT, 2021

Submitted to the Department of Mechanical Engineering and
the Center for Computational Science and Engineering
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN MECHANICAL ENGINEERING AND COMPUTATION

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2025

Authored by:	Vineet Jagadeesan Nair Department of Mechanical Engineering May 16, 2025
Certified by:	Anuradha Annaswamy Senior Research Scientist Department of Mechanical Engineering, Thesis Supervisor
Accepted by:	Youssef Marzouk Professor of Aeronautics and Astronautics Co-Director, Center for Computational Science and Engineering
Accepted by:	Nicolas Hadjiconstantinou Professor of Mechanical Engineering and Graduate Officer Co-Director, Center for Computational Science and Engineering

THESIS COMMITTEE

Thesis Supervisor

Anuradha Annaswamy
Senior Research Scientist
Department of Mechanical Engineering, MIT

Thesis Readers

Kamal Youcef-Toumi
Professor
Department of Mechanical Engineering, MIT

Andy Sun
Professor
Operations Research and Sloan School of Management, MIT

Coordination of distributed energy resources for a reliable, resilient, and affordable decarbonized grid

by

Vineet Jagadeesan Nair

Submitted to the Department of Mechanical Engineering and
the Center for Computational Science and Engineering
on May 16, 2025 in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN MECHANICAL ENGINEERING AND COMPUTATION

ABSTRACT

Rapid decarbonization of the power grid is essential to meet climate goals by reducing emissions and enabling sustainable electrification of sectors like transport and heating. This requires shifting from centralized fossil-fuel generation to variable renewables like wind and solar. The grid must also adapt to a growing number of small-scale, distributed energy resources (DERs) at the edge, such as rooftop solar, batteries, electric vehicles, and heat pumps. This thesis focuses on modeling, optimizing, and coordinating DERs to enable a flexible, resilient, and affordable grid.

First, it proposes a novel hierarchical local electricity market for low and medium-voltage distribution grids. This structure enables DER participation through decentralized and distributed optimization, respecting grid physics while preserving privacy and scalability. The market is applicable to both balanced and unbalanced radial grids using two different convex relaxations and power flow models. Grid services are also priced based on duality theory. Numerical simulations show improved dispatch efficiency, reliability, voltage regulation, and lower retail electricity rates. Second, the thesis applies game theory and mechanism design to extract flexibility from autonomous, strategic DER owners. A repeated Stackelberg game with incomplete information and intertemporal constraints yields equilibrium pricing with closed-form solutions.

Third, a distributed decision-making framework is developed to coordinate DERs for grid resilience. It mitigates cyber-physical attacks and outages, ranging from 5 to 40% of peak load, using local flexibility and grid reconfiguration, extensively validated through both software and hardware-in-the-loop simulations. Finally, the thesis addresses DER hosting capacity. New algorithms are developed that co-optimize the siting and sizing of diverse DERs under uncertainty using Monte Carlo sampling, stochastic programming, and k-means clustering for scenario reduction. Results show that intelligent DER coordination can defer grid infrastructure upgrades and support greater renewable integration and electrified demand growth. Together, these contributions provide analytical and simulation tools to improve the planning and real-time operation of future distributed, low-carbon power grids.

Thesis supervisor: Anuradha Annaswamy
Title: Senior Research Scientist
Department of Mechanical Engineering

Acknowledgments

I would like to first sincerely thank my advisor Dr. Anuradha Annaswamy, for her guidance, mentorship, and support throughout my time at MIT, both during my SM and PhD. I have learned a great deal from her and grown significantly as a researcher and academic as a result. I would also like to thank my thesis committee, Prof. Kamal Youcef-Toumi and Prof. Andy Sun, for their feedback and guidance over the past three years. They have consistently challenged and encouraged me to dive deeper into difficult research questions, and this has strengthened the novelty and strength of my research contributions.

I have had the pleasure of collaborating with many wonderful people both within MIT and beyond. In particular, I would like to thank Yue Guan for his mentorship during my SM. I would also like to thank my other MIT mentors Venkatesh Venkataramanan, Priyank Srivastava, Rabab Haider, and Morteza Vahid Ghavidel. I have enjoyed collaborating on side projects with Thomas Lee and Layla Araiinejad at MIT.

I've gotten the chance to work with many others outside MIT, including Prof. Vincent Poor at Princeton and Prof. Anurag Srivastava and Partha Sarker at West Virginia University, Laurentiu Marinovici at Pacific Northwest National Lab, and Jun Zha at LTDES. I've also enjoyed multiple international collaborations with Lucas Pereira at Tecnico Lisboa in Portugal, Jesús Rodríguez-Molina at Universidad Politécnica de Madrid in Spain, and Giulio Ferro at the University of Genoa in Italy. I would also like to thank several mentors and collaborators during my research internships - Amit Chakraborty and Biswadeep Dey at Siemens; Rob Hovsapian at the National Renewable Energy Lab; Hamed Khalilinia, Amir Saadat, and Evgeny Tsydenov at Google X; and Zachary Caruso and Mansur Mohammed at Avangrid.

One of the most rewarding parts of my PhD experience was mentoring and advising several amazing students on their research projects. I would like to thank MIT undergraduates Danielle Knutson and Jose Vargas, as well as Luca Hartmann and Peer Brigger, who were both visiting master's students from ETH Zurich.

Lastly, I want to sincerely thank my family and friends who have loved and supported me in this journey. The past six years at MIT have been some of the most challenging years of my life, and I would not have been able to get here without this strong support system. I especially want to thank my parents, Acha and Amma, my brother Vivek and Visakh, and my sister-in-law Roshni - you all are my rock. And finally, a huge thank you to my friends for always being there for me. From the many new friends I made while at MIT, to older friends from other chapters of my life in Bahrain, Dubai, India, Berkeley, and Cambridge - this degree is as much yours as it is mine.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and motivation

Rapid decarbonization of the power grid is crucial to meet climate change mitigation goals by reducing emissions from the electric sector and sustainably electrifying other sectors like transportation, heating, and industry. This will entail a shift away from centralized fossil fuel-based power generation towards renewables like wind and solar. Distributed Energy Resources (DERs) such as renewables (wind, rooftop solar), battery storage, electric vehicles (EVs), and flexible loads (such as heat pumps and data centers), will play a key role in this transition away from fossil-fuel-powered generation [1]. DER installations are expected to grow exponentially over the next decade, especially with increasingly favorable policies and regulations at the federal, state, and local levels, to meet the US goal of 100% carbon carbon-neutral electricity by 2035 and reaching a net-zero emissions economy-wide by 2050. However, current market structures are not well-suited to high levels of DER integration. Traditionally, electricity markets have been highly centralized and set up for large centralized generation sources like thermal power plants. However, generation is now increasingly distributed to the grid edge, resulting in bidirectional power exchanges and the rising prevalence of 'prosumers', i.e., consumers who also produce some electricity locally. The rise in local generation has also been accompanied by energy storage devices like batteries as well as highly controllable and flexible loads, enabled by the proliferation of devices like smart meters, smart inverters, smart chargers, smart thermostats, etc.

One of the main results of all these changes on both the supply and demand side is increased flexibility in net power injections in the distribution grid. However, there are currently inadequate mechanisms in place (if any) to leverage this flexibility in order to optimize the grid operation. Existing consumer choice programs like net energy metering (NEM) and net energy billing (NEB) generally compensate DERs for their grid services through either direct incentives or feed-in tariffs. However, these policies are inefficient since they don't price the fine-grain locational and temporal variations in the services that DERs are capable of providing, and thus are unable to meet network-level objectives under high DER penetration [2]. As the number of DERs and prosumers (energy consumers who also produce some electricity) increases, more structure is warranted to coordinate them. New operational entities, such as a distribution system operator (DSO) that oversees the

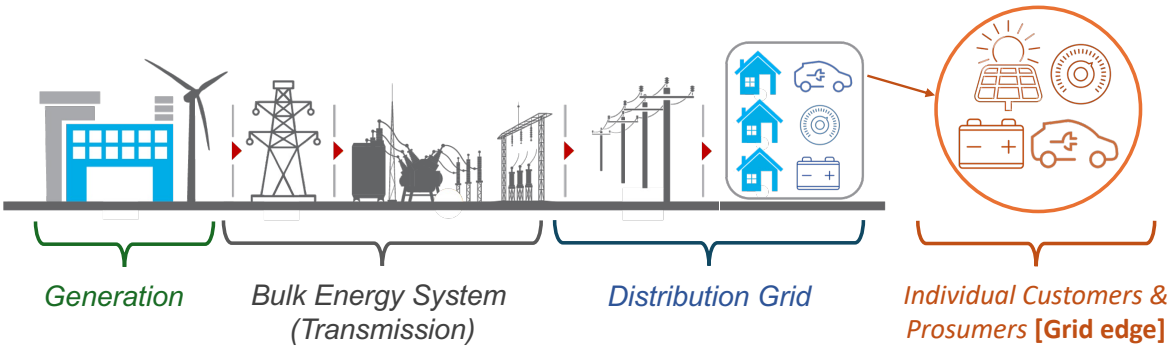

Figure 1.1: Simplified overview of the power grid.

distribution grid [3], and new types of tariffs such as distribution-level locational marginal prices (d-LMP), may be needed [4]. Furthermore, significant upgrades and retrofits will need to be made to both transmission and distribution networks in order to handle new issues like reverse power flows and increased uncertainty from renewables [5, 6].

Fig. 1.1 shows the overall structure of the power grid. Power is first produced by large centralized power generation sources like coal and natural gas plants or wind and solar farms. This is then transported over long distances by high-voltage transmission lines and other related infrastructure. This forms the large-scale transmission grid or bulk energy system. As we get closer to the end customers, the voltage is stepped down in the medium and low-voltage distribution grid. Finally, the power is distributed to individual customers and prosumers. Prosumers are electricity consumers who also produce some of their own power. My thesis specifically focuses on these last two parts of the grid: the distribution network and the grid edge. Both are transforming rapidly due to the increasing penetration of DERs. The grid edge is also seeing an increasing number of stakeholders, such as prosumers, DER aggregators, and virtual power plants. The grid edge of the past primarily consisted of passive demand. This paradigm is completely changing with increasing amounts of generation and storage resources being added at the grid edge, along with highly responsive, flexible, and controllable loads. Thus, the grid edge is becoming much more intelligent, complex, and capable. The central idea behind this thesis is how to best utilize all this flexibility present in the grid to aid in both grid planning and operation. *How can we use intelligent decision-making to enable rapid, reliable, and affordable integration of Distributed Energy Resources (DERs)?*

This PhD thesis is focused on modeling power distribution grids and electricity markets under a high penetration of DERs. In particular, we design new decision-making tools and frameworks to effectively coordinate and aggregate DERs in order to improve the reliability, efficiency, affordability, and resilience of decarbonized grids. While most of my applications have focused on distribution grids, the methods can be easily generalized to larger-scale transmission grids as well. Furthermore, the tools are flexible and can be applied to future grids with high penetrations of renewables and DERs as well as to current grids with low to medium penetrations.

1.2 Related Work

Several related works in the literature have considered the impacts of increasing DER penetration on distribution grids and how these challenges can be addressed. One active area is in the development of new optimization algorithms for DER coordination and market clearing. Particular emphasis has been on distributed and decentralized algorithms that are more scalable and computationally tractable, while also preserving privacy and enhancing security [7]. This thesis will leverage such approaches and apply them to local electricity markets. A specific example is a proximal atomic coordination algorithm - a dual ascent method that has been applied to clear our markets at the primary feeder level [8]. This thesis will build upon these tools and explore extensions aimed at further improving the performance of these algorithms, e.g., accelerated convergence and enhanced privacy [9, 10].

Another group of papers focuses on market-based mechanisms for DER coordination and aggregation. Several different proposals have been made for restructuring electricity markets, ranging from retail markets [9, 11], microgrid management [12, 13], and local flexibility markets [14] to different aggregator models [15]. Recent government regulations like FERC Order 2222 have also opened up wholesale energy markets to DERs. However, direct participation of DERs at the wholesale level may introduce 'tier-bypassing' issues [16] since intermediate grid physics constraints may be violated, which can potentially lead to instabilities. Furthermore, most of the papers proposing distribution-level markets consider aggregations of DERs on a larger scale and do not disaggregate the tariffs or schedules down to the level of individual homes, buildings, and DERs at the primary and secondary feeder levels. This thesis aims to address the problem of DER coordination across the entire distribution grid, from the substation to the primary feeder and, eventually, secondary feeder nodes. In addition to computing d-LMPs, these will be further broken down into local retail tariffs to better capture temporal and spatial variations with higher resolution. Thus, our work aims to leverage these differentiated price signals to influence prosumer behavior, similar to a transactive control framework [17]. This places our work on LEMs in the broader context of transactive energy [18, 19].

In addition to optimization theory, other solution methods inspired by game theory have also been proposed for energy trading [20], integrating demand response [21], designing optimal bids [22], and other applications in smart grids [23]. There is also a wealth of literature in applying mechanism design to model electricity market auctions, but these have largely focused on day-ahead or real-time markets at the wholesale transmission level [24, 25]. This thesis aims to build upon these existing works and design market mechanisms and bidding methods that account for the strategic behavior or even potential market manipulation of prosumer agents in distribution grids. Such issues are arguably even more critical here given the smaller scale of these local markets and the fewer participants involved, which can give individual agents more market power. Finally, another class of works looks at applications of machine learning and artificial intelligence in power systems and electricity markets [26]. Examples range from deep learning for approximating solutions to the non-convex AC optimal power flow (OPF) problem [27, 28] and graph neural networks for predicting real-time LMPs [29], to federated learning for load forecasting [30] and reinforcement learning for grid services such as voltage control [31].

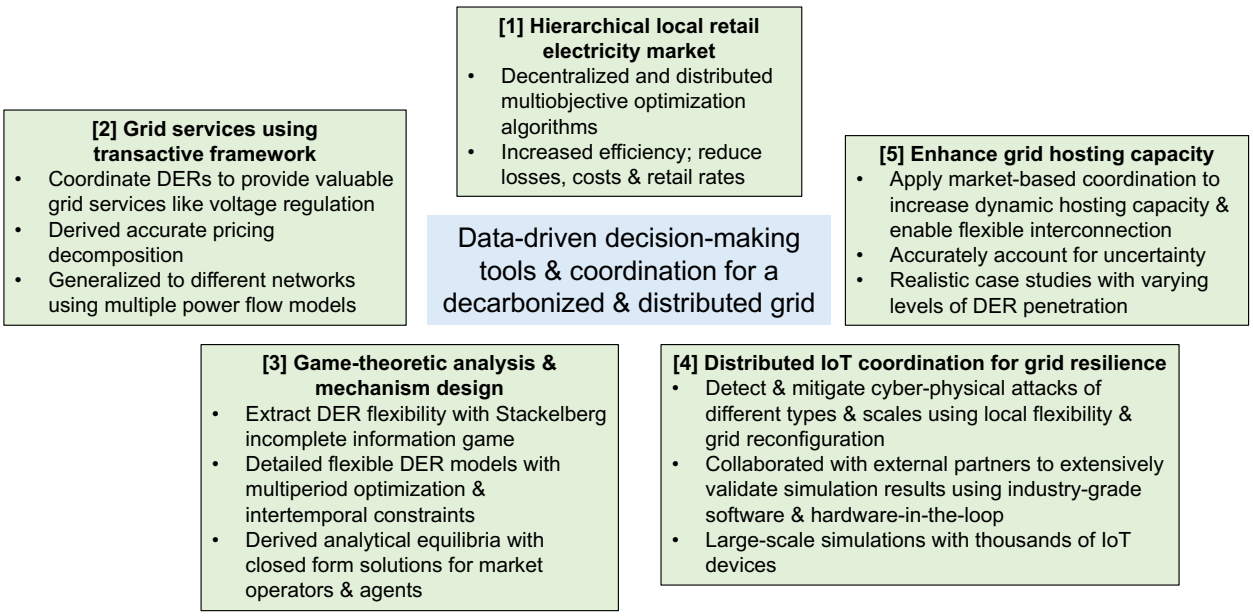

Figure 1.2: Summary of five chapters of this thesis.

1.3 Contributions of this thesis

In this thesis, I propose solutions to mainly address the following four important challenges with future power grids:

- *Problem*: DERs are generally autonomous and independently owned. Thus, utilities or grid operators can't directly control them.
  Solution: Use new markets and price signals to incentivize agents indirectly instead.

- *Problem*: Future distributed and decentralized grids will have millions of DERs, making it difficult to coordinate resources and manage the grid.
  Solution: Develop hierarchical market designs and scalable optimization tools.

- *Problem*: Renewable intermittency, uncertainty, and variability will create reliability and stability issues, along with inefficiencies.
  Solution: Use a transactive framework to provide valuable grid services and resolve these issues.

- *Problem*: Rapid growth of renewables, storage, and electrified demand will severely stress the grid, thus hindering decarbonization and raising costs.
  Solution: Coordinate DERs to dynamically increase grid capacity.

A common thread throughout this thesis is envisioning how we can redesign future electricity markets and create the necessary coordination structures needed to manage millions of independently owned, autonomous DERs. Fig. 1.3 is a schematic of how the electricity market operates today, where small-scale DERs and their agents have little to no participation in, or interaction with, the market. However, our proposal, as shown in Fig. 1.4,

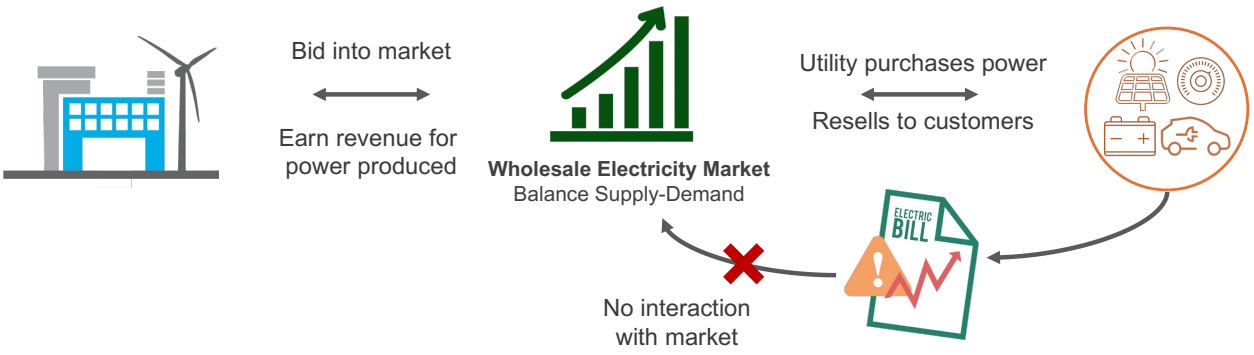

Figure 1.3: How electricity markets are structured today.

Figure 1.4: Our vision for future electricity markets.

redesigns this by introducing new local retail market structures specifically designed for the distribution grid and DERs at the grid edge.

The first chapter of this thesis in Chapter 2 develops a novel, hierarchical local electricity market (LEM) structure for low and medium-voltage distribution grids that enables seamless market participation for small-scale resources. By introducing agents and operators at different voltage levels in the network, the market can optimally coordinate and aggregate DERs to provide flexibility to the grid. The market utilizes decentralized optimization at the lower level and a distributed optimization method at the upper level. Such distributed algorithms help preserve data privacy and reduce computational burden. The hierarchical structure also allows it to accurately account for all grid physics constraints while being scalable for arbitrarily large networks with many nodes. Numerical simulations show that this results in more efficient dispatch of DERs and improves affordability by lowering retail rates for end users. In the second chapter in Chapter 3, the market structure is generalized to both balanced and unbalanced radial grids by leveraging two types of power flow models and convex relaxations to improve tractability. This framework is then applied to provide valuable grid services like voltage regulation, with simulations confirming that it improves power quality and reliability. Accurate spatially and temporally varying prices for such services are also derived using duality theory. Thus, in addition to optimizing for grid objectives like power quality, the algorithms and frameworks developed also seek to appropriately compensate DERs and prosumers for services provided.

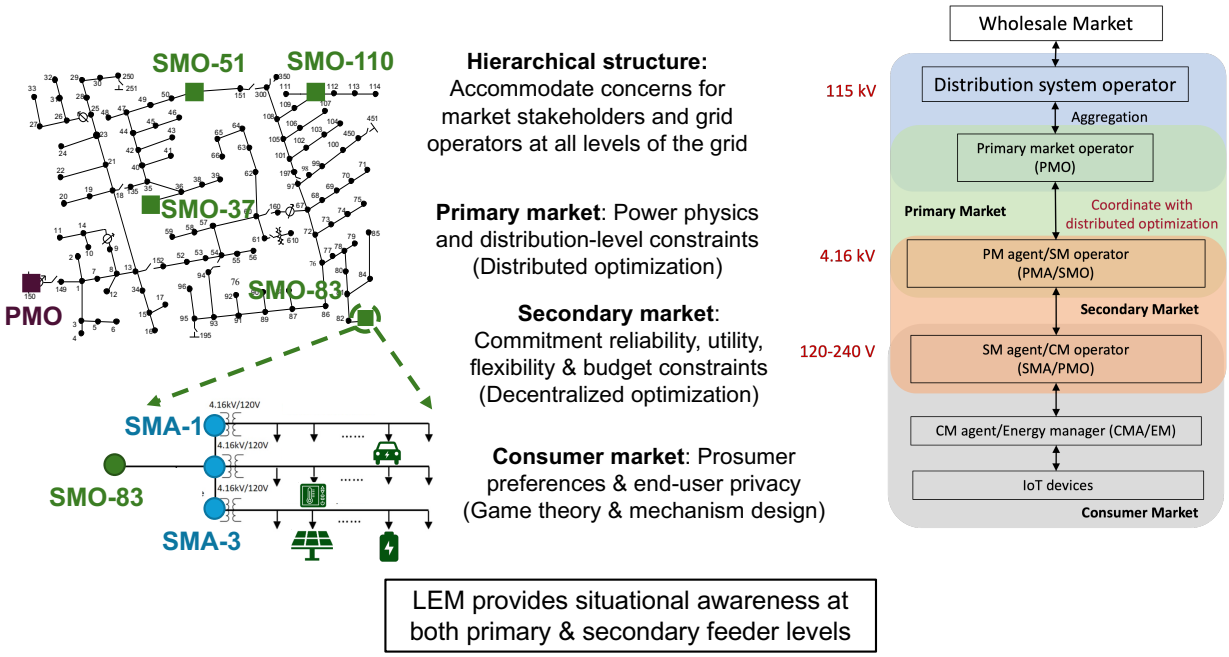

Figure 1.5: Summary of tools and methodologies for different components of our market and coordination framework.

The third chapter of this thesis in Chapter 4 applies game theory and mechanism design to develop a method to optimally extract flexibility from independently owned, autonomous agents. This is based on a repeated Stackelberg game of incomplete information that also incorporates physical device-level and intertemporal constraints of DERs. By accounting for strategic behavior, a pricing method is developed that results in an equilibrium among market participants with closed-form solutions.

The fourth chapter of this thesis in Chapter 5 builds on the market to develop a distributed decision-making framework that coordinates large numbers of trustable Internet-of-Things-enabled DERs to enhance the resilience of the grid against cyber-physical attacks and outages. Several different types of attacks on both load and generation are studied, with attack magnitudes ranging from 5 to 40% of peak load, including both grid-connected and islanded cases. Results confirm that the method can successfully mitigate all the attacks considered through a combination of locally available flexible assets and reconfiguration of the grid topology. The simulation results are also validated extensively on both software and hardware.

The final chapter of this thesis in Chapter 6 shows how intelligent and flexible DER coordination can significantly increase the hosting capacity of distribution grids to accommodate more DERs, thus supporting the rapid growth of both renewable generation and electrified demand. Three novel algorithms for analyzing hosting capacity are developed, which allow planners to co-optimize the siting and sizing of multiple types of DERs simultaneously. Monte Carlo sampling and two-stage stochastic programming are used to account for uncertainty, and scenario reduction based on k-means clustering is used to increase simulation speed. Reducing the need for infrastructure upgrades and retrofits also lowers the cost of DER integration.

Fig. 1.5 provides a holistic overview of the different components of our LEM and coordination framework. It also summarizes the diverse tools and methodologies used. These include novel approaches and algorithms we developed from scratch as well as existing methods that were adapted for our specific applications. Note that the primary and secondary markets primarily relate to Chapter 2 and Chapter 3, while the consumer market is the focus of Chapter 4. Chapter 5 leverages both the primary and consumer markets while Chapter 6 uses the primary market as one option to enable the DER coordination. Overall, this thesis provides a set of general tools for improved planning and operation of future, zero-carbon power grids.

Chapter 2

A new hierarchical local retail electricity market for the distribution grid

With the increasing penetration of DERs in the distribution system, it is critical to design market structures that enable the smooth integration of DERs. In this chapter, we propose a hierarchical local electricity market (LEM) structure with a secondary market (SM) at the lower level representing secondary feeders and a primary market (PM) at the upper level, representing primary feeders, in order to effectively use DERs to increase grid efficiency and resilience. The lower level SM enforces budget, power balance, and flexibility constraints and accounts for costs related to consumers, such as their disutility, flexibility limits, and commitment reliability, while the upper level PM enforces grid physics constraints such as power balance and capacity limits, and also minimizes line losses. The hierarchical LEM is extensively evaluated using a modified IEEE-123 bus with high DER penetration, with each primary feeder consisting of at least three secondary feeders. Data from a GridLAB-D model is used to emulate realistic power injections and load profiles over the course of 24 hours. The performance of the LEM is illustrated by delineating the family of power-injection profiles across the primary and secondary feeders, as well as corresponding local electricity tariffs that vary across the distribution grid. Through numerical simulations, the hierarchical LEM is shown to improve the efficiency of the market in terms of lowering overall costs, including both the distribution locational marginal prices (d-LMPs) as well as retail tariffs paid by customers. Together, they represent an overall framework for a distribution system operator (DSO) who can provide the oversight for the entire LEM.

2.1 Introduction and Motivation

The current electrical grid, together with the electricity market, was designed for unidirectional energy flows — from large centralized producers to final consumers. However, as the power grid becomes increasingly decentralized, the electricity market structure needs to (1) incentivize consumers appropriately for investing in DERs, (2) ensure that DERs are appropriately compensated for any grid services provided, (3) coordinate with the existing market structure at the wholesale level, (4) provide a suitable rate structure to ensure that relevant costs to various stakeholders are minimized and (5) be scalable with increasing penetration of

DERs. Local electricity markets (LEMs) have the potential to empower the consumer to take control of their energy footprint, allow transactive energy trading among members of a community, improve community resilience against wider grid events, and potentially reduce energy bills. Many of the existing consumer choice programs compensate DERs for their grid services, typically through direct incentives and feed-in tariffs. However, these policies do not price the fine-grain locational and temporal variation in the services that DERs are capable of providing, and are therefore unable to meet network-level objectives under high DER penetration [2]. As the number of DERs and prosumers (energy consumers who also produce some electricity) increases, more structure is warranted to coordinate them. New operational entities, such as a distribution system operator (DSO) that oversees the LEM [3], and new types of tariffs such as distribution-level locational marginal prices (d-LMP), may be needed [4]. The DSO supervises the entire distribution grid - our proposal could be viewed as an expansion of the current responsibilities of a DSO, which comprise grid maintenance and grid reliability, to include market oversight and regulation as well. In this sense, the role of the DSO would be analogous to that of existing independent system operators for transmission grids [32].

Local energy markets have the capability to allow electricity prices to be endogenous quantities rather than being imposed exogenously. In such a marketplace, prosumers can buy and sell energy in an open marketplace, or through an operator [33]. The introduction of DERs such as rooftop solar panels and electric vehicles (EVs) has introduced significant complexity to the management of the grid. Grid operators and utilities often rely on standard load profiles derived from historical data to model home energy usage and estimate the amount of energy required to supply and balance the grid. However, the intermittent and highly variable nature of the generation from photovoltaic (PV) panels, the demand of EVs, and the needs of other DERs can cause unpredictable swings in demand. LEMs have the potential to help solve this problem for energy retailers and other grid management entities by offering flexibility services and the opportunity for new business models. LEMs also provide an attractive alternative to Federal Energy Regulatory Commission (FERC) Order 2222, as the direct participation of DERs at the wholesale level may introduce tier-bypassing [16], which may lead to potential instabilities.

In this chapter, we propose an LEM for energy transactions at the distribution level. The LEM consists of a two-tier structure. The lower level consists of DER-coordinated assets (DCAs) located at each secondary feeder bidding into a secondary level market. DCAs are entities that aggregate and coordinate the DERs within their secondary feeder in order to bid into the SM. These DERs could include renewable generation such as rooftop solar PV as well as battery storage and/or flexible loads. We note here that our market mechanism does not inherently rely on any assumptions about ownership structures, e.g., it could be possible for a single agent to coordinate DCAs across multiple primary nodes. A secondary market operator (SMO) is assumed to oversee the market operations at this level, clearing and scheduling the DCAs. At the upper level, the SMO in turn bids into the PM as an agent representing a primary feeder node. These bids are, in turn, cleared and scheduled by a primary market operator (PMO), which represents a primary feeder. The payments made by the PMO to the agents at the primary feeder nodes, i.e., the SMOs, are denoted as d-LMPs, and those that are made in turn by the SMO to the DCAs are denoted as local retail tariffs.

Both the lower and the upper level market solutions proposed here are based on an

optimization framework, with the upper level based on a distributed proximal atomic coordination (PAC) approach [8, 9], while the lower level uses decentralized optimization at each primary feeder node. While the upper level accommodates detailed power physics, including nonlinear DistFlow-based power balance, and various capacity limits on the main decision variables, the lower level accommodates accurate forecasting of generation and consumption of various DERs with finer granularity and therefore better accuracy. The lower-level market also incorporates oversight over the DERs' actual participation in the market and any unmet commitments thereof in the form of a commitment score. Suitable accommodation is made in recognizing and reflecting any vulnerabilities that may be present in the form of security breaches at the secondary feeder level. Both the PMO and SMO are proposed to be managed by a DSO. We note that there may be multiple primary feeders and thus multiple PMOs connected to a single substation. In this case, all these PMOs would be coordinated by the DSO at the substation. However, for brevity, we assume in this chapter that there is only one primary feeder (and PMO) per substation. The overall structure of this hierarchical market is illustrated in Fig. 4.1.

Overall, the incorporation of such a hierarchical market framework for DERs allows an efficient incorporation of various expanded responsibilities in a local market. A DSO needs to take on various roles, including maintaining system reliability, facilitating transactions between agents and aggregators, and enabling energy procurement, market clearing, and scheduling. Our proposed PMO-SMO structure distributes these roles between the two tiers with greater emphasis on grid physics in the upper level and addressing consumer preferences, reliable performance of DERs, and monitoring of security breaches in the lower level. Through a modified case study of an IEEE-123 bus primary feeder test case with multiple secondary feeders at each bus, we demonstrate the functioning of the hierarchical structure and show that the LEM can coordinate and aggregate local DERs more effectively, and enable an optimal combination of local power and power drawn from the bulk grid. This, in turn, helps reduce distribution-level costs and d-LMPs. The incorporation of a commitment score helps to maintain better reliability while still extracting flexibility from customers and DERs. Finally, the time-varying local retail tariffs lead to more efficient market scheduling and lower final costs for end-users, while ensuring that DERs and consumers are correctly compensated for the flexibility services they provide to the grid.

2.1.1 Related work

Several papers have addressed the topic of LEMs and can be grouped into three broad categories: (i) local markets, (ii) hierarchical market structures, and (iii) real-world deployments of local electricity markets. Category (i) deals with papers that introduce the concept of LEMs and related solutions [12, 34–37]. Reference [34] is the earliest reference in the literature for the term "local electricity market". Reference [12] describes how microgrids provided a way of aggregating smaller resources to participate in a market structure. Reference [35] deals with the concept of having customers and smaller-sized DERs participating in a market structure. Transactive energy is also a big driver in enabling LEMs, as compensating consumer resources for services rendered is a key concept in LEMs. LEMs can in fact be viewed as specific structural realizations of a transactive energy framework, and consider the wider system impact and interaction with the wholesale energy market (WEM) [36]. Other

methodologies exist which enable or incentivize the participation of DERs in the WEM, such as net energy metering (NEM), and net energy billing (NEB) [37]. However, these solutions do not encourage full participation of the resources and are often restrictive in their implementation [2].

Category (ii) corresponds to papers that layout LEMs with a direct interconnect to the WEM [9, 35, 38–42]. This differs somewhat from papers in Category (i) wherein the local market structures were largely theoretical constructs and still evolving with standardizations yet to emerge. Reference [38] details the interaction between WEM and LEM and provides numerical results on the cost savings provided by LEM. However, the paper does not disaggregate the price at the DSO and the consumer level, it rather uses a uniform price throughout. References [39, 40] propose alternative retail market structures that interact seamlessly with a WEM. They utilize a centralized optimization framework, an objective function that aims to minimize the operational costs for the market operator or the DSO, and primarily consider the market participants to be microgrids and/or aggregators. In contrast to these references, our earlier work in [9] proposed a distributed optimization framework which was used to minimize a combination of social welfare and line losses, subject to optimal power flow (OPF) constraints based on nonlinear DistFlow. In this framework, general agents representing DERs can bid into a local market at the primary feeder level, which interacts directly with a WEM. In addition to these papers, several surveys that capture the evolving LEM landscape have been carried out [35, 41, 42]. Of these, in [35], the authors lay out the evolving market structure that will enable customer participation in a market structure. Reference [41] is a survey that details the related work in LEM design, existing theoretical tools and models studied in the context of LEM, and challenges of realizing an LEM structure. Reference [42] carries out a survey of peer-to-peer markets. Our paper is similar to those considered in [41], and is a distinct addition in the form of a two-level, hierarchical LEM that includes grid physics, accurate forecasts of DER generation, DER characteristics such as follow-through or unmet commitment, and vulnerability to security breaches.

Category (iii) corresponds to works related to real-field implementation [43–46]. A pioneer solution to integrate DERs has been implemented in the New York electricity market, in the United States [43]. This reference, as well as [44], underscores the potential that DER participation brings in terms of coordinated system operation and shows the feasibility of smaller DER agents effectively participating in electricity markets. In addition to such efforts in the United States, there are several illustrations of successful LEMs in Europe [45, 46]. Reference [45] proposes a two-stage auction-based local market mechanism to allocate physical storage rights. Reference [46] discusses the recently proposed market design rules in the current context of the German market with numerical simulations, and a novel market design called Tech4all is introduced. The book [47] is an excellent reference for an overall state-of-the-art summary on LEM.

2.1.2 Our approach

Our work proposes an LEM that connects with the WEM and belongs to Category (ii). The LEM consists of a two-level hierarchical structure, located in the distribution grid, consisting of a secondary market (SM) at the lower level and a primary market (PM) at the upper level. Together, this hierarchical structure allows for an efficient functioning of the

distribution grid, which has to achieve multiple objectives and satisfy complicated constraints, by accommodating complex grid physics such as nonlinear and possibly unbalanced power flows, at the upper level by the PM and consumer needs and constraints, in the SM.

The PM and SM are assumed to be operated by an SMO and a PMO, respectively. The SM consists of DCA bids submitted to the SMO, who clears and schedules them to determine local electricity tariffs. The underlying optimization framework determines the schedules of real and reactive power injections of these DCAs (located at each secondary feeder) as well as their optimal flexibility ranges. Also included are constraints corresponding to bid flexibilities, budget, power balance, and capacity limits. A multiobjective cost function of disutility to the DCAs, net cost to the SMO, bid-commitment reliability, and bid-flexibility is utilized.

The PM consists of bids from each SMO at node i in the primary feeder, which are aggregations of the cleared market schedules of all the secondary feeders connected to i. The objective function includes a weighted combination of social welfare and line losses. The constraints on real and reactive power injections are determined using the net flexibility range of each SMO from the SM clearing. Power balance constraints based on nonlinear DistFlow, capacity limits on P and Q injections, thermal line limits, and bounds on voltages are included. The underlying optimization framework is based on the distributed approach in [9], and employs a PAC algorithm to determine the schedules of these agents at each i. The resulting market clearings include the real and reactive power injections at each primary feeder node and the d-LMPs, which correspond to the payments made by the PMO to the SMOs (or vice versa). The net injection from the entire primary feeder is conveyed at the distribution substation to the WEM.

Together with the SM at the lower level and PM at the upper level, our LEM is used to determine local electricity tariffs and d-LMPs that correspond to the payments between SMOs and DCAs, and between PMOs and SMOs, respectively. Both these prices capture the fine-grain locational and temporal variation in a distribution grid, and form the basis of an efficient LEM that can efficiently allocate resources and accurately compensate prosumers for grid services provided. In this chapter, we primarily focus on flexibility services that DERs and flexible loads can offer. This flexibility could be achieved by DCAs through several different types of actions, some examples include (i) load shifting or curtailment, (ii) dispatching distributed generation (such as diesel generators) or battery storage, (iii) using smart inverters to curtail active power from non-dispatchable renewable resources like rooftop PV, or alter Q injections by varying the power factor. In addition to determining the prices, the LEM is also used to determine and update commitment scores at the secondary feeder level, which quantify the ability of each DCA to fulfil their contractual commitments and follow the cleared market schedules. We validate the entire LEM using a modified IEEE 123-node test network, with 14% renewable generation in terms of nameplate capacity[1] and *up to* 50% flexible consumption[2] distributed over 79 primary feeder nodes, and between 3-5 secondary feeders at each of the nodes.

[1] relative to a peak load of 3.6 MW on the entire feeder

[2] relative to the baseline or nominal load

2.1.3 Our contributions

We propose a novel local electricity market for real-time energy transactions in a distribution grid with high DER penetration. The following are its unique features and the key contributions of our work:

- Hierarchical local electricity market structure that is electrically collocated with the current distribution network.
- Effectively address multiple functions of the distribution grid by virtue of the proposed LEM's hierarchy, with grid physics considered in the PM and consumer needs and constraints in the SM.
- Systematic approach showing how DERs, along with their flexibility bids, can be coordinated and aggregated in real-time via DCAs, and how these aggregated entities can participate in retail and/or wholesale electricity markets.
- Optimization of multiple objectives, including commitment reliability, net cost, flexibility, and utility in the SM; and net costs, utility, and line losses in the PM.
- Accommodation of disparate constraints, including budget, power balance, capacity, and flexibility limits in the SM; and nonlinear ACOPF constraints in the PM.
- Generation of a novel commitment score aimed at improving the reliability of our LEM by tracking the performance of DCAs at the secondary level.
- Fine-grain pricing in the form of local retail tariffs in the SM and d-LMPs in the PM that vary with both location and time, allowing for more efficient allocations in terms of lower costs and accurate compensation of DERs. This, in turn, provides an alternative to the current practices of net-metering and/or direct participation of DERs through aggregators at the WEM level (under FERC 2222).
- Validation of the entire LEM using a modified IEEE 123-bus, with a high penetration of DERs and flexible loads.

The rest of the chapter is organized as follows. In Section 3.2 we introduce the structure of the LEM, including the lower level SMO, upper level PMO, and the interactions between the PMO, SMO, and WEM. In Section 3.3 we present numerical results on a modified IEEE-123 test feeder, with high levels of PV penetration and load flexibility. In Section 3.4 we provide concluding remarks and topics for future work.

2.2 A local electricity market (LEM)

The LEM we propose in this work consists of a two-level hierarchical structure, located in the distribution grid, with an SM at the lower level and a PM at the upper level, a schematic of which is shown in Figs. 2.2 and 4.1. These markets are assumed to be operated by an SMO and a PMO, respectively, with combined oversight of both provided by a DSO. In

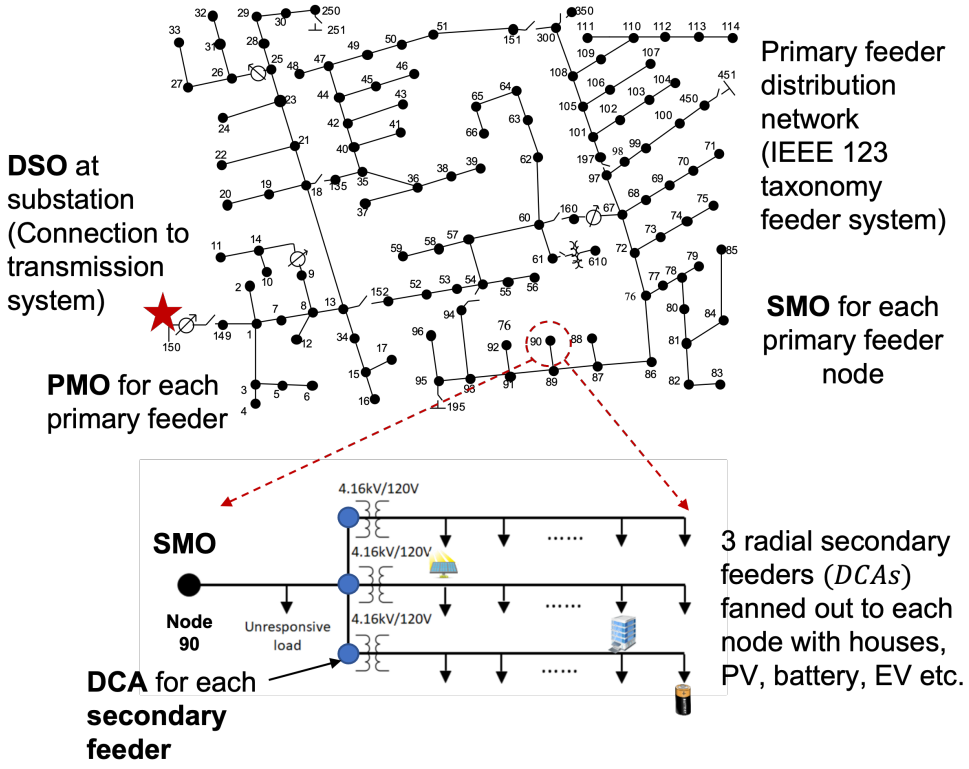

Figure 2.1: Overall schematic illustrating how the proposed LEM integrates seamlessly into the existing radial distribution network, and connects with bulk transmission.

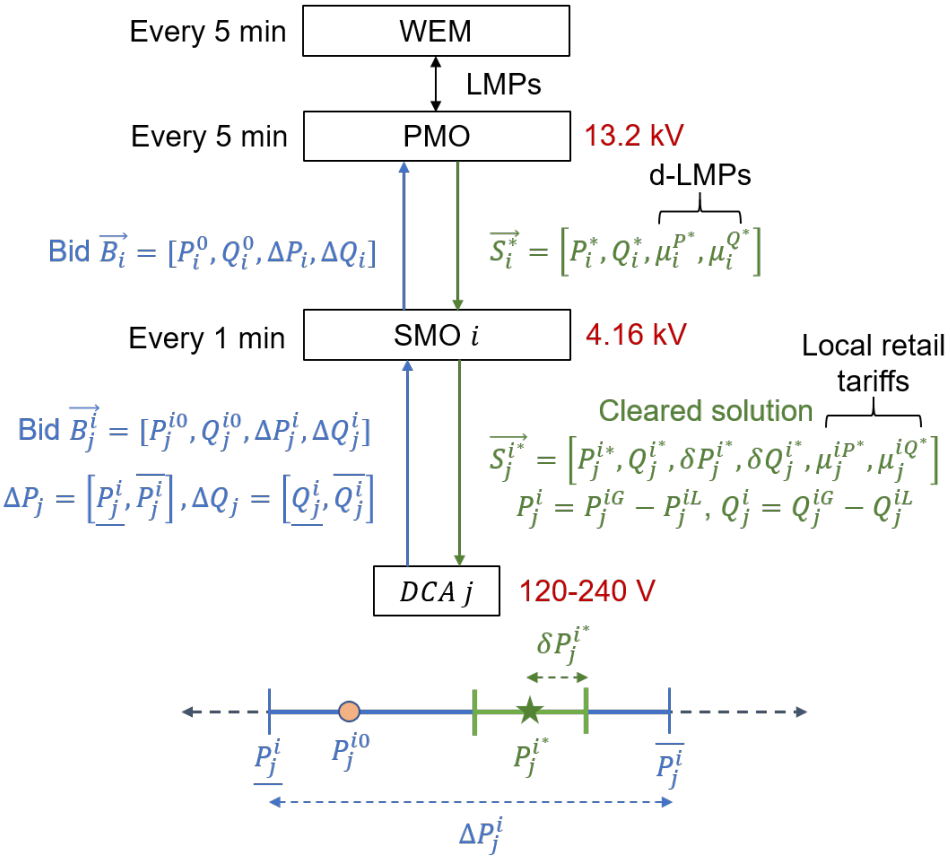

Figure 2.2: Proposed hierarchical LEM structure and optimization framework incorporating bid flexibility. The upper-level PMO coordinates with the WEM, while the SMO at the lower level oversees the DCAs. At the bottom of this figure, we show an example of a feasible flexibility bid from a DCA and the revised flexibility range that results from the SM clearing.

what follows, we provide details of the SM and PM, the interface between the two, and their timelines. Throughout this chapter, we define net injections as generation less load (i.e., $P = P^G - P^L, Q = Q^G - Q^L$), thus, net generation would be positive while net loads are taken to be negative. All bids are assumed to be based on load/generation forecasts one timestep into the future, for both the PM and SM. Thus, the cleared schedules hold for the entire duration of the next period (i.e., either Δt_p or Δt_s), until the next time market bidding and clearing occur.

2.2.1 SM structure

The following quantities are defined before we pose the optimization problem that serves as the backbone of the SM:

- $i \in \mathcal{N}_I$: Set of indices of all SMOs under consideration, downstream of a particular PMO.
- $j \in \mathcal{N}_{J,i}$: Set of indices of all DCAs under consideration, downstream of a given SMO i.
- P_j^{i0}, Q_j^{i0}: Baseline net active and reactive power injections for DCA j, which could either be forecasted by the SMO based on historical data or explicitly submitted by the DCA as part of its bid.
- $\Delta P_j^i = [\underline{P}_j^i, \overline{P}_j^i]$, $\Delta Q_j^i = [\underline{Q}_j^i, \overline{Q}_j^i]$: Bid flexibilities for each DCA. These intervals represent the range of maximum downward and upward flexibilities in P and Q injections being offered by the DCA.
- t_p and t_s: Timestamps for the PM and SM, respectively.
- $\Delta t_p, \Delta t_s$ and Δt_{WEM}: Time periods for the PM, SM, and WEM, respectively.
- $\hat{t}_p$: The most recent PM clearing prior to the current SM interval $[t_s, t_s + \Delta t_s]$.
- $P_i^*(\hat{t}_p)$ and $Q_i^*(\hat{t}_p)$: Real and reactive power setpoints, respectively, provided by the PM to SMOs at time $\hat{t}_p$.
- n_p, n_s: Number of primary clearing periods and number of secondary clearings per primary period, respectively.

The following decision variables are determined as outputs of the optimization for each DCA j bidding to SMO i:

- Decision vector $\vec{S}_j^i = [P_j^i, Q_j^i, \delta P_j^i, \delta Q_j^i, \mu_j^{iP}, \mu_j^{iQ}]$
- P_j^i, Q_j^i: Optimal power injections assigned as setpoints by the SMO i to DCA j.
- $[\delta P_j^i, \delta Q_j^i]$: Optimal symmetric flexibility ranges around the above setpoints P_j^i, Q_j^i, i.e. the DCA is directed to have net injections within these intervals $[P_j^i - \delta P_j^i, P_j^i + \delta P_j^i]$, $[Q_j^i - \delta Q_j^i, Q_j^i + \delta Q_j^i]$.

- μ_j^{iP}, μ_j^{iQ}: Local electricity tariffs.
- $C_j^i(t)$: Commitment score where $C_j^i(t) \in [0,1]$ reflects the SMO's confidence in whether the DCA j will reliably follow their committed injections within the flexibility range specified above.

The solutions correspond to power injections and net tariffs corresponding to each DCA, determined at t_s and applied over the period $[t_s, t_s + \Delta t_s]$, for all t_s. The following objective function and constraints define the underlying multi-objective optimization problem $\forall\ t_s,\ t_p\ (t_p \leq t_s)$. All quantities and variables in Eq. (5.2) are specified for the current secondary timestep t_s unless explicitly specified otherwise.

$$\min_{\vec{S}_j^i} \sum_{j \in \mathcal{N}_{J,i}} \{f_{j,1}^i, f_{j,2}^i, f_{j,3}^i, f_{j,4}^i\} \tag{2.1a}$$

$$f_{1,j}^i \succ f_{2,j}^i \succ f_{3,j}^i \succ f_{4,j}^i \tag{2.1b}$$

$$f_{j,1} = -C_j^i(P_j^{i^2} + Q_j^{i^2}),\ f_{j,2}^i = \mu_j^{iP} P_j^i + \mu_j^{iQ} Q_j^i$$
$$f_{j,3}^i = -(\delta P_j^i + \delta Q_j^i)$$
$$f_{j,4}^i = \beta_j^{iP}(P_j^i - P_j^{i0})^2 + \beta_j^{iQ}(Q_j^i - Q_j^{i0})^2$$

subject to:

$$\underline{P}_j^i + \delta P_j^i \leq P_j^i \leq \overline{P}_j^i - \delta P_j^i \tag{2.1c}$$

$$\underline{Q}_j^i + \delta Q_j^i \leq Q_j^i \leq \overline{Q}_j^i - \delta Q_j^i \tag{2.1d}$$

$$\delta P_j^i,\ \delta Q_j^i \geq 0, \tag{2.1e}$$

$$0 \leq \mu_j^{iP} \leq \overline{\mu}^{iP}, 0 \leq \mu_j^{iP} \leq \overline{\mu}^{iQ} \tag{2.1f}$$

$$\sum_{t_s}^{t_s+\Delta t_p} \sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iP}(t) P_j^i(t) \Delta t_s \leq \mu_i^{P^*}(\hat{t}_p) P_i^*(\hat{t}_p) \Delta t_p \tag{2.1g}$$

$$\sum_{t_s}^{t_s+\Delta t_p} \sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iQ}(t) Q_j^i(t) \Delta t_s \leq \mu_i^{Q^*}(\hat{t}_p) Q_i^*(\hat{t}_p) \Delta t_p \tag{2.1h}$$

$$\sum_{j \in \mathcal{N}_{J,i}} P_j^i(t_s) = P_i^*(\hat{t}_p), \quad \sum_{j \in \mathcal{N}_{J,i}} Q_j^i(t_s) = Q_i^*(\hat{t}_p) \tag{2.1i}$$

The cost functions in (5.3a) correspond to the following:

1. **Commitment $\mathbf{f_{j,1}^i}$**: This term maximizes the injections assigned to more trustworthy DCAs (i.e., C_j^i closer to 1) while minimizing the scheduling of DERs with lower commitment scores, who are relatively less likely to follow through on their contractual commitments.

2. **Net costs $\mathbf{f_{j,2}^i}$**: This term minimizes the *net* costs to the SMO for running its SM, which are composed of payments made out by the SMO to DCAs that are net generators, and denote revenue if the DCAs are loads.

3. **Flexibility $\mathbf{f^i_{j,3}}$**: These terms aim to maximize the aggregate flexibility that the SMO can extract from its DCAs, and in turn offer to the PMO.

4. **Disutility $\mathbf{f^i_{j,4}}$**: These terms aim to minimize the disutility or inconvenience caused to DCAs when they provide flexibility to the operator. Thus, our SMO is an altruistic entity that also considers welfare maximization for its DCAs. For our simulations in Section 2.4.2, the disutility coefficients were chosen as $\beta_j^{iP}, \beta_j^{iQ} \sim \mathcal{U}[0.1, 1]$.

The multiobjective optimization problem in Eq. (5.2) was formulated and solved using a hierarchical or lexicographic optimization-based approach. This method has been widely used in the literature to solve multiobjective problems [48, 49], particularly when the objectives have different units and may not be comparable in magnitude. In the hierarchical method, the different objectives are ranked in descending order in terms of their importance to the decision maker. The SMO orders their four objectives as shown in Eq. (5.3b), assigning commitment reliability as the most important goal and DCA disutility as being the least important. The SMO then solves a series of optimization problems, sequentially optimizing each of these objectives one at a time, in descending order of importance:

$$\min_{\vec{S}^i_j} F_k = \sum_{j \in \mathcal{N}_{J,i}} f^i_{j,k}(\vec{S}^i_j) \; \forall \; k = 1, 2, 3, 4 \tag{2.2}$$

$$\text{s.t. } f^i_{j,\ell}(\vec{S}^i_j) \leq (1+\epsilon) \sum_{j \in \mathcal{N}_{J,i}} f^i_{j,\ell}(\vec{S}^{i*}_j) = (1+\epsilon) F^*_\ell, \tag{2.3}$$

$$\forall \; \ell = 1, 2, \ldots, k-1, \; k > 1$$

$$\text{constraints in eqs. (2.1}c\text{) to (5.2}g\text{)} \tag{2.4}$$

At each step after optimizing the 1st objective (commitment reliability $f^i_{j,1}$), additional constraints are placed on the values of the previously optimized objectives as in Eq. (3.6). The hyperparameter ϵ controls the extent to which previous objective values are allowed to be degraded while searching for the new minima. We used $\epsilon = 0.05$ (5%) for our simulations. Using such a hierarchical approach also helps us get around the issues of our objective terms potentially being on different orders of magnitudes since we only minimize a single objective at each step, and the constraints in Eq. (3.6) look at the *relative* changes in the objective function values rather than their absolute magnitudes. Thus, we can proceed without needing to normalize any of the terms. This is especially advantageous here since most normalization methods require either prior knowledge of maximum or minimum objective values (which are not known beforehand in our case), or entail solving additional optimization problems to find these values at every iteration, which can be computationally expensive [50–52]. Finally, we note that multiobjective optimization problems in general do not have unique minima. Rather, the goal here is to find a *Pareto*-optimal set or efficient frontier of multiple possible solutions [48, 53].

The constraints in Eqs. (2.1c) to (2.1e) reflect the feasible flexibilities for each of the DCAs. Constraints in Eq. (2.1f) enforce that the cleared tariffs μ^i_j for any DCA j cannot exceed a price ceiling $\overline{\mu}^i_j$. These upper bounds may be set by an external regulatory authority such as a DSO, independent system operator (ISO), regional transmission operator (RTO), or

public utility commission (PUC). The budget constraints in Eq. (5.2e) and Eq. (5.2f) ensure that the total net payments made out by the SMO to its DCAs over all the SM clearings within each primary interval $[t_s, t_s + \Delta t_p]$ do not exceed its net revenue received from the PMO during the same period. We enforce these budget constraints separately for P and Q since proposed reactive power markets often behave quite differently compared to the energy market for real power [54]. Note that our optimization problem solves for the prices in terms of [$/kW] or [$/kVAR], so these are converted to [$/kWh] or [$/kVARh], respectively, before applying any of the budget constraints.

The constraints as written in Eq. (5.2e) and Eq. (5.2f) can be quite restrictive since they require the SMO to balance its budget for *every* PM period Δt_p. Thus, in addition to the strict budget constraints above, we also considered a softer implementation of the same. These require the SMO to balance its budget over the course of a longer time horizon (e.g., 1 day), allowing the SMO to run a deficit for some PM clearings, if needed. These are shown in Eqs. (2.5) and (2.6), where we consider a planning horizon of n_p primary clearing periods for the budget balance:

$$\sum_{0}^{n_p \Delta t_p} \sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iP}(t) P_j^i(t) \Delta t_s \leq \sum_{0}^{n_p \Delta t_p} \mu_i^{P^*}(t) P_i^*(t) \Delta t_p \tag{2.5}$$

$$\sum_{0}^{n_p \Delta t_p} \sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iQ}(t) Q_j^i(t) \Delta t_s \leq \sum_{0}^{n_p \Delta t_p} \mu_i^{Q^*}(t) Q_i^*(t) \Delta t_p \tag{2.6}$$

Since the optimization problem in Eq. (5.2) is currently framed as a single-period problem, we transformed our inherently multiperiod budget constraints to a *quasi-multiperiod* form for implementation. This was done by assuming that the SMO evenly redistributes its leftover net revenue over all the remaining SM clearings in the current budget period. For example, the quasi-multiperiod version of the strict budget constraint in Eq. (5.2e) is given by:

$$\sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iP}(t_s) P_j^i(t_s) \Delta t_s \leq \frac{\sum_0^{\hat{t}_p} \mu_i^{P^*}(t) P_i^*(t) \Delta t_p - \sum_0^{t_s - \Delta t_s} \sum_{j \in \mathcal{N}_{J,i}} \mu_j^{iP^*}(t) P_j^{i^*}(t) \Delta t_s}{\frac{\hat{t}_p + \Delta t_p - t_s}{\Delta t_s}} \tag{2.7}$$

where $\hat{t}_p$ is the most recent primary clearing time prior to this SM clearing, and the denominator is the number of secondary clearings left in the current PM interval $[\hat{t}_p, \hat{t}_p + \Delta t_p]$. A similar quasi-multiperiod form can be derived for the relaxed budget constraints in Eqs. (2.5) and (2.6) as well.

Finally, Eq. (5.2g) denotes power balance constraints for the SMO, where the sum total of net injections from the DCAs downstream needs to satisfy the net flows from the primary feeder upstream, These net injections $P_i^*(\hat{t}_p)$ and $Q_i^*(\hat{t}_p)$, at each primary feeder node (SMO) i are scheduled by the PMO. Since the PM clears less often than the SM, these values can be treated as constant for the SM optimization problem over each Δt_p. Suitable convex

relaxations were added to ensure feasibility and minimal optimality gap, and the resulting optimization problems were solved using Gurobi in Python[3].

The overall operation of the SM is summarized as follows: Starting with DCA bids $\vec{B}_j^i = [P_j^{i0}, Q_j^{i0}, \Delta P_j^i, \Delta Q_j^i]$ as inputs, the constrained optimization problem in eqs. (5.3a)-(5.2g) is solved by the SMO, with the market clearing resulting in net-injections $P_j^{i^*}(t_s)$ and local electricity tariff $\mu_j^{iP^*}(t_s)$, for real power. A similar set of injections and tariffs is derived for reactive power as well. Each of these solutions corresponds to the net-injection at DCA j, cleared at t_s, and holds for the following period $[t_s, t_s + \Delta t_s]$. Together with these net injections, the SMO also obtains, as a part of the above optimization procedure, a feasible flexibility δP_j^i in the real power injection for node j, and a similar flexibility for the reactive power injection. Together, the complete set of solutions from the SMO at secondary feeder j connected to the PMO at node i is given by $S_j^{i^*}$ (see Fig. 2.2). These market-cleared solutions are used to establish bilateral contracts between the SMO and its DCAs, with localized retail tariffs for the power injections, differentiated for each of them.

Another unique aspect of this SM structure is the commitment score $C_j^i(t)$. This is determined at every t_s, with a score decrease following every event where the DCA exhibits unmet commitment. As the goal is to have the SMO i reduce this score for such events and to reward the DCA j when they do follow through on their commitment, we propose the following recursive algorithm for a continuous update of $C_j^i(t_s)$:

$$C_j^i(t_s) = C_j^i(t_s - 1) - \frac{\widetilde{e_j^{iP}}(t_s) + \widetilde{e_j^{iQ}}(t_s)}{2} \;\forall\; j \in \mathcal{N}_{J,i} \tag{2.8}$$

$$\begin{aligned} e_j^P(t_s) = {} & [\![\hat{P}_j > \overline{P}_j^{i^*}]\!](\hat{P}_j - \overline{P}_j^{i^*}) + [\![\hat{P}_j < \underline{P}_j^{i^*}]\!](\underline{P}_j^{i*} - \hat{P}_j) \\ & + [\![\underline{P}_j^{i^*} \le \hat{P}_j \le \overline{P}_j^{i^*}]\!] \max(\hat{P}_j - \overline{P}_j^{i^*}, \underline{P}_j^{i^*} - \hat{P}_j) \end{aligned} \tag{2.9}$$

$$\widetilde{e_j^{iP}}(t_s) = \frac{e_j^{iP}(t_s)}{|P_j^{i^*}(t_s)|}, \quad \widetilde{e_j^{iQ}}(t_s) = \frac{e_j^{iQ}(t_s)}{|Q_j^{i^*}(t_s)|} \tag{2.10}$$

$$\widetilde{e_i^P}(t_s) = \frac{\mathbf{e_i^P}(\mathbf{t_s})}{\|\mathbf{e_i^P}(t_s)\|}, \; \widetilde{e_i^Q}(t_s) = \frac{\mathbf{e_i^Q}(t_s)}{\|\mathbf{e_i^Q}(t_s)\|} \tag{2.11}$$

where $\underline{P}_j^{i^*} = P_j^{i^*} - \delta P_j^{i^*}$, $\overline{P}_j^{i^*} = P_j^{i^*} + \delta P_j^{i^*}$. Here, the deviations in DCA injections are first normalized by their true setpoints in Eq. (5.16), followed by L_2 normalization in Eq. (5.18). These normalizations allow us to assess the relative performance of all DCAs under SMO i. The resulting commitment score serves as a metric of DCA reliability, and is proposed as a stand-alone indicator of DER commitment or a lack thereof. Further refinements of the SM could include the use of C_j^i for determining suitable penalties to these DCAs as C_j^i drops below unity and tightly interconnect with a corresponding ancillary market structure. Yet another component that could be coupled with C_j^i is the DCA's vulnerability to cybersecurity breaches. Currently these discussions are omitted, and will be pursued as part of future work.

[3] https://www.gurobi.com/

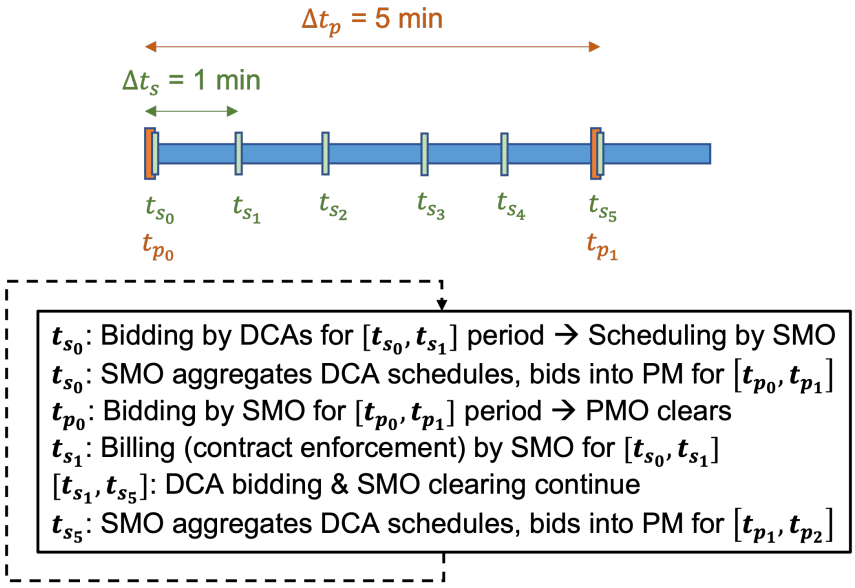

Figure 2.3: An illustrative timeline of the proposed LEM, also showing the interactions and interplay between the primary and secondary level markets.

2.2.2 Interface between SM and PM

$$P_i^0(t_p) = \sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*}(t_p),\ Q_i^0(t_p) = \sum_{j \in \mathcal{N}_{J,i}} Q_j^{i^*}(t_p) \tag{2.12}$$

$$\Delta P_i = \left[\underline{P}_i = \sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*} - \delta P_j^{i^*}, \overline{P}_i = \sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*} + \delta P_j^{i^*}\right]$$

$$\Delta Q_i = \left[\underline{Q}_i = \sum_{j \in \mathcal{N}_{J,i}} Q_j^{i^*} - \delta Q_j^{i^*}, \overline{Q}_i = \sum_{j \in \mathcal{N}_{J,i}} Q_j^{i^*} + \delta Q_j^{i^*}\right]$$

As mentioned earlier, an SM is located at each primary feeder node i, supervised by SMO i, and determines market clearings for all secondary feeders j connected to this node, with each j assumed to be represented by DCA j in the SM. The market clearing consists of $\left(P_j^{i^*}(t_s), \mu_j^{iP^*}(t_s)\right)$ at each t_s, which corresponds to power injections and local electricity tariffs for secondary feeder j. The market also returns an optimal flexibility $\delta P_j^{i^*}$ for each j. A similar set of clearing variables is associated with reactive power, too. At each primary market timestep t_p, this SM clearing is completed *before* the PM clearing at the same timestep. The SMO i in turn uses these DCA solutions from its SM in order to form its own bid into the PM at time t_p, representing primary feeder node i, that is at the upper level of our proposed LEM. The SMO aggregates power injections from all secondary feeders as its baseline injection P_i^0, at time t_p. The SMO bid also includes a corresponding aggregated flexibility ΔP_i based on the DCA-flexibilities $\delta P_j^{i^*}$. The specific bid determinations are given by Eq. (5.5). These bids and the corresponding PM clearing are assumed to occur every Δt_p, with the assumption that $\Delta t_p > \Delta t_s$.

2.2.3 PM structure

The starting point for the PM is the SMO bids P_i with the flexibility δP_i, and a similar set of quantities for reactive power, at each primary feeder node i. The PMO (possibly at a substation), which has oversight over the entire primary feeder, accepts these bids and clears the PM at every t_p at intervals of Δt_p. The PMO clears the market, with an underlying distributed optimization framework that facilitates market clearing. This framework includes appropriate cost functions such as Social Welfare and line losses, and constraints that correspond to power physics constraints modeled with nonlinear DistFlow equations (branch flow model), as well as various network-level constraints.

The underlying optimization problem for the PMO over a radial distribution grid is defined in Eq. (5.6), where P_i, Q_i, v_i are the nodal real power, reactive power, and voltages, superscripts G and L denote generation and load, and $P_{i',i}, Q_{i',i}, l_{i',i}, R_{i',i}, X_{i',i}$ denote the real power, reactive power, squared current magnitude, resistance, and reactance of the branch from node i' to i. These nonlinear constraints describe the power flow in the radial distribution grid, assuming balanced flows, small angles, and a convexification of the definition of power using second-order cone constraint programming (SOCP). The remaining constraints describe the network voltage constraints and generator and load flexibility limits.

$$\min_{y} \ f^{S-W}(y) \tag{2.13}$$

subject to:

$$v_i - v_{i'} = \left(R_{i',i}^2 + X_{i',i}^2\right) l_{i',i} - 2\left(R_{i',i}P_{i',i} + X_{i',i}Q_{i',i}\right)$$

$$P_i^G - P_i^L = -P_{i',i} + R_{i',i} l_{i',i} + \sum_{k:(i,k)\in\mathcal{T}} P_{i,k}$$

$$Q_i^G - Q_i^L = -Q_{i',i} + X_{i',i} l_{i',i} + \sum_{k:(i,k)\in\mathcal{T}} Q_{i,k}$$

$$P_{i',i}^2 + Q_{i',i}^2 \leq \overline{S}_{i',i}^2$$

$$P_{i',i}^2 + Q_{i',i}^2 \leq v_i l_{i',i}$$

$$\underline{v}_i \leq v_i \leq \overline{v}_i$$

$$\underline{P}_i^G \leq P_i^G \leq \overline{P}_i^G$$

$$\underline{P}_i^L \leq P_i^L \leq \overline{P}_i^L$$

$$\underline{Q}_i^G \leq Q_i^G \leq \overline{Q}_i^G$$

$$\underline{Q}_i^L \leq Q_i^L \leq \overline{Q}_i^L$$

The specific cost function that will be used is as below:

$$f^{S-W}(y) = \sum_j \Big[f_i^{\text{Load-Disutil}}(y) + f_i^{\text{Gen-Cost}}(y) \Big] + \xi \Big[\sum_k f_k^{\text{Loss}}(y) \Big] \tag{2.14}$$

$$f_i^{\text{Load-Disutil}}(y) = \beta_i^P (P_i^L - P_i^{L0})^2 + \beta_i^Q (Q_i - Q_i^{L0})^2 \tag{2.15}$$

$$f_i^{\text{Gen-Cost}}(y) = \begin{cases} \alpha_i^P (P_i^G)^2 + \alpha_i^Q (Q_i^G)^2, \\ \lambda_i^P P_i^G + \lambda_i^Q Q_i^G, \text{if } i \qquad \text{is PCC} \end{cases} \tag{2.16}$$

$$f_{i',i}^{\text{Loss}}(y) = R_{i',i} l_{i',i} \tag{2.17}$$

where i are the indices for SMOs participating in the PM, P_i^L, Q_i^L and P_i^G, Q_i^G are the loads and generation at node i respectively, P_i^{L0} and Q_i^{L0} are the baseline loads, $l_{i',i}$ and $R_{i',i}$ are the squared magnitude of current and resistance of line from node i' to i. The cost coefficients for load disutility of the SMO are computed as $\beta_i^P = \frac{1}{|\mathcal{N}_{J,i}|} \sum_{j \in \mathcal{N}_{J,i}} \beta_j^{iP}$, where notation $|\mathcal{N}_{J,i}|$ denotes the cardinality of the set $\mathcal{N}_{J,i}$, i.e. the number of DCAs downstream of SMO i, and likewise for reactive power disutility β_i^Q. The quadratic coefficients for generating cost are α_i^P and α_i^Q, and the wholesale price of power (LMP) from the WEM at the point of common coupling (PCC) is λ^P and λ^Q. The coefficient of power loss ξ in Eq. (5.11) reflects the trade off between minimizing line losses versus minimizing disutility and generation costs. These coefficients would be privately chosen by the SMO according to its preferences. For our simulations, we used a constant value of $\xi = 100$ for all SMOs. We chose this value in order to balance the tradeoffs of socioeconomic costs versus line losses. Using a ξ value that's too low would devalue the line losses term entirely due to relative scaling issues, and a value that's too high is also unrealistic because economic and utilitarian decisions drive the market. During actual implementation, this power loss factor would likely have to be tuned by each SMO over time, based on their operational objectives and the above-mentioned tradeoffs. Finally, the PMO interfaces with a market operator at the substation and bids into the WEM every $\Delta t_{WEM} = 5\ min$.

The linear cost term in Eq. (5.3a) for the SM accounts for the retail costs to customers and payments to DGs. However, we use a quadratic cost in Eq. (5.13) for the PM in order to account for the additional costs to the SMO for operating the SM and maintaining its distribution network. These include procuring reserves, adequate storage, and standby or auxiliary generation capacity[4] for contingencies, e.g., to meet shortfalls in case of unmet commitments from its DCAs. These *fixed* and amortized cost coefficients for the SMOs were chosen to be between $\alpha_i^{P,fixed} \in [4, 8]$ $\$/p.u.$ [55, 56], with $S_{base} = 1\ MVA$. In addition to these O&M costs, the SMO also adjusts these cost-coefficients at every t_p by a small variable amount $\alpha_i^{P,var}$, based on the weighted average of the retail tariffs across all its DCAs from the previous SM clearing. This variable cost component reflects the average cost of electricity in the SM and thus can also be interpreted as a price bid or offer by the SMO to the PMO,

[4]Resources like large-scale battery storage, diesel gensets, natural gas peaker plants, etc.

resulting from the SM optimization.

$$\alpha_i^{P,var}(t_p) = \frac{\sum_{t_p - n_s \Delta t_s}^{t_p} \sum_{j \in \mathcal{N}_{J,i}} \mu_i^{jP^*} \left| P_i^{j^*} \right|}{\sum_{t_p - n_s \Delta t_s}^{t_p} \sum_{j \in \mathcal{N}_{J,i}} \left| P_i^{j^*} \right|} \tag{2.18}$$

$$\alpha_i^{P}(t_p) = \alpha_i^{P,fixed} + \alpha_i^{P,var}(t_p) \tag{2.19}$$

With the costs and constraints as above, the OPF problem is solved using CVX in MATLAB[5] to carry out the market clearing, which consists of the power injection at each SMO i and d-LMP, the electricity price corresponding to primary node i and determined using the dual variable obtained from the OPF. These clearings are implemented through a distributed optimization algorithm denoted as PAC [8, 9], which consists of peer-to-peer communication between neighboring SMOs in an autonomous manner. This makes the computation more tractable and also reduces communication latencies. Upon reaching an agreement with its neighbors, each SMO enters into a bilateral agreement with its PMO, thereby committing to deliver or consume the decided amount of power at the d-LMP rate. Any net loads consumed by the SMO will be charged the d-LMP, and equivalently, net generation by an SMO will be remunerated at the d-LMP. All payments will be made to/from the PMO.

2.2.4 Timelines of SM and PM

The overall time scales of the secondary and primary levels of our proposed retail market in relation to the real-time market in a WEM are indicated in Fig. 2.3. Market clearings of the SM and PM are assumed to be every Δt_s and Δt_p apart, with $\Delta t_s < \Delta t_p$. For the use-case study, we assume that $\Delta t_p = \Delta t_{WEM}$, i.e., the PM and WEM are cleared together in lockstep. We also assume that $\Delta t_s = 1\ min$ and $\Delta t_p = 5\ min$, and that the SM clearing occurs arbitrarily quickly as the complexity and dimensionality of the underlying optimization problem are low.

2.2.5 Assumptions, Observations, and Extensions

The LEM proposed here has been constructed using a hierarchical structure precisely to address the distinct challenges that a distribution grid poses in comparison to a transmission grid. This hierarchical structure allows an efficient incorporation of multiple objectives and constraints simultaneously present. Since the SM is closer to the end-user in both location and time, we constructed the SM to be more consumer-centric, with costs and constraints pertaining to consumer flexibility and needs. Since the PM, relatively speaking, has a complex set of physical network topologies, we pay greater attention to the physical costs and constraints in its problem formulation. This allows the DSO, overall, to address the varied roles of reliability challenged by grid physics, and flexibility and granularity in location and time challenged by the presence of disparate consumers with varied needs and constraints. In this study, we assume that the SM is cleared more frequently compared to the PM, in order to quickly accommodate any variations that may occur locally at the DCA level. However,

[5] http://cvxr.com/cvx/

this condition is not necessary to operate our hierarchical LEM, allowing both markets to be in lockstep if need be.

In this chapter, we considered a single-period optimization problem solved by the SM at each timestep. However, our market structure can be applied to the multiperiod optimization setting as well. We are currently working on extensions of our model where the SMO and PMO perform multiperiod optimization over a planning horizon into the future, using an approach inspired by optimal control or model predictive control (MPC), similar to [57]. For instance, the core structure of the SM optimization problem would remain the same as in Eq. (5.2) but would now optimize over several timesteps T into the future, subject to similar constraints as in Eq. (2.1c)-Eq. (5.2g):

$$\min_{\vec{S}^i_j(t)} \sum_{t=\tau}^{\tau+T} \sum_{j \in \mathcal{N}_{J,i}} \{f^i_{1,j}(t|\tau), f^i_{2,j}(t|\tau), f^i_{3,j}(t|\tau), f^i_{4,j}(t|\tau)\} \tag{2.20}$$

where $(t|\tau)$ denotes predictions or estimates of quantities for future periods t made at time τ. Thus, the decision vector $\{\vec{S}^i_j(t)\}^T_{t=\tau}$ is now higher dimensional since it spans multiple SM periods. In addition to future flexibility bids from DCAs, i.e., $(P^{i0}_j(t|\tau), Q^{i0}_j(\tau), \Delta P^i_j(t|\tau), \Delta Q^i_j(t|\tau))$, the SMO also needs to predict future PM solutions $(\mu^{P^*}_i(\hat{t}_p|\tau), P^*_i(\hat{t}_p|\tau), \mu^{Q^*}_i(\hat{t}_p|\tau), Q^*_i(\hat{t}_p|\tau)$ in order to solve the SM multiperiod optimization problem. Herein lies the main challenge of extending to the multiperiod setting. We are exploring several time series forecasting tools like autoregressive integrated moving average (ARIMA), exponential smoothing, etc., for this purpose. This optimization would result in both binding spot values that apply for the very next timestep, as well as future values for further into the planning horizon that are non-binding. One of the main benefits of the multiperiod approach is that it allows us to implement more realistic formulations of the budget constraint, without having to resort to assumptions as described in Eq. (2.7). For example, the multiperiod version of the budget balance for active power would be:

$$\sum_{t \in \mathcal{T}_s} \sum_{j \in \mathcal{N}_{J,i}} \mu^{iP}_j(t) P^i_j(t) \Delta t_s \leq \sum_{t' \in \mathcal{T}_p} \mu^{P^*}_i(t'|\tau) P^*_i(t'|\tau) \Delta t_p \tag{2.21}$$

where $\mathcal{T}_s$ and $\mathcal{T}_p$ denote the set of SM and PM clearing times within the current planning horizon T, respectively. A multiperiod approach would also allow us to include inter-temporal constraints, in order to better optimize the scheduling of energy storage devices like batteries and EVs, as well as thermostatically controlled loads (TCLs) such as HVAC systems and water heaters [58].

Our focus in this work has been on real-time energy markets. We have not addressed issues such as settlements and billing, as the relevant discussions will have to necessarily include ancillary markets and reserves to deal with any unmet commitments and supply-demand imbalances, in real-time. Extensions similar to co-optimization [59] of these different markets are expected to be possible. Currently, bids into the PM are synthesized from the SM clearings through a simple aggregation process. Advanced game-theoretic approaches such as [60, 61] have the potential to generate more intelligent bids and are a topic for future research. The inclusion of the commitment score in our optimization problem is an effort to address consumer-centric constraints that could lead to unmet commitments. This may be

due to a variety of factors such as (i) willful reneging on contracts, (ii) malicious behavior due to system compromise or security breaches, and (iii) changes in environmental or weather conditions. The commitment score could also be potentially used for determining penalties and tariffs, thereby leading to a more efficient market design. Details of this effort are part of our future work as well.

2.3 Example instance of LEM

To better illustrate our proposal, we also sketch out a possible (hypothetical) instantiation of our dual-layer LEM for the city of Boston, MA, which is located in the New England (NE) region in the US. Fig. 2.4 shows the IEEE 39-bus transmission system [62], which is a synthetic representation of the entire NE region, with a peak load of 6254 MW and a total of around 7 million homes. This corresponds to ≈ 162 MW and 180,000 homes per bus. Given that the IEEE 123-node distribution feeder has a peak load of roughly 3.6 MW, we can estimate that there will be 44 such primary feeders per transmission bus and 4100 homes per feeder. Thus, the city of Boston, with a total of 300,000 homes [63], will be served by 73 primary feeders across 2 transmission buses.

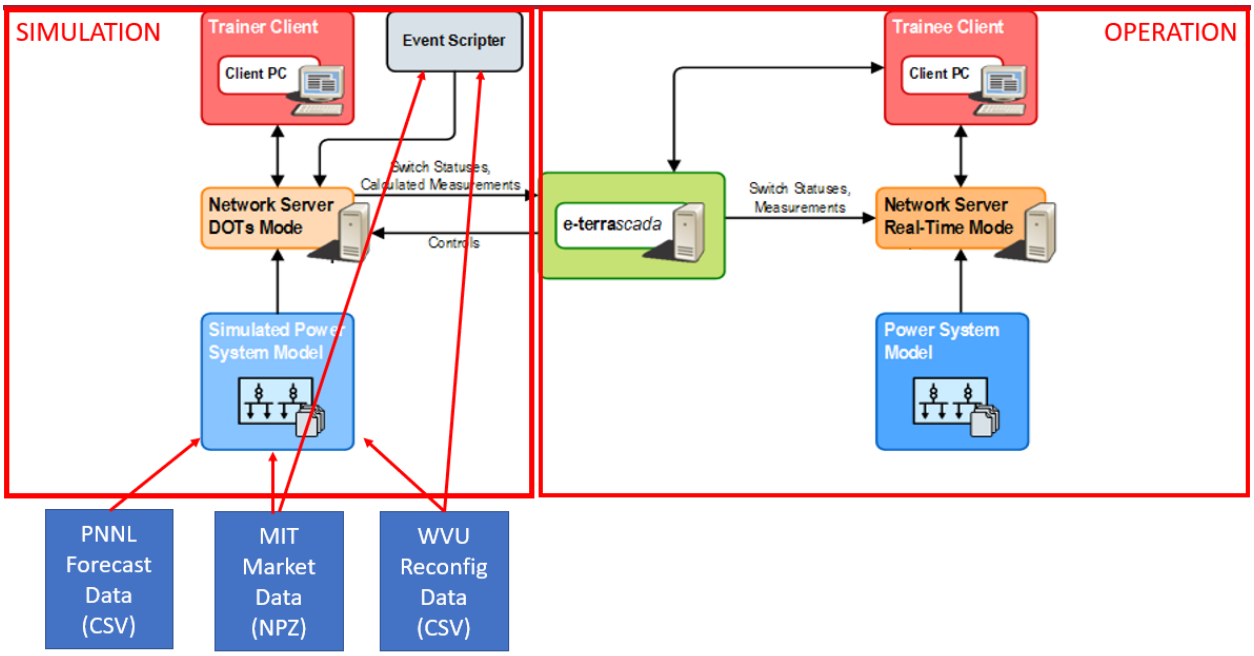

Figure 2.4: IEEE 39-bus transmission system.

Fig. 2.5 shows a breakdown of different entities to form a hierarchical LEM for Boston. Note that the main market operators and agents that are relevant for this work are marked in green.

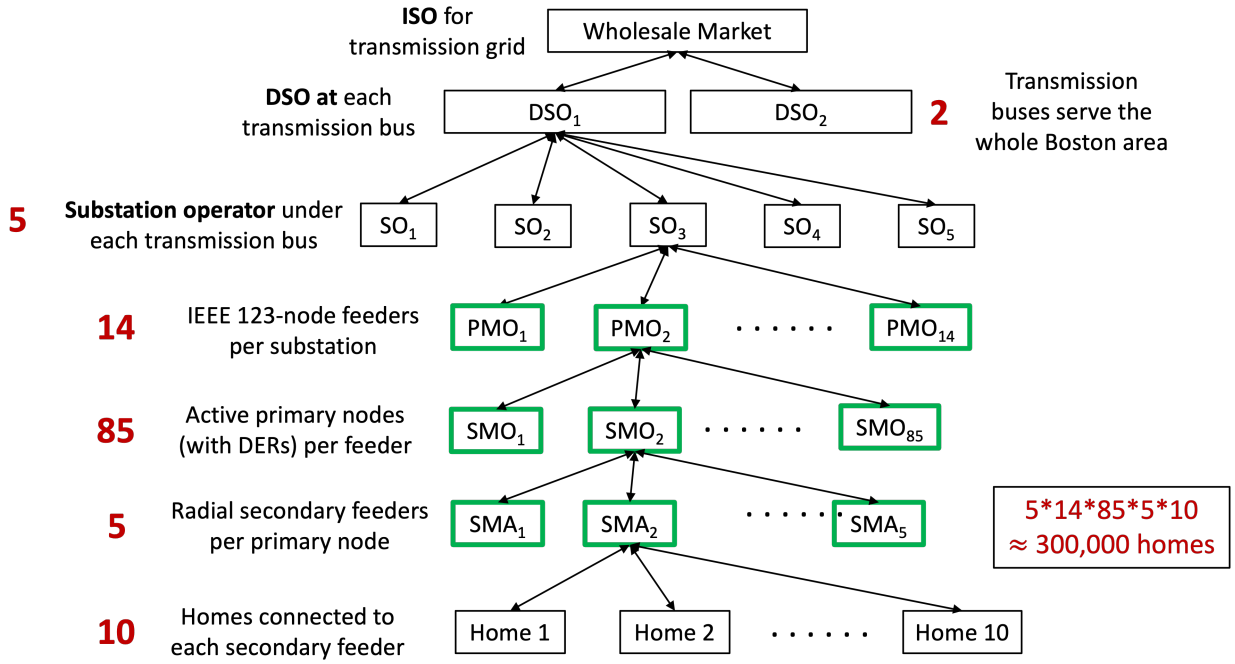

Figure 2.5: Example of hypothetical LEM for the city of Boston, MA.

2.4 Results and Discussion

2.4.1 The Use-case

The hierarchical LEM proposed in Section 3.2 is evaluated using a modified IEEE-123 test feeder. A GridLAB-D model[6] was utilized to simulate this test feeder over the course of a 24 hour period. Rooftop PV (with smart inverters) was assumed to be present at nodes 5, 20, 50, 63, and 94, with a total PV generation capacity of 510.3 kW. This corresponded to a DER (PV) penetration of about 14%, assuming that the peak load is at about 3.6 MW [64, 65]. An SMO was assumed to be present at 79 of the primary feeder nodes (i.e. $|\mathcal{N}_I| = 79$), and that flexible loads were present at all of these nodes with each DCA capable of *up to* $\pm 50\%$ deviations around their baseline injections. We assumed this maximum flexibility based on past studies forecasting demand response potentials in the US [66]. The GridLAB-D model included triplex meters to record P and Q injections every minute at each of these 79 nodes. Weather data for Boston, MA, was used to forecast PV generation, and real-time 5-minute LMPs from ISO-NE for August 28, 2021 were used as input data to the SM and PM optimization problems [67]. Since no reactive power market currently exists, we assumed the Q-LMP to be 10% of the P-LMP [68]. The price ceilings in Eq. (5.2) were set to be $\overline{\mu}^{iP}, \overline{\mu}^{iQ} = 0.2$ \$/kWh, which is almost twice the current average retail rate of 0.129 \$/kWh charged by Eversource, a utility in Massachusetts[7]. The overall test feeder was converted to a balanced 3-phase distribution network by (i) assuming switches to be at their normal

[6] https://www.gridlabd.org/

[7] https://www.eversource.com/content/ema-c/residential/my-account/billing-payments/about-your-bill/rates-tariffs/summary-of-electric-rates

positions, (ii) converting single phase spot loads to be 3-phase, (iii) assuming cables to be 3-phase transposed, (iv) converting configurations 1 thru 12 to symmetric matrices and (v) modeling shunt capacitors as 3-phase reactive power generators [9]. A PMO was assumed to be at the slack bus, at 13.2kV, with the SMOs at 4.16kV, and each DCA at 120-240V.

Each SMO was assumed to have anywhere between $|\mathcal{N}_{J,i}| \in [3, 5]$ DCAs with the actual number chosen uniformly at random. The number of DCAs at each SMO i is chosen independently. We set the baseline injections P_j^{i0}, Q_j^{i0} to be equal to the results from the GridLAB-D simulations. Since the injection data was only available up to the primary feeder node level, we artificially disaggregated the injections at each SMO amongst its DCAs, with each DCA being either a net load or net generator. The flexibility bids for the SM $\Delta P_j^i, \Delta Q_j^i$ were also randomly generated, allowing each DCA to offer flexibilities of *up to* $\pm 50\%$ away from their baseline. Thus, the upper and lower limits for the bid flexibilities were set as $\underline{P}_j^i = P_j^{i0}(1 - \underline{\Delta}_i^j), \overline{P}_j^i = P_j^{i0}(1 + \overline{\Delta}_i^j)$, where $\underline{\Delta}_i^j, \overline{\Delta}_i^j \sim \mathcal{U}[0, 0.5]$. We focus here on the results for active power only; similar trends were observed for reactive power.

2.4.2 SM scheduling

The first step in our use-case study is the SM structure, and its market clearing using the optimization problem outlined in eqs. (5.3a)-(5.2g). The bids $\vec{B}_j^i$ corresponding to these parameters are shown in Fig. 2.6a for a randomly selected SMO $i = 7$ having 3 DCAs $j = 1, 2, 3$. The interval of interest was chosen to be of a 60-min duration, with the actual hour chosen at random. The power injections $P_j^{7^*}$ obtained from solving (5.3a)-(5.2g) as well as the corresponding flexibilities, for each DCA j, are indicated in Fig. 2.6b. These two figures clearly illustrate the optimal flexibility range for each of the DCAs, reflecting the ability of the SM to incorporate the constraints of the DCAs and multiple objectives such as utility, monetary costs, and commitment reliability. The corresponding local electricity tariffs, $\mu_j^{7P^*}$ are shown in Fig. 2.6c for $j = 1, 2, 3$. Figs. 2.6b and 2.6c also illustrate the correlation between injections and prices. For instance, the tariffs for DCA 3 are consistently higher than those for 1 and 2, as DCA 3 is more heavily loaded than the other DCAs. Similarly, tariffs for DCA 1 are lower as its net generation is higher; the price fluctuations are more or less in sync with generation and demand patterns.

2.4.3 PM scheduling

The optimal injections with associated flexibilities from the SM clearing in Fig. 2.6b are aggregated across all three DCAs to form this SMO's bid $P_7^0, \Delta P_7$ into the primary level market, as described in Eq. (5.5). The resulting SMO bids are shown in Fig. 2.7a, where the solid red line indicates P_7^0 and the shaded area indicates the flexibility range $[P_7^0 - \Delta P_7, P_7^0 + \Delta P_7]$. These bids are in turn used to solve the PM OPF problem in Eq. (5.6) using the distributed PAC algorithm, where the SMO's flexible bids $\Delta P_7 = [\underline{P}_7, \overline{P}_7]$ set the feasible operational limits for the power flow constraints in (5.6). Solving this optimization problem corresponds to clearing the PM, and determines the PM schedules for the SMO. The results of the PM clearing for SMO $i = 7$ are shown in Fig. 2.7c.

Our two-tier market structure generates two sets of schedules and prices, every 1 min and every 5 min for the SM and PM, respectively, shown in Fig. 2.6 and Fig. 2.7. We further note

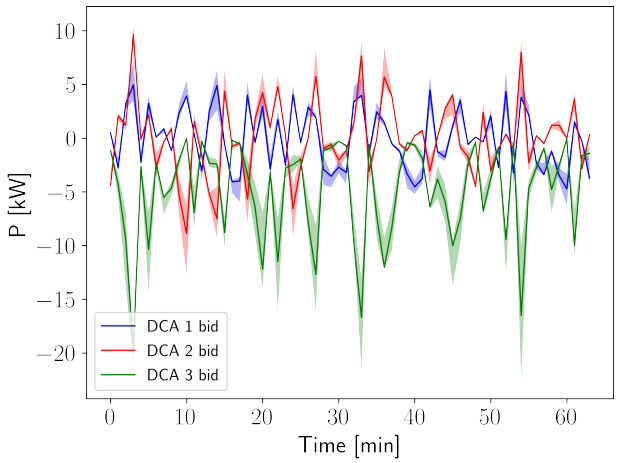

(a) Bids with flexibilities $P_j^{7,0}, \Delta P_j^7$.

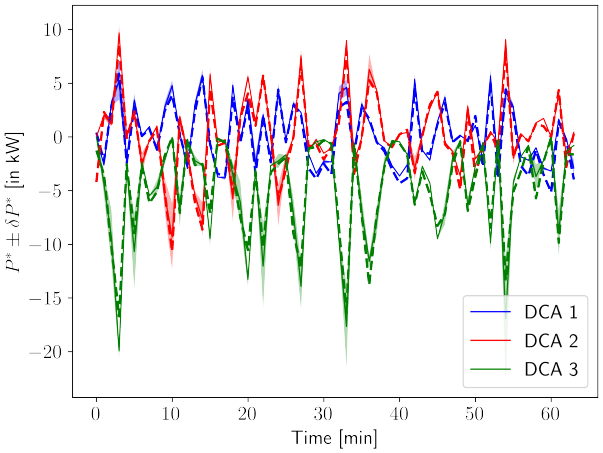

(b) DCA schedules and responses.

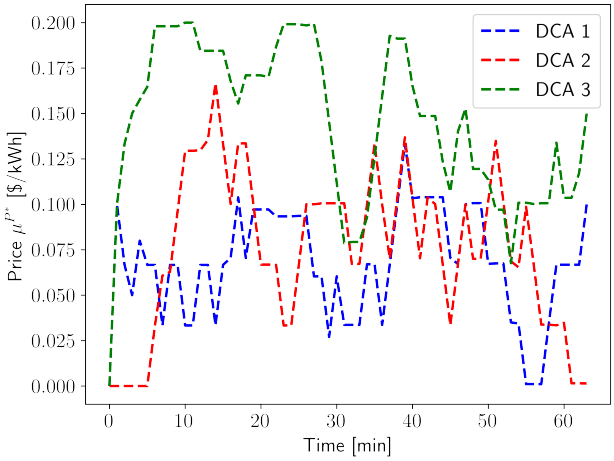

(c) Market cleared local retail tariffs.

Figure 2.6: SM bidding and clearing for primary feeder node 7, with 3 DCAs $j \in \{1, 2, 3\}$. The solid lines in Fig. 2.6a and Fig. 2.6b represent the baseline injection bids and market cleared setpoints, respectively, while the shaded regions around them are the flexibility ranges. Local retail tariffs from the SM $\mu_j^{7P^*}$ are shown in Fig. 2.6c. The SMO aggregates these PM schedules to bid into the PMO as in Fig. 2.7a. The dashed lines in Fig. 2.6b indicate the actual responses of the DCAs in response to their market-cleared schedules.

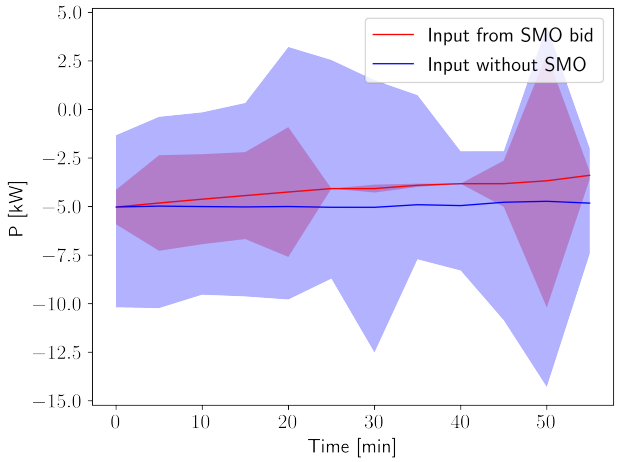

(a) Inputs from node 7 for PM clearing.

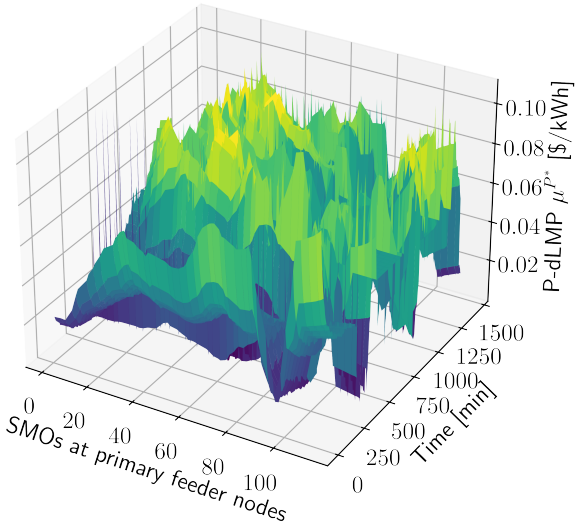

(b) d-LMPs across all SMO nodes over 1 day.

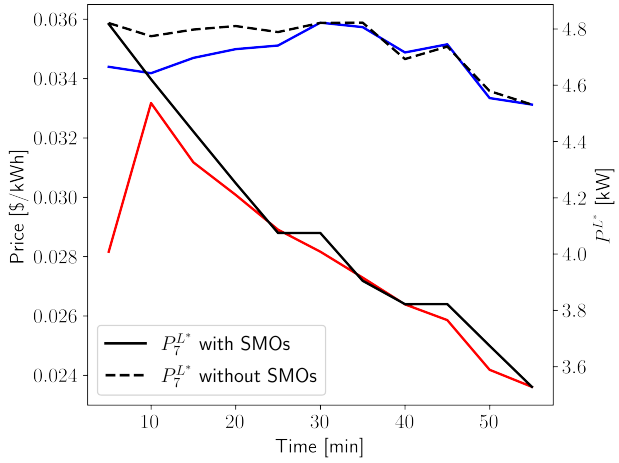

(c) PM solutions for SMO 7. The solid and dashed black lines are the load injections, while the red and blue lines are the d-LMPs with and without the SMO, respectively.

Figure 2.7: Selected solutions from the PM clearing.

that both the local electricity tariffs and the d-LMPs determined by the SM and PM display a high degree of spatio-temporal variations, as shown in Fig. 2.7b. This illustrates the need for local primary and secondary markets to capture such changes with sufficient resolution.

In order to evaluate the impact of the hierarchical structure that we have included in the LEM, we compare the performance of the PM to the case when there is no SM at the lower level. The 'without SMO' scenario consists of only a PM, with the PMO directly assuming flexibility ranges for each primary feeder node that best represent an aggregation of local generation and curtailable loads. In what follows, we compare the performance of our hierarchical LEM, i.e., the 'with SMO' scenario, with the 'without SMO' scenario. First, we compare the inputs into the PM at node 7.

Figure 2.7a shows that the PMO has a larger flexibility range that may not be accurate or realizable. The red curve in Fig. 2.7a shows that the flexibility range with SMO is narrower, and reflects the true preferences of the DCAs. Furthermore, the amount of flexibility that the SMO provides to the PMO is also impacted by other factors like the SM retail costs and the commitment scores of each of its DCAs, both of which vary with time. As a result, the 'with SMO' case is more performant as the baseline injection is optimized in comparison to the relatively ad-hoc choice in the without SMO case (the blue curve in Fig. 2.7a).

We next compare the performance of our hierarchical market across the entire primary feeder consisting of all 79 SMO nodes, over the course of the whole simulation period of 24 hours. In Fig. 2.8a, the inputs to the PMO are shown (the red curve), with all SMO solutions aggregated across all 79 primary feeder nodes $i \in \mathcal{N}_I$ and for the entire day. We see that without the additional visibility and granularity offered by the SM structure, the PM would assume much larger ranges for the injection limits in the 'without SMO' case (the blue curve) when compared to the 'with SMO' case. These are less accurate and may also be overoptimistic in terms of how much flexibility can be realistically expected from the DCAs, which in turn can cause issues in case of reneged commitments. It should be pointed out that the amounts of local generation seen in Fig. 2.7 and Fig. 2.9b are above the installed PV capacity of 510.3 kW. This is because while generating the synthetic flexibility bids for the DCAs, we allowed for the possibility of additional DERs like batteries, EVs, and curtailable or shiftable loads, present at each of these secondary feeders, which weren't explicitly modeled in the GridLAB-D simulation.

In Fig. 2.9a, we see that the d-LMPs both with and without the SMO are generally higher than the LMP, which is expected since the d-LMPs account for additional costs associated with congestion, line losses and other delivery charges incurred by the PMO and DSO in the distribution network, downstream of the substation. The d-LMP with the SMO does fall slightly below the LMP between 100-500 minutes (02:00:0700). This can be explained by the total electricity demand being low during this period which in turn occurs as the SMOs are able to curtail flexible loads to a larger extent by coordinating their DCAs more intelligently and compensate them accordingly at the local retail tariff rate. In fact, we find that the SMOs are able to achieve higher levels of load curtailment throughout the course of the day when compared to the case without SMOs. Once again, this is likely because the SMO can access additional information on DCA's preferences and effectively utilize any additional flexibility that they're willing to provide. The SM allows the SMO to more efficiently allocate resources amongst the secondary feeders at each primary feeder node, and take advantage of differences in load and generation profiles across DCAs over time, since they could potentially

complement each other.

The second observation from Fig. 2.9a and Fig. 2.9b is that the 'with SMO' case schedules lower levels of local generation mid-day compared to the 'without SMO' case. This may be due to a combination of multiple objectives utilized in the SM that include both net costs and flexibility. The optimal behavior, as a result, as predicted by the LEM, is one where more power is purchased from the main transmission grid rather than from local generation midday. This is also supported by Fig. 2.9a, which shows that such a behavior leads to lower d-LMPs and reduced distribution network costs with the hierarchical LEM than without the SMO. This is desirable since the SMOs can then reduce the retail tariff charged to their DCAs, improving affordability for customers, as seen in Table 2.1. It also ensures that DSOs aren't over-compensating prosumers with DERs. This can help avoid excessive cross-subsidies from consumers to prosumers, which is a major challenge associated with net energy metering (NEM) programs today [69, 70], and can thus produce more equitable allocations.

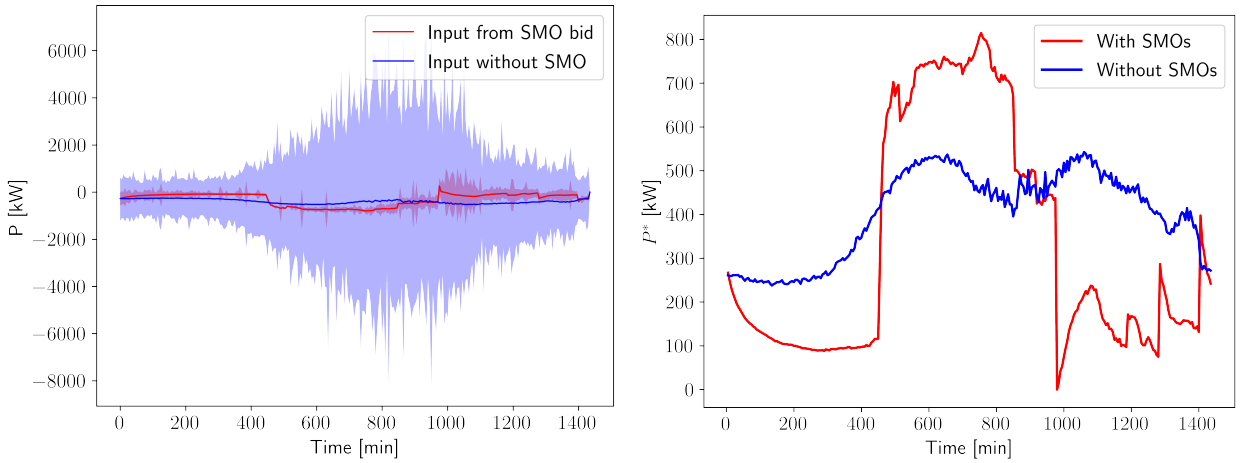

(a) Inputs to PM aggregated across all primary feeder nodes except the slack bus.

(b) PM solutions for net injections at the slack bus.

Figure 2.8: Comparison of PM bids (or inputs) and slack bus injections, with and without SM. The slack bus (node 149) is connected to the substation and distribution transformer. Positive injections here indicate that the feeder as a whole is importing power from the main grid.

Figures 2.8b, 2.9a, and 2.9b correspond to the main conclusions of the proposed LEM. In all three figures, the red curves correspond to the behavior with the LEM while the blue curves correspond to the 'without SMO' case. The red curve in Fig. 2.8b shows that the LEM schedules generation from the bulk grid more in the middle of the day and less otherwise; those in Fig. 2.9b show that it's advantageous to increase local generation in the latter part of the day and to curtail load in the earlier part of the day. The LEM determines that the IEEE-123 feeder needs to import around 700 kW between minutes 400 to minute 850, and less than 300 kW from minute 1000 onward. This behavior is significantly different from the market structure without SMOs, as the primary market alone does not have the granular customer-level information to accurately estimate the power injections and their associated flexibilities. Finally, Fig. 2.9a shows the optimal d-LMP from the LEM that enables the

Table 2.1: Summary financial metrics for our simulations under different types of market structures.

	SM + PM	PM only	No LEM
Avg. P d-LMP [$/kWh]	0.064	0.116	N/A
Avg retail tariff [$/kWh]	0.082	0.116	0.129

overall generation mix as shown in Fig. 2.8b and Fig. 2.9b, and that it is lower than what the 'without SMO' case predicts.

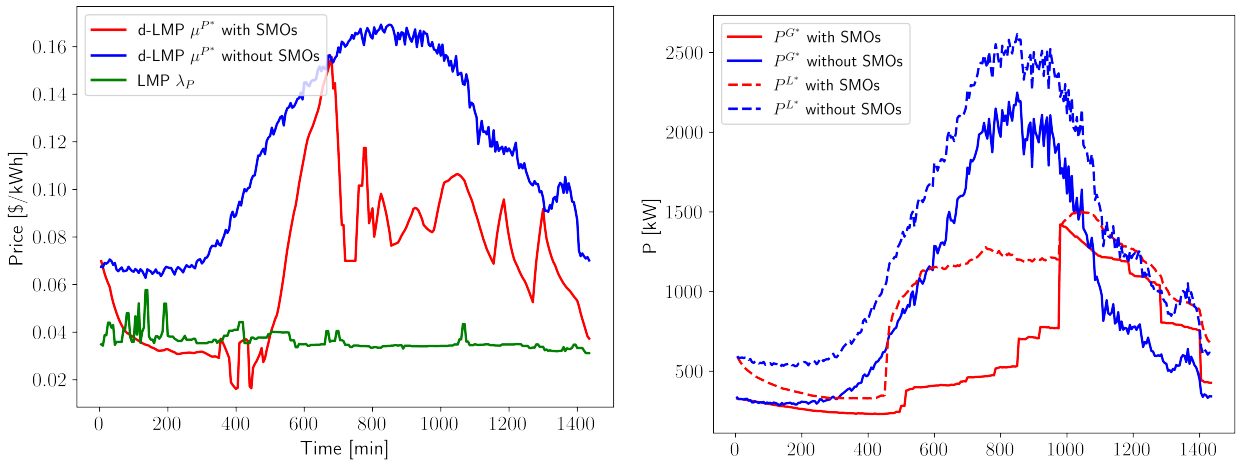

(a) d-LMPs averaged across all 79 primary nodes considered, with and without the SM, compared to the LMP from the WEM.

(b) PM total load and generation injections, summed over all primary feeder nodes except the slack bus.

Figure 2.9: Comparison of PM solutions obtained with and without SM.

2.5 Conclusion

A hierarchical local electricity market (LEM) structure was proposed in this chapter with a secondary market (SM) at the lower level representing secondary feeders and a primary market (PM) at the upper level, representing primary feeders, in order to effectively use DERs. The lower level SM enforces budget, power balance, and flexibility constraints and accounts for costs related to consumers, such as their disutility, flexibility limits, and commitment reliability, while the upper level PM enforces grid physics constraints such as power balance and capacity limits, and also minimizes line losses. The hierarchical LEM is evaluated using a modified IEEE-123 bus with high DER penetration, with each primary feeder consisting of several secondary feeders. Realistic power injections and load profiles were obtained over the course of 24 hours from GridLAB-D. The performance of the LEM was illustrated by delineating the family of power-injection profiles across the primary and secondary feeders as well as the corresponding local electricity tariffs that vary across the distribution grid.

It was shown that the overall LEM is capable of capturing fine-grain variations across the primary feeders and even further across secondary feeders so that the power injections and corresponding variable tariffs accurately charge or compensate DERs, capture consumer flexibilities, DER capabilities and constraints, as well as constraints and costs stemming from power physics.

The proposed hierarchical LEM represents the first step in formulating a market structure that allows disparate DER assets to participate and be appropriately compensated. Several other steps need to be executed to develop a complete retail market with various products. First, multiperiod extensions of our optimization frameworks at both the SM and PM levels need to be carried out. Next, the co-optimization of our SM and PM in the real-time market needs to be addressed, along with other markets for ancillary services and reserves. In doing so, we also hope to fully address the issues around settlement and billing in our hierarchical LEM, as well as the handling of unmet or reneged commitments in real-time. Also required is the development of advanced game-theoretic approaches that could be used by the SMO to generate its bids into the PM from the SM solutions, instead of the simple direct sum aggregation used at present. Similarly, the possibility of strategic bidding by DCAs in the SM and methods to counter this need to be examined as well. This will also help guide the design of a consumer-level market within each secondary feeder, forming the final tier of the proposed hierarchical LEM. Finally, more realistic distribution-level test cases and datasets are planned to be developed to validate the overall LEM, leveraging both simulations as well as real-world data [65, 71].

Chapter 3

Application of transactive framework for grid services and pricing

We propose a hierarchical local electricity market (LEM) at the primary and secondary feeder levels in a distribution grid to optimally coordinate and schedule distributed energy resources (DERs) and provide valuable grid services like voltage control. At the primary level, we use a current injection (CI)-based model that is valid for both radial and meshed, balanced and unbalanced, multi-phase systems. The primary and secondary markets leverage the flexibility offered by DERs to optimize grid operation and maximize social welfare. Numerical simulations on an IEEE-123 bus modified to include DERs show that the LEM successfully achieves voltage control and reduces overall network costs while also allowing us to decompose the price and value associated with different grid services so as to accurately compensate DERs.

3.1 Introduction and Motivation

With the increasing penetration of DERs such as renewables, storage, and flexible loads in the distribution system, it is critical to design market structures that enable their smooth integration at the grid edge - to balance variable supply and demand and increase the utilization of clean, renewable energy sources while maintaining affordability, reliability, and resilience. LEMs have the potential to empower consumers to take control of their energy footprint, allow transactive energy trading among members of a community, improve community resilience against wider grid events, and potentially reduce energy bills. This paper focuses on one such market, a LEM at the retail level. This market is designed for DER-rich distribution systems and includes a hierarchical structure, so as to facilitate DER integration, increase market participation of customers and prosumers, and utilize their assets to provide valuable grid services.

Local energy markets have the capability to allow electricity prices to be endogenous quantities rather than being imposed exogenously. In such a marketplace, prosumers can buy and sell energy in an open marketplace, or through an operator [72]. The normal practice adopted by grid operators and utilities is to rely on standard load profiles from historical data. This is challenged because of the intermittent and highly variable nature of the generation

from solar photovoltaic (PV) panels, demand from electric vehicles (EVs), and the needs of other DERs, as they can cause unpredictable swings in demand and/or generation. LEMs have the potential to help solve this problem for energy retailers and other grid management entities by offering flexibility services and creating opportunities for new business models. The LEM we propose has a two-tier structure, with a primary market (PM) at the upper level and a secondary market (SM) at the lower level. The SM consists of DER-coordinated assets located at each secondary feeder bidding into the market. These DCAs could consist of rooftop PVs, battery storage, or flexible loads. They could have multiple independent owners and thus may not allow a single agent to represent all of them in an aggregated manner. A secondary market operator (SMO) oversees market operations at this level to clear and schedule DCAs. At the upper level, the SMO has a dual role as an agent in the PM, as they are at a node in the primary feeder. All of these SMO agents bid into the PM. These are, in turn, cleared and scheduled by a primary market operator (PMO), which represents an entire primary feeder.

In our earlier work, we showed the market bidding and clearing process for such an LEM [73]. The distribution grid therein, however, was simple, radial, and balanced. In this paper, we relax these assumptions to model unbalanced, multi-phase, and meshed networks using a current-injection-based linear model for solving AC optimal power flow (OPF) at the primary level, employing McCormick envelopes (MCE) convex relaxations. More importantly, we show in this paper that our LEM can provide valuable services of Volt-VAr control (VVC) and voltage regulation in an unbalanced distributed grid. A distributed proximal atomic coordination (PAC) algorithm is used for PM clearing, which preserves privacy, reduces communication requirements, and improves computational tractability. We also introduce 3-phase pricing at both the SM and PM, to motivate how we can determine the value of such grid services in real-time energy markets based on an optimization framework. We disaggregate the distribution-locational marginal prices (d-LMPs) and local retail tariffs among different SMOs and DCAs, respectively, and decompose their components arising from economic objectives like maximizing social welfare and minimizing costs versus grid objectives like minimizing line losses and voltage profile deviations, while satisfying several power flow constraints. We show that the resulting LEM leads to effective VVC, along with efficient pricing and market compensation through spatial-temporal price differentiation.

3.1.1 Related work

Several works in the literature have proposed local energy markets for the future power grid with high DER penetration [34, 35, 38, 47] and transactive energy systems more broadly. However, these have not comprehensively studied the potential for such retail electricity markets to provide distribution grid services. Prior work on market design for grid services has largely focused on transmission for ancillary services like frequency regulation. Some have considered the possibility of active distribution systems to provide services like voltage support to the transmission grid [74] or the use of reactive power (VAr) assets for voltage balancing and regulation [75]. However, they have not studied how these services interact with market structures or how they would be priced. Accurate pricing mechanisms and incentives are critical to coordinate distributed agents and prosumers, where centralized dispatch or direct control is not possible.

While some recent works have explored distribution-level pricing, they have focused either solely on P [76–78] or notions of Q pricing [79]. A few other works have considered pricing for other grid services like voltage control [80], but these have relied on using simple power flow models like *LinDistFlow* which ignore line losses, or second-order conic program-based convex relaxations (DistFlow) that are restricted to radial, balanced networks [81]. These approximations are generally not valid for low to medium-voltage distribution systems, which have significant losses and are generally unbalanced and multi-phase.

3.1.2 Our contributions

We propose the use of our novel LEM architecture for providing grid services, specifically voltage regulation (or Volt-VAR control) in distribution grids. We build upon our prior work that proposed the hierarchical LEM structure consisting of the SM and PM [73], where we used a nonlinear DistFlow (branch flow) model for solving ACOPF at the primary level based on a second-order conic convex relaxation [76]. However, this was restricted to radial and balanced systems. We now extend this market structure to unbalanced, multi-phase, and meshed networks by instead using a Current-Injection based linear model for ACOPF and PM clearing [82]. We also introduce notions of three-phase pricing at both the SM and PM levels, and motivate how we can determine the value of such grid services in real-time energy markets based on an optimization framework, in order to accurately compensate DERs and prosumers. This also extends prior work on hierarchical hybrid Volt-VAR control [76] by coordinating DERs through market mechanisms and price-based transactive control rather than via direct control of these devices by the grid operator. Finally, we disaggregate the d-LMPs and local retail tariffs among different SMOs and DCAs, respectively, and break down their components arising from economic objectives like maximizing social welfare versus grid objectives like minimizing line losses and voltage profile deviations.

3.1.3 Our contributions

We propose a novel LEM for real-time energy transactions in a distribution grid with high DER penetration. The following are its unique features and the key contributions of our work:

- Hierarchical LEM structure that is electrically collocated with the current distribution network.
- Effectively address multiple functions of the distribution grid by virtue of the proposed LEM's hierarchy, with grid physics considered in the PM and consumer needs and constraints in the SM.
- Systematic approach showing how DERs, along with their flexibility bids, can be coordinated and aggregated in real-time via DCAs, and how these aggregated entities can participate in retail and/or wholesale electricity markets.
- Optimization of multiple objectives, including commitment reliability, net cost, flexibility, and utility in the SM; and net costs, utility, and line losses in the PM.

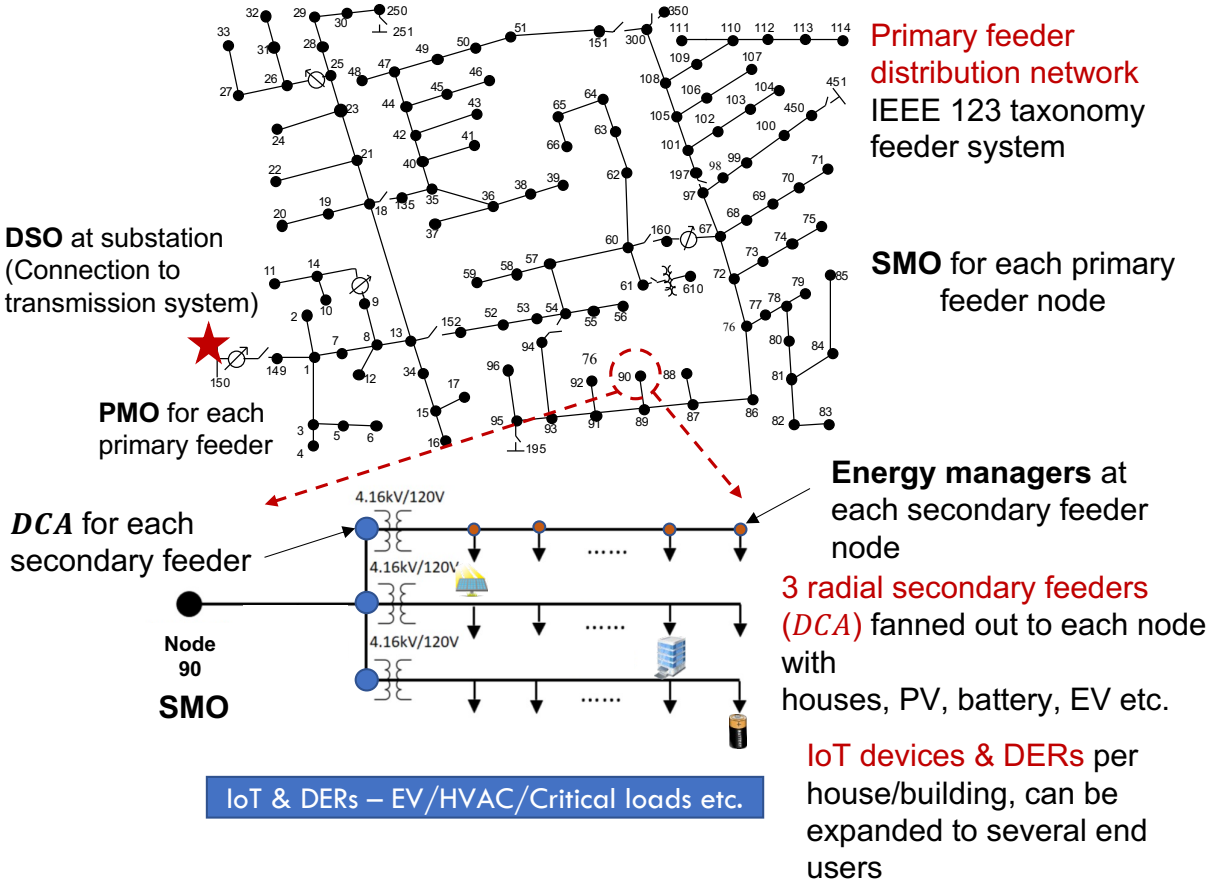

Figure 3.1: Hierarchical LEM co-located with distribution grid.

- Accommodation of disparate constraints, including budget, power balance, capacity, and flexibility limits in the SM; and nonlinear ACOPF constraints in the PM.
- Generation of a novel commitment score aimed at improving the reliability of our LEM by tracking the performance of DCAs at the secondary level.
- Fine-grain pricing in the form of local retail tariffs in the SM and d-LMPs in the PM that vary with both location and time, allowing for more efficient allocations in terms of lower costs and accurate compensation of DERs. This, in turn, provides an alternative to the current practices of net-metering and/or direct participation of DERs through aggregators at the wholesale energy market (WEM) level (under FERC 2222).
- Validation of the entire LEM using a modified IEEE 123-bus, with a high penetration of DERs and flexible loads.

The rest of the chapter is organized as follows. In Section 3.2 we introduce the structure of the LEM, including the lower level SMO, upper level PMO, and the interactions between the PMO, SMO, and WEM. In Section 3.3 we present numerical results on a modified IEEE-123 test feeder, with high levels of PV penetration and load flexibility. In Section 3.4 we provide concluding remarks and topics for future work.

3.2 Our methodology

We use the LEM structure proposed in our previous work [73], with an SM at the lower level and a PM at the upper level, as in Fig. 3.2. These markets are operated by an SMO and PMO, respectively, with combined oversight of both by a distribution system operator. The upper-level PMO coordinates with the WEM, while the lower-level SMO oversees the DCAs. Both the SM and PM clearing use flexibility bids submitted by the DCA and SMO, respectively. Fig. 3.2b illustrates flexibility bids from a DCA and SMO, in the SM and PM, respectively. The SM clearing results in a revised flexibility range for each DCA, while the PM clearing results in setpoints for each SMO. These solutions are used to set bilateral contracts between SMOs and DCAs (in the SM) and the PMO and SMOs (in the PM). The following quantities are used to define our SM optimization. For simplicity, we have ignored subscripts and superscripts that indicate the SMO and DCA being considered. For e.g., $P_j^{i,\phi}$ corresponds to DCA j under SMO i while $P_{i,\phi}$ corresponds to SMO i, for phase ϕ.

- $j \in \mathcal{N}_{J,i}$: Set of indices of all DCAs under SMO i.
- P^0, Q^0: Baseline active and reactive power injections.
- $\Delta P = [\underline{P}, \overline{P}$, $\Delta Q = [\underline{Q}, \overline{Q}]$: Bid flexibilities for each DCA, giving the range of maximum downward and upward flexibilities in P and Q injections offered by the DCA.
- t_p and t_s: Timestamps for the PM and SM, respectively.
- $\Delta t_p, \Delta t_s$ and Δt_{WEM}: Time periods for the PM, SM, and WEM, respectively. Here, we assumed that the SM is cleared more frequently, every $\Delta t_s = 1$ minute, while the PM is cleared every $\Delta t_p = 5$ minutes.
- $\hat{t}_p$: PM clearing prior to current SM interval $[t_s, t_s + \Delta t_s]$.
- $P^*(\hat{t}_p)$, $Q^*(\hat{t}_p)$: Setpoints provided by the PM to SMOs.

3.2.1 Secondary market model

The SMO i solves the multiobjective optimization problem shown in Eq. (3.1) to schedule the DCAs j based on their flexibility bids. We model both net load and net generator DCAs at the secondary feeder level separately, indicated by superscripts $K \in \{L, G\}$, respectively. All power injections are three-phase (and possibly unbalanced) variables for each DCA or secondary feeder. Terms without a superscript K, G, or L refer to net injections (generation minus load), i.e., $P_j^i = P_j^{iG} - P_j^{iL}$. All variables are specified for the current secondary timestep t_s unless stated otherwise.

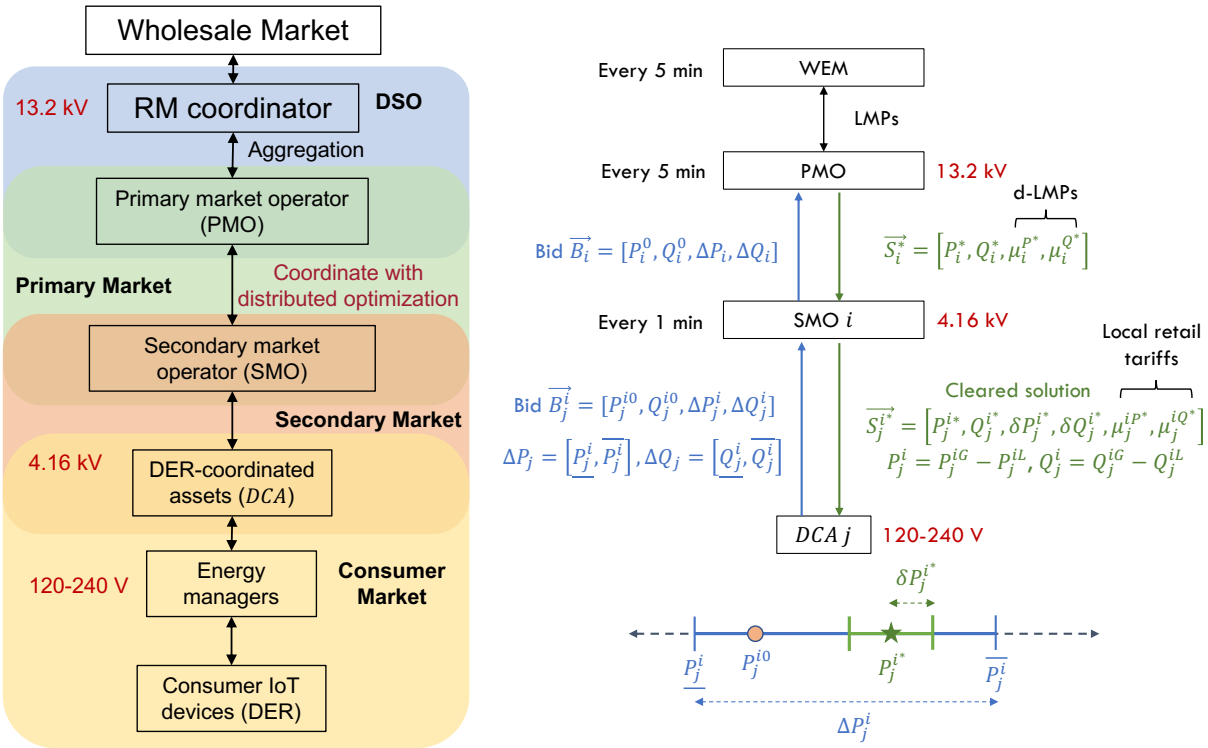

(a) LEM agents.

(b) LEM bids and solutions.

Figure 3.2: Overall structure of the hierarchical LEM.

$$\min \sum_{j \in \mathcal{N}_{J,i}} \{f_{j,1}^i, f_{j,2}^i, f_{j,3}^i, f_{j,4}^i\} \tag{3.1a}$$

$$f_{1,j}^i \succ f_{2,j}^i \succ f_{3,j}^i \succ f_{4,j}^i, \quad K \in \{L, G\}, \quad \Phi = \{a, b, c\} \tag{3.1b}$$

$$f_{j,1} = -C_j^i \left(\sum_{\phi \in \Phi} (P_j^{iK,\phi} - P_j^{iK0,\phi})^2 + (Q_j^{iK,\phi} - Q_j^{iK0,\phi})^2 \right)$$

$$f_{j,2}^i = \sum_{\phi \in \Phi} \mu_j^{iP,\phi} P_j^{i,\phi} + \mu_j^{iQ,\phi} Q_j^{i,\phi}, \; f_{j,3}^i = -\sum_{\phi \in \Phi} (\delta P_j^{iK,\phi} + \delta Q_j^{iK,\phi})$$

$$f_{j,4}^i = \beta_j^{iP} \sum_{\phi \in \Phi} \left(P_j^{iK} - P_j^{iK0}\right)^2 + \beta_j^{iQ} \sum_{\phi \in \Phi} \left(Q_j^{iK} - Q_j^{iK0}\right)^2$$

subject to:

$$P_j^{iK,\phi} - \delta P_j^{iK,\phi} \geq \underline{P}_j^{iK,\phi} \; Q_j^{iK,\phi} - \delta Q_j^{iK,\phi} \geq \underline{Q}_j^{iK,\phi} \tag{3.1c}$$

$$P_j^{iK,\phi} + \delta P_j^{iK,\phi} \leq \overline{P}_j^{iK,\phi}, \; Q_j^{iK,\phi} + \delta Q_j^{iK,\phi} \leq \overline{Q}_j^{iK,\phi} \tag{3.1d}$$

$$\delta P_j^{iK,\phi}, \; \delta Q_j^{iK,\phi} \geq 0, \; 0 \leq \mu_j^{iP} \leq \overline{\mu}^{iP}, 0 \leq \mu_j^{iP} \leq \overline{\mu}^{iQ} \tag{3.1e}$$

$$\sum_{t_s}^{t_s + \Delta t_p} \sum_{j \in \mathcal{N}_{J,i}} \sum_{\phi \in \Phi} \left(\mu_j^{iP,\phi}(t) P_j^{i,\phi}(t) + \mu_j^{iQ,\phi}(t) Q_j^{i,\phi}(t) \right) \Delta t_s$$

$$\leq \sum_{\phi \in \Phi} \left(\mu_i^{P^*,\phi}(\hat{t}_p) P_i^{\phi^*}(\hat{t}_p) + \mu_i^{Q^*,\phi}(\hat{t}_p) Q_i^{\phi^*}(\hat{t}_p) \right) \Delta t_p \tag{3.1f}$$

$$\sum_{j \in \mathcal{N}_{J,i}} P_j^{i,\phi}(t_s) = P_i^{\phi^*}(\hat{t}_p), \quad \sum_{j \in \mathcal{N}_{J,i}} Q_j^{i,\phi}(t_s) = Q_i^{\phi^*}(\hat{t}_p) \tag{3.1g}$$

The following decision variables are determined as outputs of the optimization for each DCA j bidding to SMO i, determined at t_s and applied over the period $[t_s, t_s + \Delta t_s]$:

- $P_j^{i,\phi}, Q_j^{i,\phi}$: Optimal power injection setpoints.
- $[\delta P_j^{i,\phi}, \delta Q_j^{i,\phi}]$: Optimal symmetric flexibility ranges around setpoints, i.e. DCAs are directed to have net injections within these intervals $[P_j^{i,\phi} - \delta P_j^{i,\phi}, P_j^{i,\phi} + \delta P_j^{i,\phi}]$.
- μ_j^{iP}, μ_j^{iQ}: Local electricity tariffs.

The commitment score $C_j^i(t) \in [0, 1]$ reflects the SMO's confidence in whether the DCA j will reliably follow their committed injections within the flexibility range specified above, with higher values C_j^i indicating more reliable assets. This is updated based on the deviations of actual DCA responses ($\hat{P}_j^i$) from their scheduled setpoints ($P_j^{i^*}$) and ranges ($[\underline{P}_j^{i^*}, \overline{P}_j^{i^*}]$) in the SM (see [73] for details).

The cost functions in (3.1) correspond to the following:

1. **Commitment reliability $\mathbf{f_{j,1}^i}$**: Maximize flexibility of injections assigned to more trustworthy DCAs (i.e., C_j^i closer to 1) while minimizing the flexible scheduling of DERs with lower commitment scores, who are less likely to abide by their contractual commitments.
2. **Net costs $\mathbf{f_{j,2}^i}$**: Minimize *net* costs to the SMO for running its SM, comprising payments from the SMO to DCAs that are net generators, and denote revenue from the DCAs that are loads.
3. **Flexibility $\mathbf{f_{j,3}^i}$**: Maximize aggregate flexibility that the SMO can extract from its DCAs, and offer to the PMO.
4. **Disutility $\mathbf{f_{j,4}^i}$**: Minimize inconvenience to DCAs when they provide flexibility to the operator. Thus, our SMO is an altruistic entity that also considers welfare maximization for its DCAs. For our simulations, disutility coefficients were $\beta_j^{iP}, \beta_j^{iQ} \sim \mathcal{U}[0.1, 1]$.

The constraints in Eqs. (5.3c) and (5.3d) reflect the feasible power injections for each of the DCAs, which are determined by their capacity limits and flexibility bids. Note that while generators can submit both upward and downward flexibilities from their baseline injection values, loads can only offer downward flexibility, i.e. $\overline{P}_j^{iL} = P_j^{i0}$. Constraints in Eq. (5.3e) enforce that the cleared tariffs μ_j^i for any DCA j cannot exceed a price ceiling $\overline{\mu}_j^i$. The budget constraint Eq. (5.3f) ensures that the total payments by the SMO to its DCAs for all SM clearings within each primary interval are less than its net revenue received from the PMO during the same period, in order to remain solvent. Notice that all the terms in the objective function are either linear or convex, except for the net bilinear cost term $f_{j,2}^i$, which is not convex in general. Similarly, all the constraints are linear except for the bilinear budget balance inequality constraint Eq. (5.3f). In our prior work [73], we solved the full nonconvex optimization problem in order to obtain retail tariffs directly as outputs of the SM clearing. However, in this work, we consider a simpler, convexified version of the problem. We achieve this by removing the budget constraint Eq. (5.3f) from the original problem and

instead enforcing exact budget balance ex-post by setting the retail tariffs for each of the DCAs after the fact. This also implies that the SMO is no longer explicitly minimizing costs in the SM (i.e., we remove the $f^i_{j,2}$ term from the objective) and is instead just breaking even, essentially acting as a purely non-profit entity and market maker. By transforming the inequality constraint in Eq. (5.3f) to a strict equality, the SMO can derive the following localized real-time retail tariffs for each DCA j at every secondary timestep t_s using their cleared P and Q schedules:

$$\mu^{iP}_j(t_s) = y^{iP}_j(t_s)\frac{|R_{PM}|}{|P^{i^*}_j(t_s)|\Delta t_s},\ \mu^{iQ}_j(t_s) = y^{iP}_j(t_s)\frac{|R_{PM}|}{|Q^{i^*}_j(t_s)|\Delta t_s}$$
$$R_{PM} = \left(\mu^{P^*}_i(\hat{t}_p)P^*_i(\hat{t}_p) + \mu^{Q^*}_i(\hat{t}_p)Q^*_i(\hat{t}_p)\right)\Delta t_p \tag{3.2}$$

Here R_{PM} is the net revenue to the SMO i from the most recent PM clearing, while P^*, Q^* indicate net injections. In order to balance its budget, the SMO sets the price multipliers $y^{iP}_j(t_s)$, $y^{iQ}_j(t_s)$ for each of the DCAs and at each time t_s using the following heuristics, which can be derived from the exact budget balance equality constraint. If the DCA j is a net generator in terms of P or Q injections (summed over all its non-zero phases) i.e. $j \in S^{P,Q}_G(t_s) = \{j \in \mathcal{N}_{J,i} : P^{i^*}_j(t_s) > 0 \text{ or } Q^{i^*}_j(t_s) > 0\}$, the price multipliers are:

$$y^{iP,Q}_j(t_s) = \begin{cases} 0, & \text{if } |S^{P,Q}_G(t_s)| = 0 \\ \frac{1}{2|S^{P,Q}_G(t_s)|} & \text{if } |S^{P,Q}_L(t_s)| = 0 \\ 1 & \text{otherwise} \end{cases} \tag{3.3}$$

If DCA j is a net load i.e. $j \in S^{P,Q}_L(t_s) = \{j \in \mathcal{N}_{J,i} : P^{i^*}_j(t_s) < 0 \text{ or } Q^{i^*}_j(t_s) < 0\}$, the price multipliers are:

$$y^{iP,Q}_j(t_s) = \begin{cases} 0, & \text{if } |S^{P,Q}_L(t_s)| = 0 \\ \frac{1}{2|S^{P,Q}_L(t_s)|} & \text{if } |S^{P,Q}_G(t_s)| = 0 \\ \frac{1+2|S^{P,Q}_G(t_s)|}{2|S^{P,Q}_L(t_s)|} & \text{otherwise} \end{cases} \tag{3.4}$$

where $|S^{P,Q}_{G,L}(t_s)|$ denotes the cardinality of the set of DCA generators or loads at time t_s. The retail tariffs are assumed to be identical across all phases for each DCA, and are set based on the net power injections summed over all its non-zero phases, i.e., $P^{i^*}_j(t_s) = \sum_\phi P^{i,\phi^*}_j(t_s)$. Note that all the above prices are solved for in terms of [\$/kWh] or [\$/kVARh]. Finally, Eq. (5.2g) denote power balance constraints for the SMO, requiring that injections from the DCAs downstream must equal the net flows $P^*_i(\hat{t}_p)$ and $Q^*_i(\hat{t}_p)$ from the primary feeder upstream, as scheduled by the PMO. Since the PM clears less often, these values can be treated as constant for the SM optimization over each Δt_p.

The multiobjective optimization problem Eq. (3.1) is solved using a hierarchical approach [49, 83] since our objective terms considered in Eq. (5.3a) have different units and may not be comparable in magnitude. The SMO ranks their objectives as shown in Eq. (5.3b), setting commitment reliability as their most important goal and disutility as being the least important. The SMO then solves a series of optimization problems, sequentially optimizing

these objectives one at a time, in descending order of importance. At each step, constraints are added on how much the previous objective value can be degraded, controlled by the parameter $\epsilon = 5\%$:

$$\min_{\vec{S}_j^i} F_k = \sum_{j \in \mathcal{N}_{J,i}} f_{j,k}^i(\vec{S}_j^i) \; \forall \; k = 1, 2, 3 \tag{3.5}$$

$$\text{s.t. } f_{j,\ell}^i(\vec{S}_j^i) \leq (1+\epsilon) \sum_{j \in \mathcal{N}_{J,i}} f_{j,\ell}^i(\vec{S}_j^{i^*}) = (1+\epsilon) F_\ell^*, \tag{3.6}$$

$$\forall \; \ell = 1, 2, \ldots, k-1, \; k > 1, \text{ and constraints } Eq. \text{ (5.3c)-(5.2g)}$$

3.2.2 Primary market model

Before each PM clearing, SMO i aggregates schedules of its DCAs j from the last SM clearing, and uses this to bid into the PM. Its bid is similar to the DCA, with baseline injections and flexibilities both determined by SM solutions:

$$P_i^0(t_p) = \sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*}(t_p), \; \Delta P_i = \left[\underline{P_i^\phi}, \overline{P_i^\phi}\right]$$
$$\left[\sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*} - \delta P_j^{i^*}, \sum_{j \in \mathcal{N}_{J,i}} P_j^{i^*} + \delta P_j^{i^*}\right] \tag{3.7}$$

The PMO accepts these bids and then clears the PM at every t_p at intervals of Δt_p. The PMO clears using a dual ascent-based distributed optimization algorithm, known as PAC, that facilitates market clearing [8]. Compared to centralized optimization, such a distributed approach helps preserve privacy, reduces the computational burden, and improves scalability for large networks since each SMO only needs to locally exchange information with neighboring agents, while still approaching globally optimal solutions. The PM optimization solves the ACOPF problem using a CI model as proposed in [82].

Current injection model

The primal decision variables for each SMO i obtained by solving the optimization problem $x = [P_i^\phi, Q_i^\phi, V_i^{\phi,R}, V_i^{\phi,I}, I_i^{\phi,R}, I_i^{\phi,I}]$ consists of (i) active ($P_i^{\phi^*}$) and reactive ($Q_i^{\phi^*}$) power setpoints (ii) real and imaginary components of nodal voltages ($V_i^{\phi,R^*}, V_i^{\phi,I^*}$) and current injections ($I_i^{\phi,R^*}, I_i^{\phi,I^*}$). Note that these are solved for each non-zero phase $\phi \in \mathcal{P} = \{a, b, c\}$.

The CI-OPF problem formulation is given by:

$$\min_x f^{obj}(x) \tag{3.8a}$$
$$I^R = \mathrm{Re}(\mathrm{Y}V),\ I^I = \mathrm{Im}(\mathrm{YV}) \tag{3.8b}$$
$$P_i^\phi = V_i^{\phi,R} I_i^{\phi,R} + V_i^{\phi,I} I_i^{\phi,I} \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{3.8c}$$
$$Q_i^\phi = -V_i^{\phi,R} I_i^{\phi,I} + V_i^{\phi,I} I_i^{\phi,R} \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{3.8d}$$
$$(I_{ij}^R)^2 + (I_{ij}^I)^2 \leq \overline{I_{ij}}^2 \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{3.8e}$$
$$\underline{V_i^\phi}^2 \leq (V_i^{\phi,R})^2 + (V_i^{\phi,I})^2 \leq \overline{V_i^\phi}^2 \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{3.8f}$$
$$\underline{P_i^\phi} \leq P_i^\phi \leq \overline{P_i^\phi},\ \underline{Q_i^\phi} \leq Q_i^\phi \leq \overline{Q_i^\phi} \tag{3.8g}$$

where Y is the 3-phase bus admittance matrix for the network, and V and I are matrices of nodal voltages and currents, respectively. Problem Eq. (5.8) is nonconvex due to bilinear constraints Eqs. (5.8c) and (5.8d), and the ring constraint Eq. (5.8f) on voltage magnitudes. We obtain a convex relaxation by using McCormick envelopes (MCE), which represent the convex hull of a bilinear product $w = xy$ by using upper and lower limits on x, y. Thus, we replace the bilinear equality with a series of linear inequalities, denoted as $\mathrm{MCE}(w) = \{w = xy : x \in [\underline{x}, \overline{x}], y \in [\underline{y}, \overline{y}]\}$:

$$MCE(w, \underline{x}, \overline{x}, \underline{y}, \overline{y}) = \begin{cases} w \geq \underline{x}y + x\underline{y} - \underline{x}\underline{y} \\ w \geq \overline{x}y + x\overline{y} - \overline{x}\overline{y} \\ w \leq \underline{x}y + x\overline{y} - \overline{x}\underline{y} \\ w \leq \overline{x}y + x\underline{y} - \underline{x}\overline{y} \end{cases} \tag{3.9}$$

We introduce auxiliary variables for each of the four bilinear terms $\{a_i^\phi, b_i^\phi, c_i^\phi, d_i^\phi\} = \{V_i^{\phi,R} I_i^{\phi,R}, V_i^{\phi,I} I_i^{\phi,I}, V_i^{\phi,R} I_i^{\phi,I}, V_i^{\phi,I} I_i^{\phi,R}\}$ allowing us to convert constraints Eqs. (5.8c) and (5.8d) to linear constraints with MCE constraints on each of the auxiliary variables. We also need additional constraints on the nodal current injections and nodal voltages in order to define the MCE constraints. These voltage and current bounds can be determined by applying a suitable preprocessing method using the nodal P and Q limits from the SMO bids. The resulting bounds will also implicitly satisfy constraints Eq. (5.8e) and Eq. (5.8f). Thus, we can replace constraints Eqs. (5.8c) to (5.8f) with the following set of constraints in order to obtain the relaxed CI-OPF problem, which reduces to a linear program that can be solved easily. However, we do incur the overhead of computing the tightest possible V and I bounds to obtain a good convex relaxation, which in turn ensures

that the relaxed solutions are feasible for the original problem.

$$P_i^\phi = a_i^\phi + b_i^\phi, \;\; Q_i^\phi = -c_i^\phi + d_i^\phi \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{3.10a}$$

$$\underline{I_i^{\phi,R}} \leq I_i^{\phi,R} \leq \overline{I_i^{\phi,R}}, \; \underline{I_i^{\phi,I}} \leq I_i^{\phi,I} \leq \overline{i^{\phi,I}} \tag{3.10b}$$

$$\underline{V_i^{\phi,R}} \leq V_i^{\phi,R} \leq \overline{V_i^{\phi,R}}, \; \underline{V_i^{\phi,I}} \leq V_i^{\phi,I} \leq \overline{V_i^{\phi,I}} \tag{3.10c}$$

$$a_i^\phi \in MCE(V_i^{\phi,R} I_i^{\phi,R}, \underline{V_i^{\phi,R}}, \overline{V_i^{\phi,R}}, \underline{I_i^{\phi,R}}, \overline{I_i^{\phi,R}}) \tag{3.10d}$$

$$b_i^\phi \in MCE(V_i^{\phi,I} I_i^{\phi,I}, \underline{V_i^{\phi,I}}, \overline{V_i^{\phi,I}}, \underline{I_i^{\phi,I}}, \overline{I_i^{\phi,I}}) \tag{3.10e}$$

$$c_i^\phi \in MCE(V_i^{\phi,R} I_i^{\phi,I}, \underline{V_i^{\phi,R}}, \overline{V_i^{\phi,R}}, \underline{I_i^{\phi,I}}, \overline{I_i^{\phi,I}}) \tag{3.10f}$$

$$d_i^\phi \in MCE(V_i^{\phi,I} I_i^{\phi,R}, \underline{V_i^{\phi,I}}, \overline{V_i^{\phi,I}}, \underline{I_i^{\phi,R}}, \overline{I_i^{\phi,R}}) \tag{3.10g}$$

Objective functions for voltage control

After converting all quantities to per-unit (p.u.), we considered a weighted linear combination of several convex objective functions for the PM clearing using CI-OPF, where the weight ξ controls the relative tradeoff between the first 2 'socio-economic' objectives versus the last 2 'electrical' objectives:

$$\begin{aligned} f^{obj}(x) = &\sum_{\phi \in \mathcal{P}} \sum_{i \in \mathcal{N}_I} \left[f_i^{\text{Load-Disutil},\phi}(x) + f_i^{\text{Gen-Cost},\phi(x)} \right] \\ &+ \xi \sum_{\phi \in \mathcal{P}} \left[\sum_{(i,k) \in \mathcal{T}} f_{ik}^{\text{Loss},\phi}(x) + \sum_{i \in \mathcal{N}_I} f_i^{\text{Volt},\phi}(x) \right] \end{aligned} \tag{3.11}$$

The first term minimizes disutility due to load flexibility:

$$f_i^{\text{Load-Disutil},\phi}(x) = \beta_i^P (P_i^{L,\phi} - P_i^{L0,\phi})^2 + \beta_i^Q (Q_i^{L,\phi} - Q_i^{L0,\phi})^2$$

The second term minimizes generation costs. These are set by the locational marginal price (LMP) Λ_i^P, Λ_i^Q for the primary feeder node at the point of common coupling (PCC) at the substation. For SMOs at all other primary feeder nodes, these depend on some fixed coefficients α_i^P, α_i^Q that represent costs to the SMO for running its SM:

$$f_i^{\text{Gen-Cost},\phi}(y) = \begin{cases} \Lambda_i^P P_i^{G,\phi} + \Lambda_i^Q Q_i^{G,\phi}, \text{if } i & \text{is PCC} \\ \alpha_i^P P_i^{G,\phi} + \alpha_i^Q Q_i^{G,\phi} & \text{otherwise} \end{cases}$$

The third term minimizes line losses in the network for more efficient operation. These are determined by the

$$f_{ij}^{\text{Loss},\phi}(x) = R_{ij} |I_{ij}^\phi|^2 = R_{ij} \left(I_{ij}^{\phi,R^2} + I_{ij}^{\phi,I^2} \right)$$

where $\mathcal{T}$ is the set of network branches, R_{ij} are branch resistances and I_{ij}^ϕ are branch current flows. These can be readily obtained from the nodal currents I_i since $I = A^\intercal I_{branch}$, where A is the three-phase graph incidence matrix.

The fourth term is the voltage regulation term that we specify to perform voltage control. This penalizes voltage deviations from some desired nominal values, in order to achieve a desired profile:

$$f_i^{\text{Volt},\phi}(x) = \left(V_j^{\phi,R} - \tilde{V}_j^{\phi,R}\right)^2 + \left(V_j^{\phi,I} - \tilde{V}_j^{\phi,I}\right)^2$$

In this study, we regulated voltage about setpoints $\tilde{V}_j^{\phi,R} = 1,\ \tilde{V}_j^{\phi,I} = 0$, to track a nominal magnitude $|\tilde{V}_j^{\phi}| = 1$ p.u.

3.2.3 Pricing

Both the SM and PM result in localized, real-time prices for each DCA and SMO, respectively, which allows us to capture the high degree of spatial and temporal variation in prices. In this study, we focus on the pricing results for SMOs in the PM - please refer to [73] for detailed results on localized retail tariffs for DCAs in the SM. We can derive PM prices by inspecting dual variables (λ) corresponding to different sets of linear equality constraints in the PM CI-OPF problem Eq. (5.8). The Lagrangian for the primal problem Eq. (5.8) is:

$$\begin{aligned}\mathcal{L} &= f^{obj}(x) + \lambda_P^\intercal P_{balance} + \lambda_Q^\intercal Q_{balance} \\ &\quad + \lambda_I^\intercal (I - YV) + \lambda_{ineq}^\intercal (RHS_{ineq} - LHS_{ineq})\end{aligned} \tag{3.12}$$

where $P_{balance}$ and $Q_{balance}$ refer to the active and reactive power balance equations Eqs. (5.8c) and (5.8d) respectively, and $I = YV$ enforces the linear Ohm's law constraint from Eq. (5.8b). The last term in the Lagrangian corresponds to all the remaining inequality constraints from Eq. (5.10b)-5.10g. However, our focus here is only on the duals of equality constraints Eqs. (5.8b) to (5.8d) for pricing purposes. Note that the dual variable λ_I is in terms of current, which can be converted to an equivalent value in terms of voltage:

$$\begin{aligned}&\lambda_I^\intercal (I - YV) \equiv \lambda_V^\intercal (ZI - V) = \lambda_V^\intercal (Y^{-1}I - Y^{-1}YV) \\ &= \lambda_V^\intercal Y^{-1}(I - YV) \implies \lambda_I^\intercal = \lambda_V^\intercal Y^{-1} \implies \lambda_V = Y^\intercal \lambda_I\end{aligned} \tag{3.13}$$

where $Z = Y^{-1}$ is the 3-phase network impedance matrix. These dual variables can be interpreted as prices for different services in the distribution grid. Thus, we propose the vector of dual variables above $\boldsymbol{\lambda} = [\lambda_P, \lambda_Q, \overline{\lambda_V}]$ as the d-LMP where $\overline{\lambda_V} = \text{Re}(\lambda_V)$ is the real part of the complex dual variable. In particular, λ_P and λ_Q represent the P and Q d-LMP components for active and reactive power. The P-dLMP or energy price λ_P is similar to the notion of in the transmission system and WEM. Such a structure of P and Q components in a d-LMP has also been proposed in [80], but we introduce the voltage support price $\overline{\lambda_V}$ in this paper for the first time. These d-LMPs represent the overall grid services from DERs by providing real power, reactive power, and voltage support. We note that $\overline{\lambda_V}$ can be interpreted as a price for voltage control or regulation, because it reflects the effects of perturbations in the Ohm's law constraint, on our objective function, as shown below:

$$\begin{aligned}&\frac{\partial \mathcal{L}}{\partial V} = \frac{\partial f^{obj}(x)}{\partial V} - \lambda_I^\intercal Y = \frac{\partial f^{obj}(x)}{\partial V} - \lambda_V \\ &\text{At optimality } \frac{\partial \mathcal{L}}{\partial V^*} = \frac{\partial f^{obj}}{\partial V^*} - \lambda_V = 0 \implies \frac{\partial f^{obj}}{\partial V^*} = \lambda_V^*\end{aligned}$$

Type	Number	Capacity
DERs	380	1,745.8 kVA (≈44%)
PVs	207	880.84 kVA
Batteries	173	865 kVA
Spot loads	85	3,985.7 kVA
Houses	1008	4-10 kW (variable)
Flexible loads	1-2 per house	10-50% flexibility (variable)

Table 3.1: Specifications of modified IEEE 123-node feeder.

Thus, $\overline{\lambda_V}$ intuitively represents the costs of satisfying voltage constraints on the distribution grid (in terms of degrading the objective) and can be interpreted as the value of this voltage control grid service. Similarly, λ_P and λ_Q are costs associated with meeting power balance.

3.3 Results and Discussion

3.3.1 Numerical simulations

We conducted a co-simulation of both the SM and PM on a modified IEEE-123 node feeder with high DER penetration comprising rooftop solar PV systems, batteries, and flexible loads. The specifications of the modified network are shown in Section 3.3.1. The network was simulated using GridLAB-D in order to obtain realistic profiles for baseline power injections of SMOs and DCAs, as well as primary-level nodal voltages. Synthetic flexibility bids were then generated by randomly assigning flexibilities between 10-30% for each of the DCAs. Simulations were conducted for a 24 h period, using weather data from San Francisco, CA, on August 2, 2022, along with 5-min LMP data from the California Independent System Operator (CAISO). The SM was cleared every 1 min, while the PM was cleared every 5 min, in lockstep with the WEM.

The workflow for the co-simulation is shown in **??**. In particular, we feed in the aggregated solutions from the SM clearing to form the SMOs' bids into the PM. These bids, which are in terms of active and reactive power flexibility ranges, are then preprocessed to give the corresponding V and I bounds needed for the MCE relaxation. The relaxed CI-OPF problem is then solved to clear the PM. We used the Gurobi solver for both the SM and PM optimization problems. In order to accelerate our simulations, we parallelized the SM clearing using MIT's Supercloud high-performance computing cluster [84] and Python's message passing interface (MPI). At every secondary timestep t_s (1 min), we solved the optimization problems for all 85 SMOs in parallel across multiple processors, making the problem much more computationally tractable and providing ≈ 80X speedup in solution runtimes. The d-LMPs and nodal voltage solutions are 3-phase variables, but in the following sections, we calculate their mean values averaged over all the non-zero phases that are present at each node.

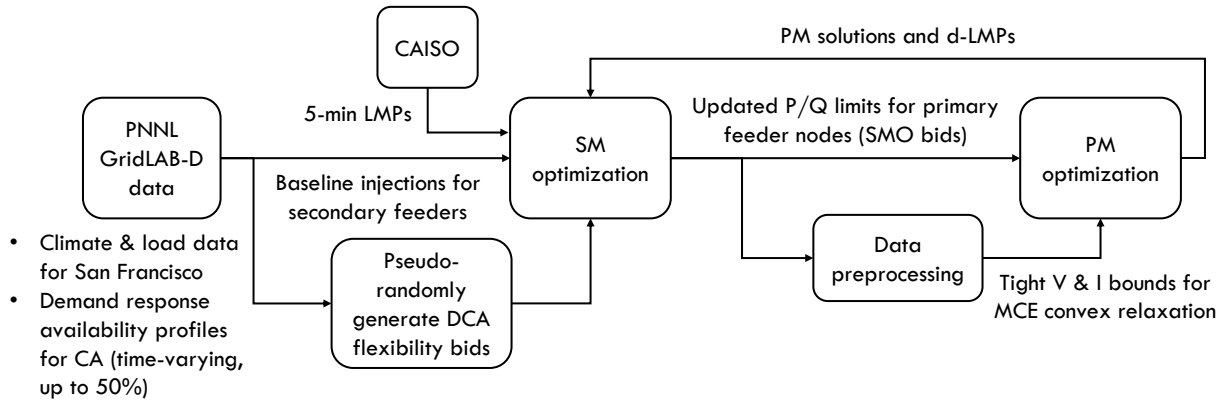

Figure 3.3: Workflow for SM and PM co-simulation.

3.3.2 Effects of the LEM on voltages

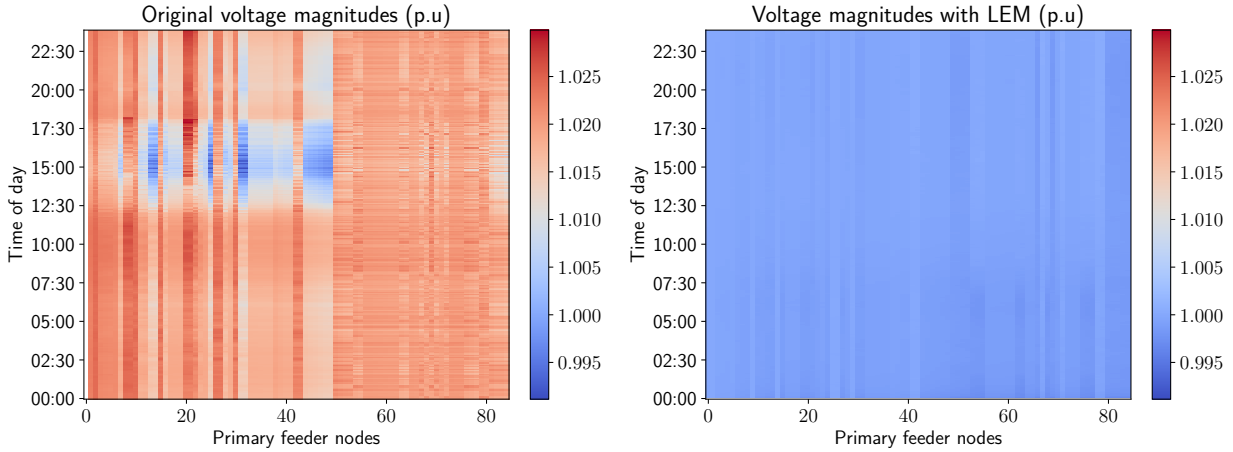

(a) Without the LEM.

(b) With the LEM.

Figure 3.4: Primary level nodal voltage magnitudes with and without the LEM, at nodes with SMOs and over time.

We find that the LEM does indeed significantly improve the overall voltage profile by making it more uniform and bringing the voltage magnitudes closer to the desired 1 p.u. setpoint, as seen in Fig. 3.4. In Fig. 3.4a we notice overvoltage (i.e. $|V| > 1$ p.u.) issues throughout most of the 24 h simulation period, but these are generally more pronounced during daylight periods of the day with higher PV output. Overvoltage problems are also more frequent and severe for specific primary nodes that correspond to SMOs and DCAs with greater local generation capacity from solar PV and/or batteries. Undervoltages (i.e., $|V| < 1$ p.u.) are less common and occur during the afternoons, likely due to higher demand spikes from heating, ventilation, and cooling (HVAC) loads. The LEM is able to effectively coordinate DERs in order to mitigate both under- and overvoltage issues throughout the day and across all nodes in the primary feeder, as seen in Fig. 3.4b. This is achieved through

smarter scheduling and dispatch of resources - these actions may include (but are not limited to) controlled battery charging or discharging, power factor control using smart inverters, as well as shifting or curtailment of flexible loads and appliances. This results in more uniform spatial and temporal voltage distributions.

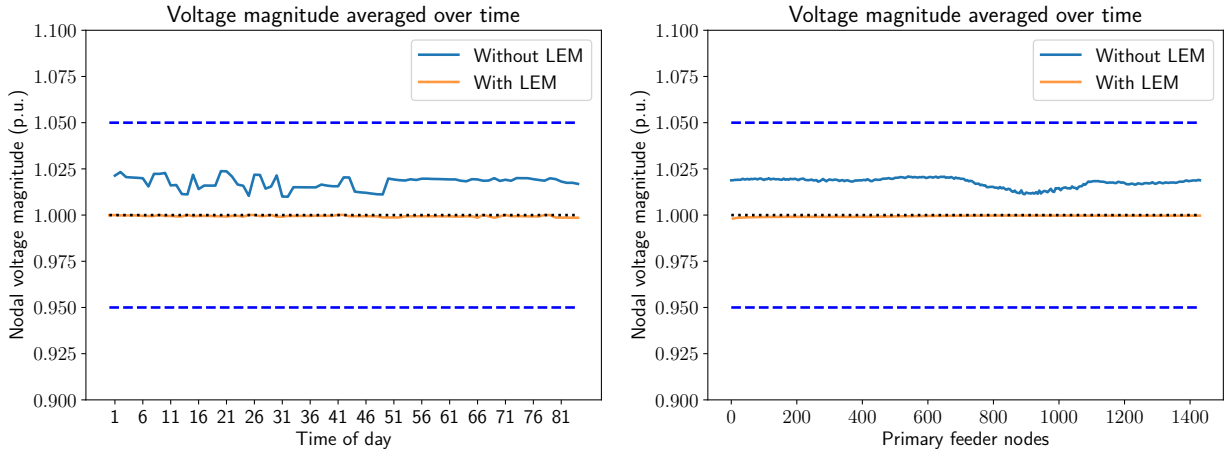

(a) Nodal averages over time.

(b) Daily average across nodes.

Figure 3.5: Primary level nodal voltage magnitudes with and without the LEM, at nodes with SMOs and over time.

The voltage profile improvements are also evident from Fig. 3.5, where both the spatial (in Fig. 3.5a) and temporal (in Fig. 3.5b) mean voltage magnitudes are almost exactly equal to the desired 1 p.u. with the LEM in place, as opposed to the consistently higher mean voltages observed without the LEM. The voltages are also well within the American National Standards Institute (ANSI) safe operating voltage limits of $[0.95, 1.05]$ $p.u.$.

3.3.3 dLMP results

Fig. 3.6 summarizes the PM pricing results and decomposition of the d-LMPs into the three components of P, Q, and V support prices. In Fig. 3.6a, temporal variations of the d-LMP components are shown over the whole day, when averaged over all the SMO nodes. At all times, the mean d-LMP over the primary feeder is higher than the LMP at the substation or PCC. This makes intuitive sense since the d-LMP accounts for additional costs and losses in the distribution grid downstream of the transmission grid, which are not included in the LMP. This also allows the DSO and PMOs to recoup their own costs for running the retail markets while participating in the WEM. Another interesting result is that throughout the day, the P and V-dLMP components contribute to the bulk of the d-LMP, while the Q-dLMP only makes up a small portion of the price. This makes sense since nodal Q injections are much smaller in magnitude compared to P injections across the distribution feeder, and is also in line with other works that have suggested, for instance, that Q-dLMPs should roughly be $\approx$ 10% of the corresponding P-dLMPs [68]. Another reason for the small Q price contribution could be that reactive compensation plays a key role in maintaining grid voltages, so some of its effects may already be taken into account by the V-dLMP.

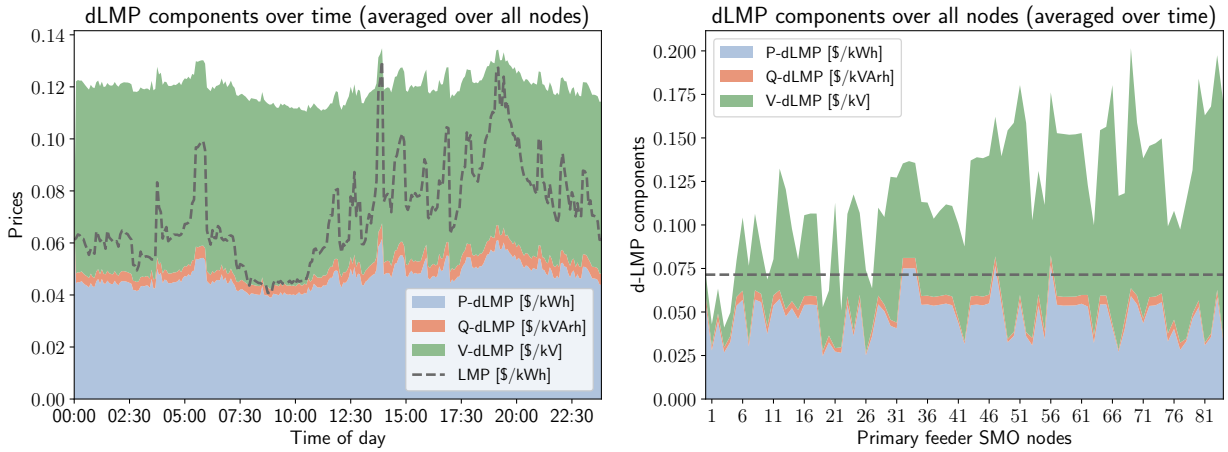

(a) d-LMP components over time, averaged over all primary level nodes along with the LMP.

(b) d-LMP components over primary level nodes, averaged over the entire 24 h simulation period.

Figure 3.6: Variations in d-LMPs for over nodes and time.

In Fig. 3.6b, the spatial node-to-node variations of the time-averaged dLMPs are shown, along with the average LMP for the day. We see again that the combined average P, Q, and V-dLMPs are higher than the average LMP at most nodes, except for a few of them (< 10). The relative breakdowns of P versus Q-dLMPs are roughly similar across the network, but the contributions of the V-dLMP differ quite significantly for different nodes. For e.g., the V-dLMP is relatively much larger for node 71 in Fig. 3.6b, indicating that it may be more challenging to meet grid physics constraints and support voltages at these specific nodes, while solving the PM clearing and CI-OPF problem. Further analysis is necessary to fully interpret and explain this trend, this will be explored more as part of future work. Both plots in Fig. 3.6 also show that the costs associated with voltage support are significant and must also be adequately accounted for in retail markets, rather than focusing solely on P and Q energy prices. In both Figs. 3.5a and 3.5b, the combined P and Q-dLMPs without including the V-dLMP are consistently lower than the LMP. This is in agreement with other related works such as [78, 80] - this indicates that distribution level costs involve not just those associated with satisfying power balance, but also other constraints like Ohm's law (Eq. (5.8b)).

The locational-temporal variations of the normalized P, Q, and V-dLMPs are shown in Fig. 3.7, for all 85 primary feeder nodes with SMOs and over the 24 h period. We observe that there's a great deal of variability in these prices, which further motivates the crucial need for new retail market structures such as our LEM in order to capture these variations. This would allow us to accurately compensate different resources depending on both the time of day as well as their geographic locations within the distribution system. Another important observation is that our combined d-LMP is significantly lower than the current retail rate charged by utilities and other load-serving entities (LSEs) throughout the day and across all primary nodes. Since current retail rates only include active power, we calculate an

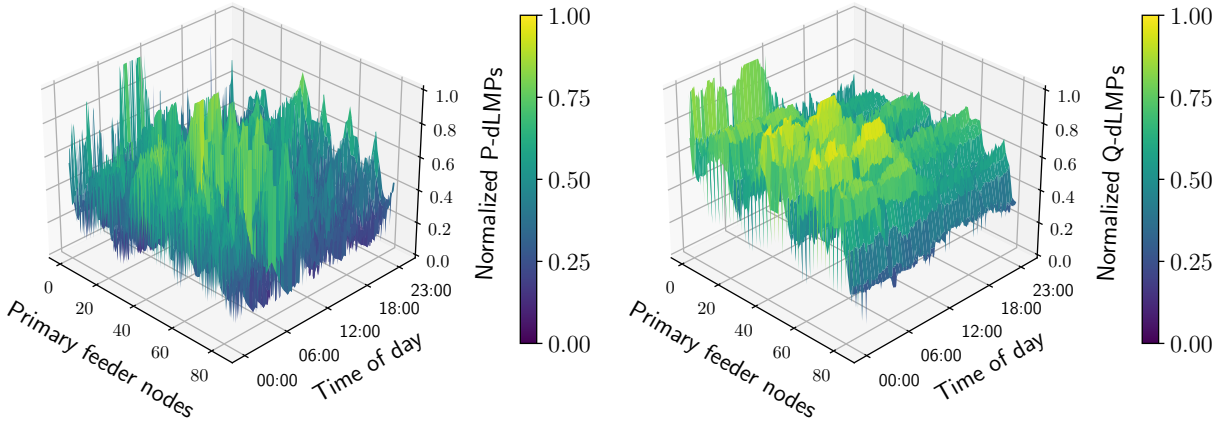

(a) P-dLMP

(b) Q-dLMP

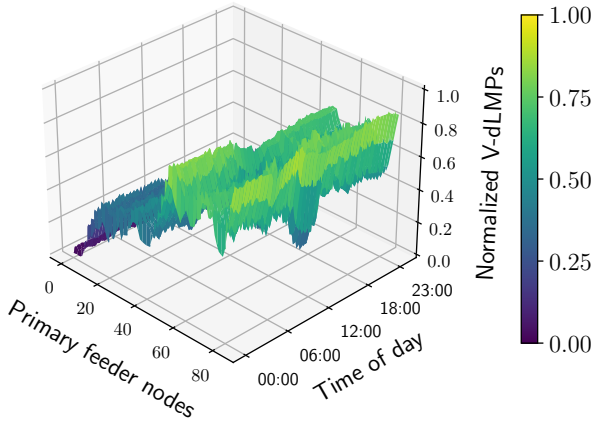

(c) V-dLMP

Figure 3.7: Distributions of d-LMP components over all SMO nodes and over the 24 h simulation period.

equivalent rate λ_{eq} in \$/kWh for our LEM as a weighted average of all 3 dLMP components:

$$\begin{aligned} \lambda_{eq} &= (\lambda_P^* P^* + \lambda_Q^* Q^* + \overline{\lambda_V}^* \Delta V^*)/P^* \\ \Delta V^* &= |V^{R^*} - 1| + |V^{I^*}| \end{aligned} \tag{3.14}$$

where ΔV^* are the deviations of voltages from the nominal values. The average bundled tariff for Pacific Gas & Electric (PG&E) customers in August 2022 was 33.72 ¢/kWh, compared to the mean equivalent rate $\overline{\lambda_{eq}} = 5.38$ ¢/kWh in our LEM, averaged over the day and the whole network. This represents a $\approx$ 84% reduction, indicating the LEM is able to coordinate and schedule DERs more effectively to reduce network-wide costs. These tariffs are likely to increase further as higher DER penetration places more stress on distribution grids, but our LEM can help mitigate these challenges [85].

However, it should also be noted that in this paper, we have only included costs for operating the primary market while meeting power flow constraints imposed by grid physics.

In reality, the DSO incurs additional costs such as maintenance costs, infrastructure expenses, and delivery charges, as well as profit margins imposed by LSEs. In addition, it has to recoup its costs for importing power from the WEM and transmission grid. For similar reasons, the retail rates charged by the SMOs to their DCAs may be higher than the breakeven tariffs determined by Eq. (3.2). These additional costs may reduce the chaptered margin of improvement from 64% to a certain extent. The final dLMPs and retail rates also represent the value provided by the PMOs and SMOs (as well as the DSO that oversees both) in terms of serving demand as well as by facilitating market participation for DERs. They allow DERs to actively bid into retail and wholesale markets and get appropriately compensated for the services they provide to the grid. This also motivates recent regulations like FERC order 2222, which opened up WEM participation to DERs [86], as well as the push towards performance-based rate regulation, which evaluates actual utility performance when establishing rates as an alternative to calculating rate plans based on utility capital investments [87].

3.4 Concluding Remarks

In this paper, we applied a hierarchical LEM to provide distribution grid services. We solved ACOPF at the primary level using a current injection model valid for balanced and meshed networks. We also accurately decomposed prices for such services by deriving dLMP components corresponding to P, Q, and voltage support. Numerical simulations show that the LEM successfully mitigates under- and over-voltage issues across the network throughout the day. We found V-dLMP does contribute significantly relative to the P and Q-dLMPs, and the costs associated with meeting critical grid physics constraints must be taken into account for pricing. The LEM also captures the high spatial and temporal price variations, which enables setting time-varying, differentiated local tariffs for each node. Finally, by optimally coordinating DERs, the LEM achieves much lower costs for grid operators and lower rates for customers, relative to the status quo. For future work, we will compare our LEM rigorously against other proposed markets, extend to other grid services like conservation voltage reduction, and test performance on larger networks with more severe voltage issues.

Chapter 4

A game-theoretic, market-based approach to extract flexibility from distributed energy resources

In this paper, we propose a market design based on game theory to optimally utilize the flexibility of distributed energy resources (DERs) such as solar PV, batteries, electric vehicles, and flexible loads [88]. Market agents perform multiperiod optimization to determine their feasible flexibility limits for power injections while satisfying all constraints of their DERs. This is followed by a Stackelberg game between the market operator and the agents. The market operator, as the leader,r aims to regulate the aggregate power injection around a desired value by leveraging the flexibility of their agents, and computes optimal prices for both electricity and flexibility services. The agents follow by optimally bidding their desired flexible power injections in response to these prices. We show the existence of an equilibrium among the market agents between all agents and the operator, along with simulation results for a small example system.

4.1 Introduction

In order to meet climate change mitigation goals, we need to decarbonize the power grid rapidly, and clean electricity is also essential to electrify and decarbonize other sectors like heating, transportation, and heavy industry. This will entail a transition away from fossil fuels towards clean distributed energy resources such as battery storage (BS), electric vehicles (EVs), and rooftop solar photovoltaic (PV) systems. Smart inverters will enable greater controllability of inverter-based resources (IBRs) like BS and PV. In addition, we expect greater demand response capabilities through smart meters and loads like heating, ventilation, and air conditioning (HVAC) units that can provide more load flexibility. Intelligent coordination and energy management of these DERs can provide valuable grid services and help improve network operation and efficiency, in terms of metrics like operational costs, losses, and power quality through voltage and frequency regulation. However, such resources are often owned by many different independent and autonomous agents, and thus cannot be dispatched or controlled directly by operators. This motivates the use of electricity markets and price

signals to control and coordinate DERs indirectly.

We propose a consumer-level market (CM) structure consisting of consumer market operators (CMOs) and consumer market agents (CMAs), as seen in Fig. 4.1. CMAs represent individual homes or buildings, while CMOs oversee and aggregate all the CMAs under their purview. Each CMA operates several types of DERs, which offer flexibility in terms of power injections. For instance, PV can be curtailed if needed during periods of excess solar output in the middle of the day. BS and EV can be charged (or discharged) to absorb (or inject extra power). HVAC units can also provide load flexibility by varying temperature setpoints through smart thermostats. In addition to these flexible resources, we also assume that each CMA has another portion of the load that is fixed or inflexible.

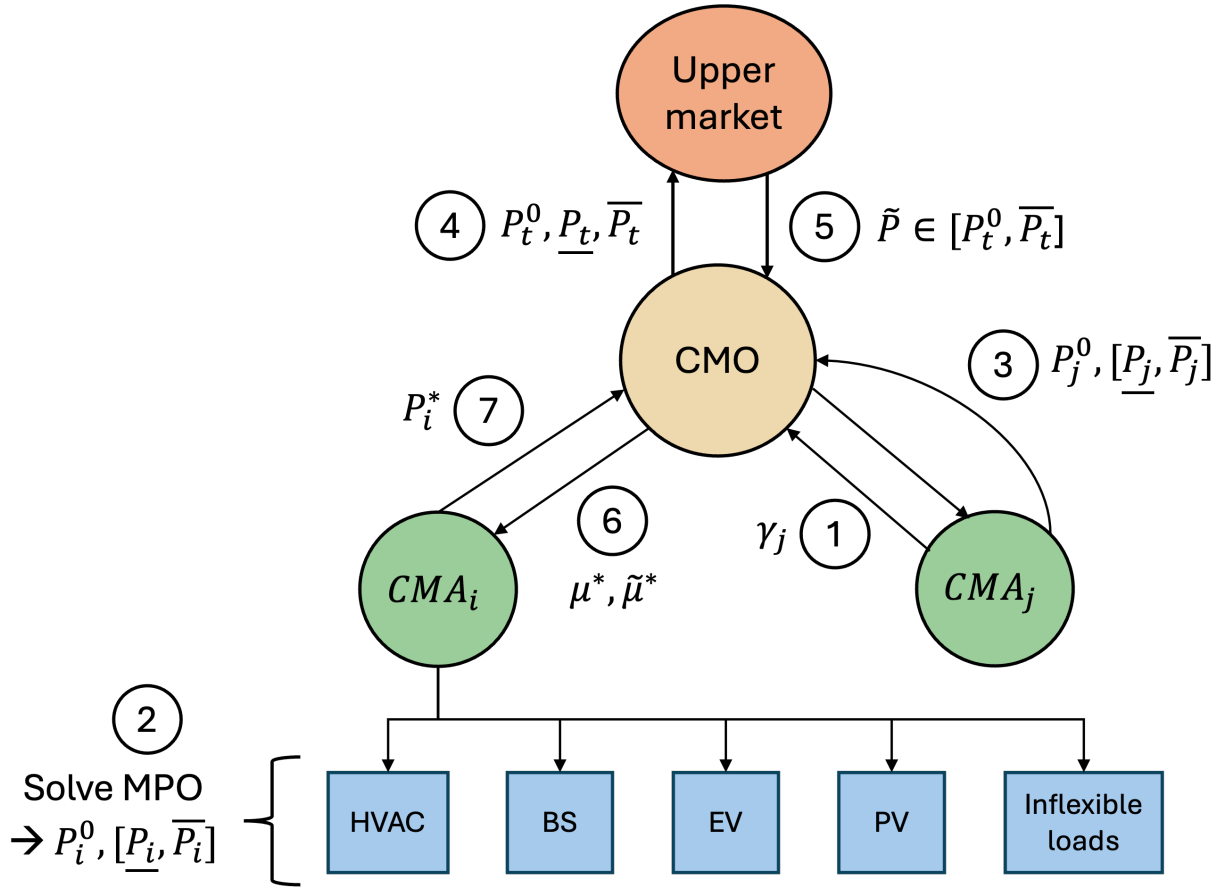

Figure 4.1: Overall schematic of the CM (see **??**). All variables will be described in the following sections.

4.1.1 Prior work and contributions

Several works have proposed new kinds of local electricity markets for the future grid rich in DERs ([47]), including retail markets ([89]) and peer-to-peer markets ([42]). We build upon this work to analyze how game theory and mechanism design can be used to inform the development of such market structures, especially those that are closest to the end-users (electricity consumers and prosumers).

There is also rich literature studying applications of game theory ([90, 91]), and mechanism design ([92]) in this domain. Common modeling approaches include Stackelberg games ([93] and coordinated (or coalitional) games ([94, 95]). Some works have also proposed distributed algorithms to solve such games ([96, 97]) or utilized Vickrey–Clarke– Groves mechanisms ([98]). Our work seeks to extend this by (i) more accurately modeling the physical dynamics and constraints of DERs, (ii) proposing a new approach to aggregate and maximize flexibility, and (iii) using different types of tariffs to charge or compensate agents.

4.2 DER modeling

Our starting point is a multiperiod optimization (MPO) problem solved by each CMA in order to determine its desired power setpoint and maximum feasible flexibility range. This MPO accounts for all the device-level constraints of individual DERs, including time-coupled state constraints for the BS, EV, and HVAC. The simulation timestep is Δt, the total simulation time is $\mathcal{T} = [0, T]$ and the planning horizon for the MPO is given by $\mathcal{H} = [t_H, t_H + (H-1)\Delta t] \subset \mathcal{T}$. We assume that the market-clearing timestep is also equal to Δt. Here, H is the length (number of timesteps) of the planning horizon. Note that throughout this paper, P denotes net active power injections (generation minus load). Thus, $P > 0$ implies net generation while loads correspond to $P < 0$. We do not consider reactive power in this work.

4.2.1 BS model

The state of charge (SOC) dynamics of the battery are:

$$SOC_i^{BS}(t+1) = (1 - \delta_{BS}^i) SOC_i^{BS}(t) - \frac{P_i^{BS}(t)\Delta t \eta_i^{BS}}{\overline{E}_i^{BS}} \tag{4.1}$$

$$\underline{P}_i^{BS} \leq P_i^{BS}(t) \leq \overline{P}_i^{BS} \tag{4.2}$$

$$\underline{SOC}_i^{BS} \leq SOC_i^{BS}(t) \leq \overline{SOC}_i^{BS} \tag{4.3}$$

$$SOC_i^{BS}(0) = SOC_i^{BS}(T) \tag{4.4}$$

where $\delta_i^{BS}, \eta_i^{BS}$ and $\overline{E}_i^{BS}$ are the BS self-discharge rate, round-trip efficiency, and maximum capacity, respectively. We also have a terminal constraint to ensure that the state of charge at the start and end of the simulation must be equal. During BS operation, the CMA also aims to minimize the cycling cost to avoid excessive charge and discharge cycles, which can degrade the battery's lifetime.

$$f_i^{BS}(P_i^{BS}) = \alpha_{cyc} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (P_i^{BS}(t+1) - P_i^{BS}(t))^2 \tag{4.5}$$

where we sum over all the timesteps in the planning horizon starting at time t_H.

4.2.2 EV model

The EV also has SOC constraints on its battery similar to the BS. In addition, we place restrictions on EV availability. We assume the EV is not present at the building or home during the period $[t_1, t_2]$, e.g., between 9am and 5pm when the owner is at work.

$$P_i^{EV}(t) = 0 \; \forall \; t \in [t_1, t_2] \tag{4.6}$$

We can impose a similar cycling cost on the EV to extend its lifetime. We also add a tracking objective that the EV owner would like to achieve a certain desired SOC (SOC_i^*) by a specific time (t^*), say the owner needs the EV to be 90% charged by 9am before work. Thus, the EV objective function is

$$\begin{aligned} f_i^{EV} = \alpha_{cyc} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (P_i^{EV}(t+1) - P_i^{EV}(t))^2 \\ + \xi_{ev}(SOC_i^{EV}(t^*) - SOC_i^{EV^*})^2 \end{aligned} \tag{4.7}$$

4.2.3 HVAC model

The HVAC dynamics describe how the power drawn affects the indoor air temperature in the home or building. In this work, we consider the HVAC unit to be a heat pump (HP), which can serve as either a heating or cooling device depending on the ambient temperature. If $T_i^{out}(t) > T_i^{in}(t)$, the temperature dynamics of the HP in cooling mode (i.e., when it acts as an air conditioner) are ([99]):

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)\left(T_i^{out}(t) + \rho_i P_i^{HP}(t)\right)$$

where $\theta_i = e^{\frac{-\Delta t}{R_i^{th} C_i^{th}}} \approx 1 - \frac{\Delta t}{R_i^{th}, C_i^{th}}, \rho_i = R_i^{th}\eta_i$ and $R_i^{th}, C_i^{th}, \eta_i$ are the equivalent thermal resistance, thermal capacitance, and coefficient of performance of the system, respectively. The temperature dynamics in heating mode, when $T_i^{out}(t) < T_i^{in}(t)$, are:

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)\left(T_i^{out}(t) - \rho_i P_i^{HP}(t)\right)$$

The HP operation is also subject to operational limits on power draw and indoor temperature:

$$\underline{P}_i^{HP} \leq P_i^{HP}(t) \leq 0 \tag{4.8}$$

$$\underline{T}_i^{in} \leq T_i^{in}(t) \leq \overline{T}_i^{in} \tag{4.9}$$

$$\tag{4.10}$$

Note that here the lower limit is determined by the maximum power consumption rating of the HVAC unit, i.e. $\underline{P}_i^{HP} = -P_{rated,i}^{HP}$ since the HP always acts as a load. The CMA would also like to track a desired temperature setpoint to maximize the thermal comfort of occupants.

$$f_i^{HP} = \xi_{ac} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (T_i^{in}(t) - T_i^{in^*})^2 \tag{4.11}$$

In addition to HVAC units, using a similar approach, we may also consider other types of thermostatically controlled loads (TCLs) such as water heaters (WH).

4.2.4 PV model

The maximum PV generation output is determined by the forecasted, time-varying solar irradiance profile $\alpha^{PV}(t)$ along with its maximum rated capacity $\overline{P}_i^{PV}$, which can be curtailed if needed.

$$0 \leq P_i^{PV}(t) \leq \alpha^{PV}(t)\overline{P}_i^{PV} \tag{4.12}$$

The objective here is to minimize the amount of clean power that is curtailed.

$$f_i^{PV} = \xi_{pv} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (\alpha^{PV}(t)\overline{P}_i^{PV} - P_i^{PV}(t))^2 \tag{4.13}$$

In addition, the CMA would also like to utilize as much of the PV output as possible (when it's available) to charge the BS and EV, by minimizing the following objective:

$$f_i^{util} = (P_i^{PV} + P_i^{BS} + P_i^{EV})^2 \tag{4.14}$$

4.3 Game theoretic formulation

We propose modeling the CM as a *Stackelberg* game, where the CMO (as the leader) acts first by announcing prices and the CMAs (as followers) then respond to these prices by bidding the desired flexibility in power injections they would each like to provide. This is a game of *incomplete information* (or a *Bayesian* game) since CMAs do not know the private types of other players and cannot observe their actions or bids. This can also be viewed as a *hierarchical* game since the CMAs are all coupled indirectly through the CMO. In addition to their flexibility bids, each player's reward (or utility) depends on the prices, which in turn are influenced by others' bids as well. The CMAs compete with each other in a *non-cooperative* fashion to provide flexibility to the CMO, and the game is repeated at every market-clearing time instance.

4.3.1 CMA DER-coordination problem

Each CMA solves an MPO to determine the optimal desired power setpoints for each DER ($P_i^{d^*}(t)$) as well as the symmetric upward or downward flexibility ($\delta_i^{d^*}(t)$) they can provide around it. In this multiobjective optimization, the CMA aims to (i) maximize the total flexibility it can provide, (ii) maximize the PV utilization for charging, (iii) minimize the costs associated with providing flexibility (as specified in Section 6.5), and (iv) maximize the total net power injection into the grid i.e. maximize net generation (import) or minimize net load (import). This is subject to all the DER-level constraints and dynamics described in the preceding sections, in addition to operational upper and lower limits with flexibility. The

MPO formulation is:

$$\min_{P_i^d(t),\delta_i^d(t)} \sum_{t\in\mathcal{H}}\sum_{d\in\mathcal{D}} -\delta_i^d(t) + f_i^d(P_i^d) + f_i^{util}(P_i^d) - P_i^{total}(t) \tag{4.15}$$
$$\text{s.t. } \underline{P_i^d(t)} \le P_i^d(t) - \delta_i^d(t),\ P_i^d(t) + \delta_i^d(t) \le \overline{P_i^d(t)}$$
$$\text{All device-specific state constraints for each DER } d \in \mathcal{D}_i$$
$$P_i^{total}(t) = \sum_{i\in\mathcal{D}_i} P_i^d(t) - P_i^{fixed}(t),\ \epsilon_1|P_i^d| \le \delta_i^d \le \epsilon_2|P_i^d|$$

where $\mathcal{D}_i \subseteq \{BS, EV, HVAC, PV\}$ is the set of all DERs owned and managed by CMA i. Note that we also place upper and lower limits on the flexibility of each DER (with $\epsilon_1 < \epsilon_2$). Since the absolute values are non-convex, we used an exact big-M reformulation to represent $|P_i^d| = P_i^{d,abs}$ by introducing additional binary variables. Note that we only need to do this $d = BS, EV$ since for the other devices we know $|P_i^{HP}| = -P_i^{HP}$ and $|P_i^{PV}| = P_i^{PV}$.

$$P_i^d = P_i^{d,+} - P_i^{d,-},\ |P_i^d| = P_i^{d,+} + P_i^{d,-},\ z_i^d \in \{0,1\}$$
$$0 \le P_i^{d,+} \le z_i^d \overline{P_i^d},\ 0 \le P_i^{d,-} \le (1 - z_i^d)|\underline{P_i^d}| \tag{4.16}$$

Thus, the overall MPO is a mixed integer quadratic program (MIQP). At each market clearing t, after solving the MPO in Eq. (4.15), the CMA aggregates the solutions and flexibilities across all their DERs to determine their total baseline power injections P_i^0 and feasible flexibility range $[\underline{P}_i, \overline{P}_i]$ that they could provide to the CMO. Note that this range is their maximum possible flexibility based on device-level costs and physical constraints of DERs. It does not account for the CMA's utility, welfare, or strategic behavior. This will be considered in Section 4.3.2.

$$P_i^0 = \sum_{d\in\mathcal{D}_i} P_i^{d^*},\ \underline{P}_i = P_i^0 - \sum_{d\in\mathcal{D}_i} \delta_i^{d^*},\ \overline{P}_i = P_i^0 + \sum_{d\in\mathcal{D}_i} \delta_i^{d^*}$$

4.3.2 CMA welfare maximization problem for game

Having determined $P_i^0, [\underline{P}_i, \overline{P}_i]$, the CMA solves the following constrained optimization to maximize its social welfare U_i^{cma} and thus determine its optimal bid into the CM. We consider the CMA as a for-profit, strategic aggregator of DERs.

$$\max_{P_i} U_i^{cma}(P_i, \widetilde{\mu}, \mu) = \widetilde{\mu}(P_i - P_i^0) + \mu P_i - \gamma_i(P_i - P_i^0)^2$$
$$\text{s.t. } \underline{P}_i \le P_i \le \overline{P}_i, P_i \ge P_i^0 \tag{4.17}$$

where $\gamma_i > 0$ is the CMA's disutility preference parameter and denotes their type. We assume in this work that the CMO always requests only *upward* flexibility in net power injections from the CMAs, i.e., increasing net injections either by reducing demand (via load shifting or curtailment) or increasing generation. This is in line with demand response programs today. This ensures that the first term in the objective function is always non-negative, since $(P_i - P_i^0) \ge 0$. This first term is the compensation given by the CMO to CMAs for any

flexibility they provide to the market, at the rate of $\widetilde{\mu}(t)$. The second term represents the net cost to the CMA, which constitutes payments to the CMO (in case the CMA is a net load $P_i < 0$) or revenue received from the CMA (in case the CMA is a net generator $P_i > 0$). This depends on the electricity price $\mu(t)$. The last term is the disutility caused to the CMA by flexibility, due to deviations from its ideal (nominal) injection P_i^0. Thus, the CMA aims to maximize its flexibility compensation and minimize disutility and net costs while satisfying the injection limits from stage I. Note that U_i^{cma} is concave for the maximization problem, thus, the corresponding minimization problem is a convex quadratic program (QP).

In this game, each CMA's action or strategy is given by its flexible power injection bid $P_i^*(t)$. The baseline electricity price $\mu(t)$ and flexibility price $\widetilde{\mu}(t)$ are determined by the CMO's strategy. Note that the prices vary over time but are common for all CMAs. By applying the Karush–Kuhn–Tucker (KKT) conditions, we can analytically derive the optimal solutions for Eq. (4.17):

$$P_i^*(\widetilde{\mu}, \mu) = \begin{cases} \overline{P}_i & \text{if } \frac{\gamma_i(\overline{P}_i - P_i^0)}{\widetilde{\mu}+\mu} < \frac{1}{2} \\ P_i^0 + \frac{\widetilde{\mu}+\mu}{2\gamma_i} & \text{otherwise} \end{cases} \tag{4.18}$$

4.3.3 CMO optimization problem

The goal of the CMO is to track the setpoint $\widetilde{P}(t)$ for the total power injection, as determined by the market above it. It does so by leveraging the flexibility bids of their CMAs. The CMO thus aims to extract as much flexibility as possible from the CMAs while accounting for their preferences and managing costs. The CMO can then offer this aggregate flexibility to the higher-level market. This upper market could be a local electricity market (LEM) for the distribution grid, as proposed in our previous work ([100]). Alternatively, the CMO could also participate directly in the wholesale energy market (WEM) by bidding its flexibility at the main transmission grid level. It has to respect the power bounds of each CMA as specified by the bids, in addition to a budget balance. This requires that the CMO exactly breaks even at each market-clearing time step. We assume here that the CMO is a not-for-profit entity that serves just as a coordinator, but we will generalize to the case where the CMO is profit-making as part of future work.

At each market clearing t, given the flexible power injection bids $P_i^*(t)$ and the other inputs $P_i^0(t), [\underline{P}_i(t), \overline{P}_i(t)]$ from all CMAs i, the CMO solves the following feasibility optimization problem to match the desired $\widetilde{P}$ and set both the prices to satisfy the budget balance constraint while respecting each CMA's flexibility bids. It's crucial to carefully set these prices so as to encourage the desired power injections from the CMAs, while accurately valuing and compensating any flexibility they provide. The budget constraint also ensures that the CMO meets its flexibility commitment $\widetilde{P}$ in the most cost-effective manner. We assume a lossless power balance between the CMO and CMAs, as well as between the CM

and the LEM above.

$$\begin{aligned} \min_{\widetilde{\mu}(t),\mu(t)} -U^{cmo}(P_i(t),\mu(t),\widetilde{\mu}(t)) &= \left(\sum_{i\in\mathcal{C}} P_i(t) - \widetilde{P}(t)\right)^2 \\ \text{s.t. } P_i^0(t) \le P_i(t) &\le \overline{P}_i(t)\ \forall\ i \\ \widetilde{\mu}(t)\sum_{i\in\mathcal{C}}(P_i(t)-P_i^0(t)) + \mu(t)\sum_i P_i(t) &= \pi(t)\sum_i P_i(t) \end{aligned} \tag{4.19}$$

Where $\mathcal{C}$ is the set of all CMAs under the purview of the CMO, and $\pi(t)$ is the time-varying electricity price set by the LEM above. This would be the locational marginal price (LMP) if the CMO participates directly in the WEM. Alternatively, this could be the d-LMP or retail rate if the CMO is part of another market in the distribution grid. The budget balance constraint implies that the net cost to the CMO for running the CM (i.e., total payment to CMAs) must equal the net revenue it receives from the LEM, or conversely, the net total revenue it receives from all the CMAs must equal the total payment it makes to the LEM. Summing up terms n the budget balance in Eq. (4.19) over all CMAs i gives:

$$\widetilde{\mu}(P_t - P_t^0) + \mu P_t = \pi P_t \tag{4.20}$$

where $P_t = \sum_i P_i$ and $P_t^0 = \sum_i P_i^0$. Further, we assume that at the start of each market operation timestep (i.e. before price scheduling and bidding), the CMAs directly chapter their solutions from the MPO (Eq. (4.15) to the CMO, i.e. P_i^0 and feasible $[\underline{P}_i(t), \overline{P}_i(t)]$. The CMO then aggregates these and chapters the total baseline injection P_t^0 as well as the upper and lower bound of total available flexibility $\underline{P}_t = \sum_i \underline{P}_i, \overline{P}_t = \sum_i \overline{P}_i$ to the upper LEM. This ensures that the LEM gives a command for the desired power setpoint that is feasible for the CM to provide, i.e., $\underline{P}_t \le \widetilde{P} \le \overline{P}_t$. We further assume that the LEM may only request upward flexibility, implying that $\widetilde{P} \ge P_t^0$. $\widetilde{P}$ and P_t^0 will also have the same sign since a net load (or generator) will remain a net load (or generator) even after flexibility provision.

4.3.4 Computing optimal prices

We now show how the CMO can set optimal prices for electricity and flexibility. Firstly, we note that CMO always sets prices such that $\mu + \widetilde{\mu} < 2\gamma_i(\overline{P}_i - P_i^0)\ \forall\ i$. According to Eq. (4.18), it follows that $P_i^* \neq \overline{P}_i\ \forall\ i$. The CMO sets this condition to ensure that no single CMA i is disproportionally affected, and we avoid relying excessively on any of the CMAs for providing flexibility. Thus, summing up the optimal flexible power injections across all CMAs in Eq. (4.18), we get the following

$$P_t^* = \sum_{i\in\mathcal{C}} P_i^* = P_t^0 + \gamma_t \frac{\mu + \widetilde{\mu}}{2} = \widetilde{P} \tag{4.21}$$

where $\gamma_t = \sum_{i\in\mathcal{C}} \frac{1}{\gamma_i}$ and we also used the fact that at the optimum of Eq. (4.19), we have $P_t^* = \widetilde{P}$ since we know the setpoint $P_t^0 \le \widetilde{P} \le \overline{P}_t$ is feasible to achieve for the CMO. The CMO now needs to compute the optimal $\mu^*, \widetilde{\mu}^*$ that causes all the CMAs to bid P_i^* such

that $\sum_i P_i^* = \widetilde{P}$ exactly. Combining Eq. (4.20) and Eq. (4.21), we can derive the baseline electricity price:

$$\mu^* = \frac{\pi \widetilde{P}}{P_t^0} - \frac{2(\widetilde{P} - P_t^0)^2}{\gamma_t P_t^0} \tag{4.22}$$

Based on Eq. (4.20), we can then also obtain the flexibility price based on the electricity tariff as follows:

$$\widetilde{\mu}^* = \frac{P_t^*(\pi - \mu^*)}{P_t^* - P_t^0} = \frac{\widetilde{P}(\pi - \mu^*)}{\widetilde{P} - P_t^0} = \frac{\widetilde{P}\left(2(\widetilde{P} - P_t^0) - \pi\gamma_t\right)}{\gamma_t P_t^0} \tag{4.23}$$

4.3.5 Positivity of prices

We also need to verify certain conditions to ensure that the prices $\mu^*(t), \widetilde{\mu}^*(t) > 0 \; \forall \; t$. From Eq. (4.23), we see that:

$$\begin{aligned} \widetilde{\mu}^* > 0 \implies & \frac{\widetilde{P}\left(2(\widetilde{P} - P_t^0) - \pi\gamma_t\right)}{\gamma_t P_t^0} > 0 \\ \implies 2(\widetilde{P} - P_t^0) - \pi\gamma_t > 0 \;\; \because \frac{\widetilde{P}}{P_t^0} > 0 \implies & \widetilde{P} > P_t^0 + \frac{\pi\gamma_t}{2} \end{aligned} \tag{4.24}$$

Given that $\widetilde{P} > P_t^0$, we note that this condition generally holds true for most realistic values of γ_t, and π since the magnitude of the term $\frac{\pi\gamma_t}{2}$ is small relative to P_t^0.

From Eq. (4.22), we can also derive conditions for $\mu^* > 0$:

$$\frac{\pi \widetilde{P}}{P_t^0} - \frac{2(\widetilde{P} - P_t^0)^2}{\gamma_t P_t^0} > 0 \implies \frac{\pi \widetilde{P}\gamma_t - 2(\widetilde{P} - P_t^0)^2}{\gamma_t P_t^0} > 0$$

We notice that $\mu^* > 0$ for any negative values of $\widetilde{P}, P_t^0$, i.e., if the CMO as a whole is a net load. However, if the CMO is a net generator, i.e., $\widetilde{P}, P_t^0 > 0$, we require the following conditions to hold true:

$$\begin{aligned} \mu^* > 0 > \implies (\widetilde{P} - P_t^0)^2 < \frac{\pi \widetilde{P}\gamma_t}{2} \implies & a_1 < \widetilde{P} < a_2 \\ a_1 = P_t^0 + \frac{\pi\gamma_t}{4} - \frac{1}{2}\sqrt{\frac{\pi\gamma_t}{2}\left(4P_t^0 + \frac{\pi\gamma_t}{2}\right)} \\ a_2 = P_t^0 + \frac{\pi\gamma_t}{4} + \frac{1}{2}\sqrt{\frac{\pi\gamma_t}{2}\left(4P_t^0 + \frac{\pi\gamma_t}{2}\right)} \end{aligned} \tag{4.25}$$

Combining Eq. (4.24) and Eq. (4.25), we get the following on $\widetilde{P}$ for positivity of $\mu^*, \widetilde{\mu}^*$:

$$\begin{cases} \widetilde{P} > P_t^0 + \frac{\pi\gamma_t}{2} & \text{if } \widetilde{P}, P_t^0 < 0 \\ \max\left(a_1, P_t^0 + \frac{\pi\gamma_t}{2}\right) \leq \widetilde{P} \leq a_2 & \text{if } \widetilde{P}, P_t^0 > 0 \end{cases} \tag{4.26}$$

During our numerical simulations, the CMO as a whole was generally a net load (i.e. $\widetilde{P}, P_t^0 < 0$) and we found that the conditions above in Eq. (4.26) consistently held true for all time periods, ensuring that the prices remained positive throughout the day.

4.3.6 Overall process for market operation

We summarize the steps for our CM operation as follows (also see Fig. 4.1):

1. All CMAs chapter their type γ_i to the CMO only at the start of the game, and these remain the same for all subsequent market clearings. In this work, we assume that the CMAs chapter their type truthfully.

2. CMAs solve MPO in Eq. (4.15) to determine ideal baseline injection P_i^0 and maximum feasible flexibility intervals $[\underline{P}_i, \overline{P}_i]$.

3. CMAs communicate their baseline power injection P_i^0 and flexibility range $[\underline{P}_i, \overline{P}_i]$ to the CMO. We argue that the CMAs are incentivized to be truthful about P_i^0 to maximize their utility (see Eq. (4.17)) since this is their ideal power injection if no flexibility is needed.

4. The CMO chapters aggregated $P_t^0, \underline{P}_t, \overline{P}_t$ to the LEM.

5. The LEM requests a feasible flexible setpoint $\widetilde{P} \in [P_t^0, \overline{P}_t]$ from the CMO.

6. The CMO announces optimal prices $\mu^*, \widetilde{\mu}^*$ to all its CMAs, to track the regulation signal $\widetilde{P}$.

7. CMAs respond to the prices with their optimal bids for desired flexible power injections P_i^*, following which the market is cleared and settled.

We note here that we still need to verify the incentive compatibility (IC) of our market design to ensure that the CMAs chapter their true types and bid according to their actual preferences and flexibility capabilities. We will study this more rigorously as part of future work. Moreover, we know from the revelation principle that it is possible, if needed, to come up with a truthful, IC mechanism to implement our optimal bidding and pricing functions (derived above) and achieve the same Bayesian equilibrium outcome and welfare.

4.3.7 Equilibrium

We now show that the optimal prices $\mu^*, \widetilde{\mu}^*$ set by the CMO will induce optimal bids P_i^* from all CMAs that lead to an equilibrium in pure strategies.

Theorem 1. *The above set of bids and prices $\{P_i^* \; \forall \; i, \mu^*, \widetilde{\mu}^*\}$ corresponds to a unique Nash equilibrium (NE) amongst the CMAs and thus also a unique Stackelberg equilibrium (SE) between all the CMAs and the CMO.*

Proof. Note that the welfare of each CMA i also depends others' actions and is coupled through the prices:

$$U_i^{cma}(P_i, \widetilde{\mu}(P_i, P_{-i}), \mu(P_i, P_{-i})) \equiv U_i^{cma}(P_i, P_{-i}, \widetilde{\mu}, \mu)$$

where P_{-i} denotes the flexible power injection bids of all the CMAs other than CMA i. Given the electricity and flexibility prices $\mu^*, \widetilde{\mu}^*$ set by the CMO, the CMAs bid their welfare-maximizing flexible power injections P_i^* according to Eq. (4.18). Thus, this strategy profile $\{P_i^*\}$ corresponds to a unique NE among the CMAs since:

$$U_i^{cma}(P_i^*, P_{-i}^*, \widetilde{\mu}^*, \mu^*) \geq U_i^{cma}(P_i, P_{-i}^*, \widetilde{\mu}^*, \mu^*) \forall P_i \in [\underline{P}_i, \overline{P}_i]$$

Similarly, given that all the CMAs submit their best response flexible injection bids P_i^*, the optimal prices $\mu^*, \widetilde{\mu}^*$ from Eqs. (4.22) and (4.23) result in a unique SE between the CMO and the CMAs. This is because $\mu^*, \widetilde{\mu}^*$ is the unique pair of prices that balances the budget constraint in Eq. (4.19) and exactly matches the total power injection $\sum_i P_i$ to equal $\widetilde{P}$, thus implying that:

$$U^{cmo}(P_i^*, P_{-i}^*, \mu^*, \widetilde{\mu}^*) = 0 \geq U^{cmo}(P_i^*, P_{-i}^*, \mu, \widetilde{\mu}) \ \forall \mu, \widetilde{\mu}$$

□

Note: Nash's theorem states the existence of a mixed strategy Nash equilibrium (NE) for any finite game (with discrete strategy sets), and Glicksberg's theorem extends this to the infinite setting with continuous strategies, as in our case ([101]). However, in reality, computing these equilibria can be very challenging. We thus needed to make several assumptions and simplifications to be able to readily show the existence and uniqueness of the NE, and derive it analytically.

4.4 Simulation results

We now briefly describe some simulation results. We simulated a hypothetical CM with 1 CMO and 3 CMAs, each of whom has all the types of DERs (BS, EV, HVAC, PV) along with some fixed load. We used outdoor air temperature and solar PV production data for San Francisco, California, and price data for $\pi(t)$ from CAISO. Fig. 4.2a shows the total power injections for the CMO, aggregated over all 3 CMAs. We see that the CMO remains a net load throughout the day, with the net load lower mid-day during peak solar PV output. The LEM requests varying amounts of load flexibility or demand response throughout the day. Fig. 4.2b shows the original LEM price π along with the CM electricity price μ^* and flexibility price μ^*. We notice that the CM prices are about an order of magnitude higher than the LEM price. Thus, for the CMO to provide the required flexibility to the LEM, it has to increase its prices and also compensate CMAs and DERs sufficiently, which raises costs. We are currently exploring approaches to mitigate these price impacts. One possibility is to use varying prices $\mu_i, \widetilde{\mu}_i$ for each CMA i rather than a common rate. Although this makes the equilibrium analysis more challenging, it could lead to more equitable and fair tariffs ([102]).

Fig. 4.3a shows the power injections of various DERs owned by CMA 1 over the course of the day. As expected intuitively, we see that the BS charges using the solar PV output during the middle of the day, while the EV is only available to charge over a small fraction of this period. On the other hand, both the BS and EV generally discharge through other parts

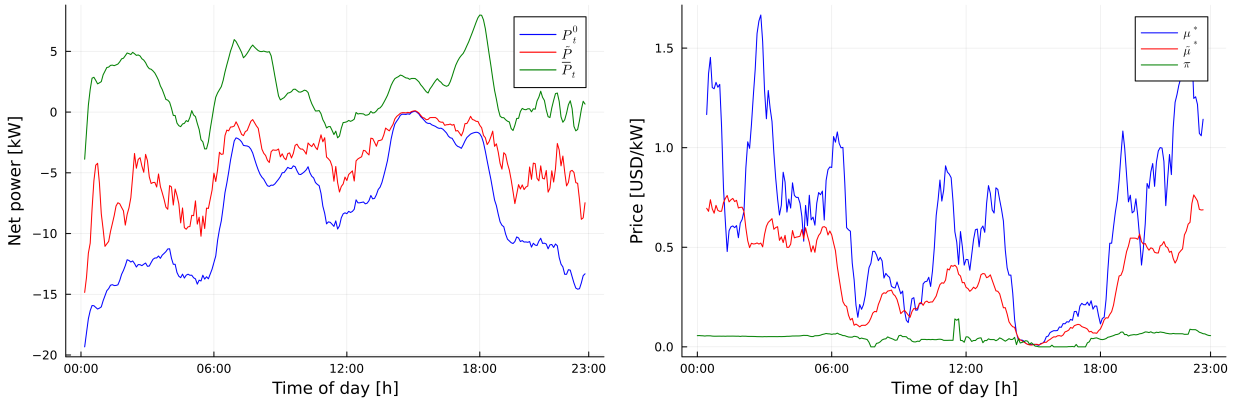

(a) Total CMO power injections.

(b) Prices over the day.

Figure 4.2: CMO power injections and prices.

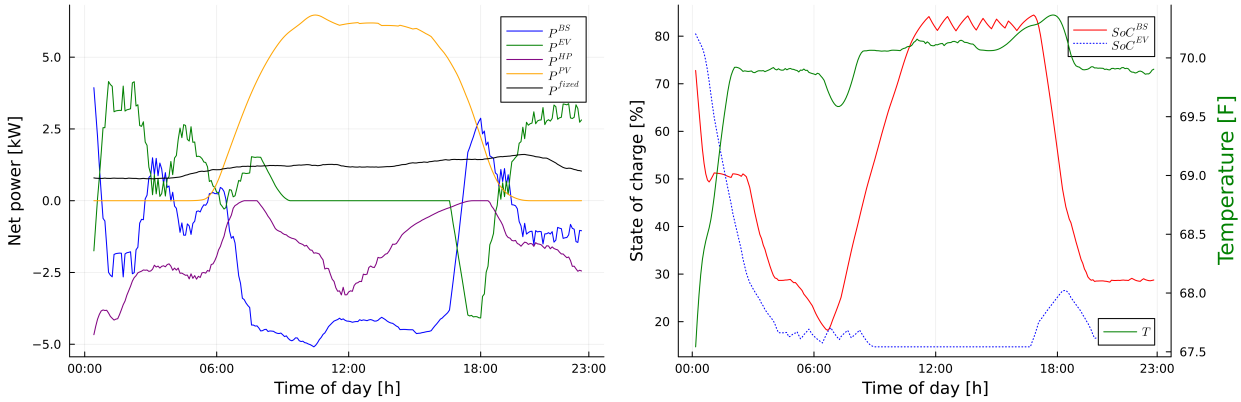

(a) DER power injections.

(b) Temperature and state of charge profiles.

Figure 4.3: DER power injections and states for CMA 1.

of the day (evening through early morning) to provide significant flexibility. These are also reflected in the SOC profile of the BS and EV over the day, shown in Fig. 4.3b. We see that the power consumption of the HVAC unit or heat pump (HP) also varies quite significantly throughout the day, confirming that temperature setpoint control can indeed provide load flexibility to the CMA, while still maintaining the temperature profile close to the desired comfortable value of 70 F, as seen in Fig. 4.3b. The HP generally serves as a cooling device throughout the day, except from late night to early morning when it acts as a heater instead due to lower outdoor temperatures.

4.5 Conclusions and future work

In this paper, we proposed a consumer-level market structure for a CMO to coordinate and aggregate several CMAs, each of whom operates several DERs. In our hierarchical structure, the CMAs first coordinate their DERs to determine their flexibility capabilities while

accounting for all device-specific constraints. This is followed by a Stackelberg, incomplete information game between the CMO, who sets prices, and CMAs who respond to the prices with flexibility bids. We derive analytical solutions for prices and bids that lead to an equilibrium among all market participants, along with simulation results for a small instantiation of the CM.

As part of future work, we will conduct simulations on larger systems. We also plan to extend our work to more realistic settings by relaxing some of our assumptions about common knowledge and information availability. To do so, we will leverage more advanced game-theoretic approaches and possibly other notions of equilibria. In more realistic and complex settings, we will also explore numerical approaches to compute approximate equilibrium solutions.

Chapter 5

Market-based DER coordination to strengthen grid resilience

The electricity grid has evolved from a physical system to a cyber-physical system with digital devices that perform measurement, control, communication, computation, and actuation. The increased penetration of distributed energy resources (DERs) that include renewable generation, flexible loads, and storage provides extraordinary opportunities for improvements in efficiency and sustainability. However, they can introduce new vulnerabilities in the form of cyberattacks, which can cause significant challenges in ensuring grid resilience. Fig. 5.1 shows an example of possible attack surfaces for a distribution grid. This is a modified version of the standard IEEE 123-node test feeder, which was modified to have high penetrations of DERs, in collaboration with PNNL.

We propose a framework in this paper for achieving grid resilience through suitably coordinated assets, including a network of Internet of Things (IoT) devices. A local electricity market is proposed to identify trustable assets and carry out this coordination. Situational Awareness (SA) of locally available DERs with the ability to inject power or reduce consumption is enabled by the market, together with a monitoring procedure for their trustability and commitment. With this SA, we show that a variety of cyberattacks can be mitigated using local trustable resources without stressing the bulk grid. The demonstrations are carried out using a variety of platforms with a high-fidelity co-simulation platform, real-time hardware-in-the-loop validation, and a utility-friendly simulator. Fig. 5.2 provides a holistic overview of our market-based distributed coordination approach for attack mitigation and grid resilience, known as EUREICA.

5.1 Introduction

The electricity grid is going through a rapid transformation in an effort toward deep decarbonization. Large synchronous generators powered by fossil fuels such as oil, natural gas, and coal are being phased out in favor of solar and wind-based generation. While the latter enables the necessary move towards a reduced carbon footprint, it brings two major challenges in ensuring the reliable and resilient delivery of electricity to the end-user. The first of these is the temporal signature of these renewables – the amount of generation varies

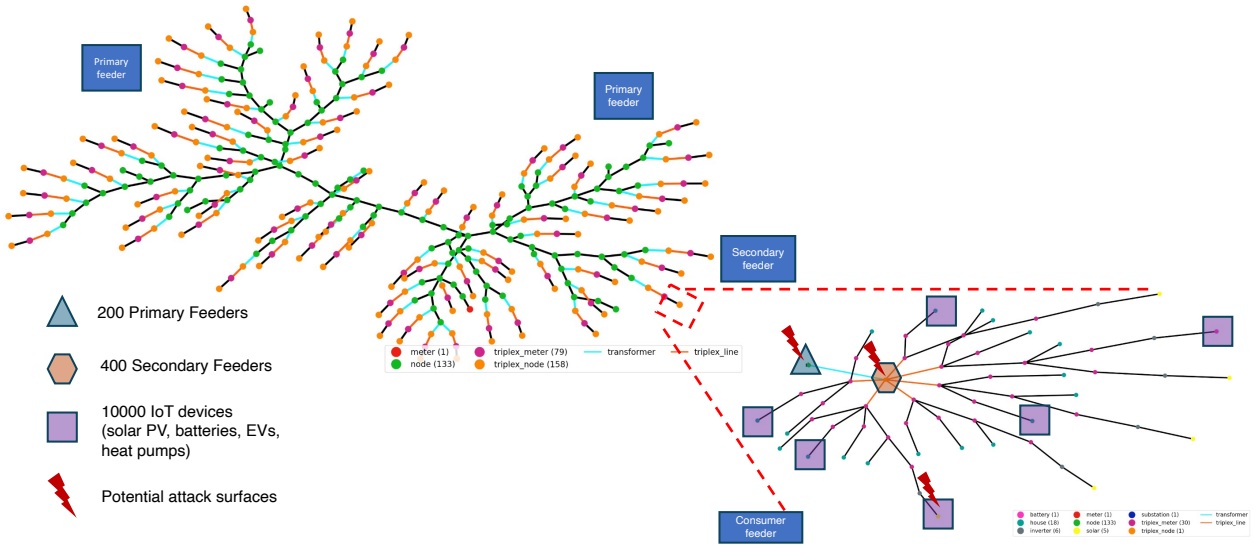

Figure 5.1: Example of a DER-rich distribution feeder with several potential attack surfaces.

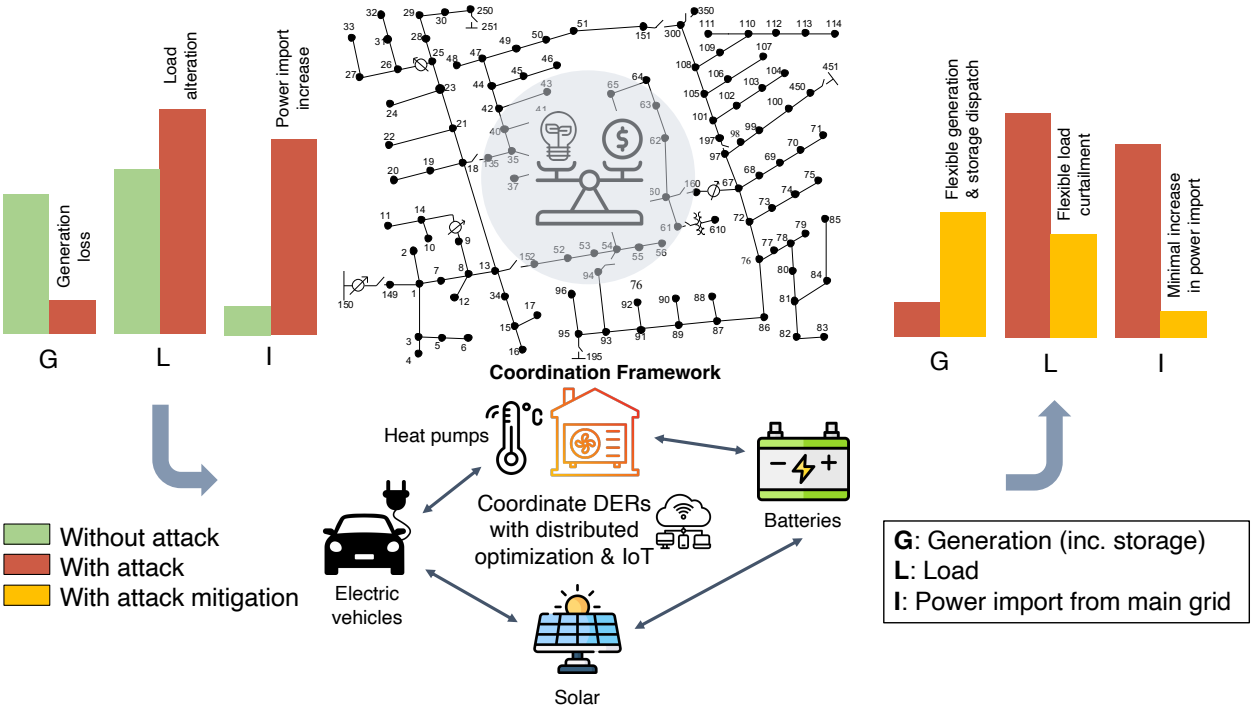

Figure 5.2: Summary of EUREICA attack mitigation framework.

with time, both in terms of intermittency and uncertainty. The second is that these are distributed and large in number. A strong enabler of the scale of the DERs is IoT, which denotes a network of physical devices such as water heaters (WHs), air-conditioners, and electric vehicles (EVs), as they enable automated and fast operation of various loads. And their pervasiveness brings in complexities of heterogeneity, decentralization, and scale. In order to ensure the reliability of the grid despite these challenges, a precise coordination of these DERs, both in space and time, has to be carried out. In particular, the power balance of generation and consumption has to be ensured at all locations and at each instant. These challenges are being overcome using a pervasive cyber layer that senses, communicates, coordinates, and enables the requisite power injection and consumption throughout the grid.

In addition to reliability, an essential property of the electricity grid is its resilience [103]. This central property, which denotes the ability of the grid to withstand and recover quickly to supply critical loads following a major disruption, such as an outage, a natural calamity, a cyberattack, or a cascading failure, is paramount, even with increased penetration of DERs. In this context of ensuring resilience, the very transformations that enable deep decarbonization, including the development of cyber-grid infrastructure, the adoption of IoT devices, the use of dynamic renewable energy sources, and increased electrification of transportation, could also introduce new vulnerabilities. Cyberattacks can disclose, deceive, or disrupt crucial information, thereby causing significant damage, ranging from small outages to brownouts and blackouts. Recent reports [104–107] indicate the ubiquity, ease, and scale of cyberattacks on sensitive industrial environments, including supervisory control and data Acquisition (SCADA), operational technology (OT), and industrial control systems (ICS), underscoring the importance of ensuring resilience to such adversaries.

By and large, most of the information for power grid operations flows through utility-controlled communication networks, which are more reliable and resilient than commercial networks and utilize commercial telecommunications services for other informational needs, such as accessing the internet and communicating with customers. Such a tight separation is challenged by the increased information flow, which becomes necessary with a stronger presence of a cyber-layer, which in turn is necessitated due to increased coordination and automation at the grid edge. What has remained as tightly closed systems thus far may have to relax their boundaries, introducing complexities in the underlying communication. While air gaps and protections will always be important and included, imperfect protections are inevitable as complexity increases. With the increased penetration of instrumentation and automation, motors and generators may be manipulated by adversaries to open and close at will. Another point to be noted is that with increased complexities due to intermittent and uncertain generation and consumption, utilities alone cannot cater to all needs; public and private partnerships may be necessary. It is therefore extremely important to design an appropriate cyberinfrastructure that ensures that the lights stay on, despite increased communication, which may be between disparate stakeholders. The focus of this paper is on such a distributed decision-making framework.

Given the size and complexity of the problem of cyberattacks, providing a complete resilience framework for the entire power grid is a tremendously difficult task. In this paper, we propose a first step, of providing SA to the grid operators in a distribution grid, with SA corresponding to the knowledge of local DERs in terms of their location and the amount of power generation that they are able to provide, as well as a resilience score (RS) that the

operators can make use of to provide resilience. We argue that this first step, of providing SA, is enabled through a local electricity market (LEM) structure that consists of operators at different voltage levels in a distribution grid. This market structure is proposed to be local, across the distribution grid, electrically co-located with primary and secondary circuits, with operators scheduling all DERs at the corresponding nodes in a given region. The market will also include IoT-coordinated assets (ICAs), with the assumption that the ICA will have computing capability and the ability to exchange information. The overall framework, EUREICA (Efficient, Ultra-REsilient, IoT-Coordinated Assets), is the innovation in the proposed cyberinfrastructure and will be shown to lead to SA made available to operators placed hierarchically at various locations, thereby providing an important first step in ensuring resilience.

5.2 Brief literature review

LEMs have been addressed in several studies including [16, 35, 41, 42, 47, 73], with real-field implementations beginning to be reported [43, 46], all of which show the feasibility of a local market structure, and its advantages compared to alternate solutions that are designed to encourage full participation of DERs [89, 107]. The LEM structure that we propose in this paper builds on that in [73]. The resilience of the electricity grid to cyberattacks has been explored in a very large number of studies (see [105, 108–110] and references therein), with new results appearing continuously. Broadly, these approaches can be categorized into detection and isolation of the attack [111], prevention of the attack, and resilience in the presence of attacks. For large-scale attacks such as those described in [106, 112, 113], these methods are inadequate; it may be near-impossible to identify the attacker, but rather that an attack has occurred. Prevention of the attack can be enabled through varying levels of access and authorization [114] and monitoring, isolation, and protection at the component level [106]. However, as the scale, location, and number of IoT devices in particular, and DERs in general, grow, it becomes exceedingly difficult to completely prevent attacks. Ensuring resilience, especially in the face of large-scale attacks, for a large-scale system such as the electricity grid, is exceedingly difficult; current literature has either focused on systems at a small scale or with low levels of renewables. The EUREICA framework that we propose will provide SA that detects that an attack has occurred, and with this SA, deploys trustable ICAs in order to mitigate the impact of the attack, and ensure grid resilience through a distributed decision-making strategy.

The distributed decision-making in EUREICA is enabled through an LEM, a schematic of which is shown in Fig. 5.3. The same market structure [73], which has been shown to lead to grid reliability [115] and provide grid services such as voltage support [116] in addition to overall power balance, is demonstrated in this paper to ensure grid resilience against cyberattacks using local trustable DERs. In particular, the results will show that local resilience is attainable through SA of locally available ICAs that have the ability to inject power or reduce consumption, as well as a procedure for monitoring their trustability and commitment. The demonstrations are carried out using a variety of platforms such as (i) GridLAB-D, which enables the simulation of distribution grids with high fidelity, (ii) the advanced research on integrated energy systems (ARIES) platform that includes a real-time

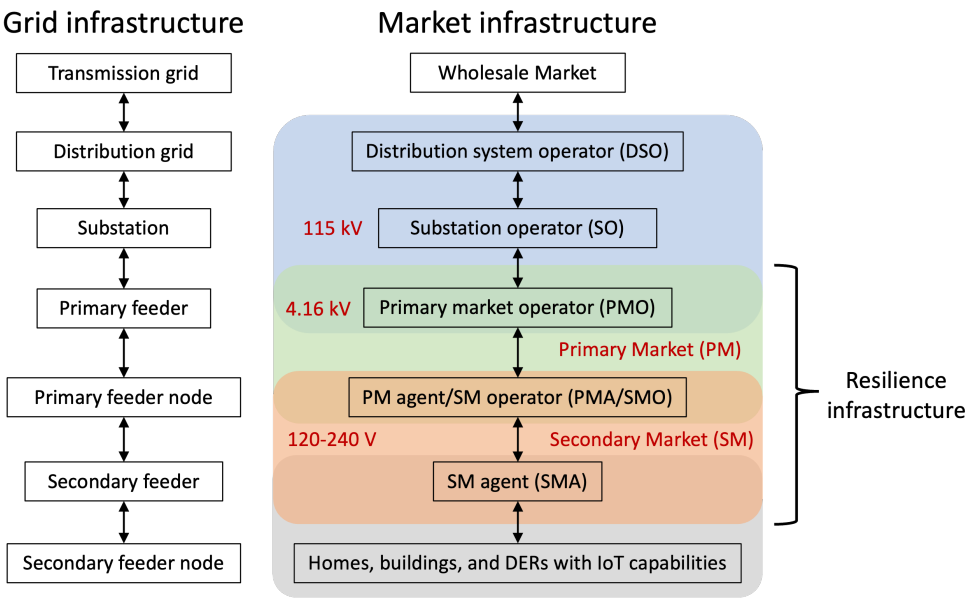

Figure 5.3: A Hierarchical LEM for a Distribution Grid. The resilience infrastructure utilizes the dual market layer consisting of PM-SM.

digital simulator (RTDS) and enables hardware-in-the-loop (HIL) validation, and (iii) General Electric's advanced distribution management system (ADMS) [117], distribution operations training simulator (DOTS), and DER integration middleware (DERIM).

5.3 Problem Statement

In this section, we delineate the problem statement, which pertains to vulnerabilities that can occur in a distribution grid that is seeing an increasing penetration of DERs. As a result, vulnerabilities in the form of cyberattacks can occur, where a variety of devices can be denied service, disrupted, or forced to disclose their identity due to adversaries tampering with communication. As the starting point, we briefly describe the approach that we take to ensure grid resilience and outline the scenarios that we will explore to demonstrate how grid resilience can be achieved using our approach.

A typical path of electricity delivery to end users traverses generation, transmission, and distribution. Distribution substations connect to the transmission system and gradually step down the voltage from 44kV or higher to 33kV (denoted as a primary network), then down to 11.2/4.6kV (denoted as a secondary network), and further down to 110V or 220V, depending on the specific region in the world. While the 20th century witnessed distribution systems operating as simple distribution lines as vehicles for sharing the electricity from transmission networks, today's distribution systems are increasingly becoming heavily integrated with distributed energy resources, that correspond to resources that are located closer to the load, including renewable generation, some of which may be behind the meter [107], batteries, and flexible consumption units. This, in turn, is causing distribution systems to become more independent and to be required to take on increased responsibilities for services such as

grid reliability and grid resilience. Other examples of DERs are distributed photovoltaics (DPVs) like rooftop solar systems, combined heat and power plants, electric vehicles, and diesel generators. DERs vary in size, from DPV systems that range between 1-1000kW in size to larger ground-mounted solar farms that range to several MW. With technological advances in power electronics and associated smart inverters as well as protection systems, fewer restrictions are being placed on the size and locations of the DERs, providing an opportunity for them to play stronger and more central roles in grid reliability and resilience [118].

Over the past years, DERs have been shown to be increasingly useful in providing key grid services such as volt-var control [119]. The central idea in these explorations is that key information is exchanged, in a distributed manner, between suitable individual components in the primary and secondary networks, coordinated both in space and time, thereby allowing local control over power injection and reduction of load at key locations and instants. Such a correct operation of the complete distribution network is predicated on this key information reaching the recipients in a secure manner. This sets the stage for malicious attacks that can disconnect and disrupt the overall grid by impairing key components.

Several attacks on power systems have been recently reported [106, 112, 120–125] on the central control systems, key nodes in the distribution grid, or at the devices at the end-user level. Those at the device end, denoted as MadIoT (Manipulation of Demand via IoT) attacks, correspond to a botnet at a secondary network node that causes the corresponding load to change abruptly. If this node corresponds to a high-wattage device, and the attack is coordinated through malware that simultaneously corrupts a large number of these devices, an argument can be made that it can cause frequency instabilities, line failures, and subsequently a severe disruption on the overall power grid. Building on the results in [124, 125], the results in [112] show that even with realistic load profiles, a strategically coordinated attack can show a better success rate than in [124, 125], requiring fewer compromised IoT devices without triggering well-established protection systems. The well-known attack studied in [106], on the other hand, is at the central control system level, which was a well-planned strategic attack that led to a power outage affecting 250,000 customers over a significant period of time. The question we address in this paper is: *How can we use a cyberinfrastructure with IoT-Coordinated Assets (ICA) to support grid resilience against cyber-attacks?*

The specific approach that we propose to circumvent the anomalous scenario consists of two steps: (1) Enable improved visibility over the grid and net power injections available at various nodes through a hierarchical market structure with operators at the primary network and secondary network nodes; (2) Enable the market operators to determine an RS computed through monitoring of various features of the communication network. Steps (1) and (2) together provide SA to the grid operators (as shown in Eq. (5.1)). Our central thesis is that through this SA, operators can determine that an attack has occurred and take appropriate steps to mitigate the impact of the attack in a timely manner. The system operators and resilience managers are suitably co-located with the electrical assets so as to respond quickly through a distributed decision-making framework. The framework, therefore, avoids the computational pitfalls of a centralized architecture while still being underpinned by a substrate of communication, sensing, and actuation. The overall solution is also well-placed to integrate with the existing grid operational and market structures, helping accelerate its adoption in the field.

Attack Number	Attack type	Attack surface	Grid connection	Power flow model	Grid model	Scale of attack [kW]
1a	LA	PMA	Grid-connected	Current injection	Unbalanced, 3-phase	36
1b	DG	PMA	Grid-connected	Current injection	Unbalanced, 3-phase	45
1c	DG	SMA	Grid-connected	Current injection	Unbalanced, 3-phase	157
2a	DG	PMA	Grid-connected	Branch flow	Balanced, single-phase	261
2b	DG	PMA	Grid-connected	Branch flow	Balanced, single-phase	650
3	DG	PMA	Islanded	Current injection	Unbalanced, 3-phase	2500

Table 5.1: Summary of attack scenarios and use-cases, LA = load alteration attack, DG = distributed generator attack.

5.4 Situational awareness

Formally, we define SA at an operator x as the tuple

$$\mathrm{SA}_x = \{\mathrm{ICA}_x, \mathrm{RS}_x\}, \tag{5.1}$$

where ICA_x stands for the IoT-coordinated assets and denotes the generator and/or consumption flexibilities of DERs under the purview of agent $x \in \{\mathrm{SMA}, \mathrm{SMO}\}$, and RS_x denotes their resilience scores, to be defined in Section 5.5.5. We will show, in what follows, that RS_x can be determined based on the asset's market performance and security against possible attacks.

We now show how the LEM made of the PM-SM layers will allow the computation of SA. The operation of a distribution grid is challenging due to its scale, complex topology, and presence of various active DER assets and fixed load nodes. We separate this complex task by having the PM focus on grid-specific costs and constraints while the SM focuses on consumer-centric costs and constraints. We assume that PM and SM clear once every 5 minutes and 1 minute, respectively. The main reason for this separation of timescales is that the SM typically needs to monitor fewer assets than the PM, and is closer to DER devices (such as rooftop solar and batteries) and therefore may need to operate at a faster timescale than a PM. The starting point for both markets is the submission of bids by the corresponding agents. Bids for the SM are submitted by the SMAs exogenously, whereas bids for the PM are computed by the SMOs via the SM.

Accordingly, we describe the operation of the SM before going into the details of the PM.

5.5 Secondary market

The operation of the SM consists of three sequential stages: bidding, clearing, and monitoring. We denote $\mathcal{N}$ to be the set of all SMOs in the network and $\mathcal{N}_i$ to be the set of all SMAs under a given SMO $i \in \mathcal{N}$.

Fig. 5.4 shows the inputs and outputs for different levels of the hierarchical LEM. For both the SM and the PM, the inputs consist of the baseline power injections and flexibility bids, while the outputs are the market schedules (setpoints for power injections) and their associated flexibility ranges, along with the corresponding electricity prices of tariffs.

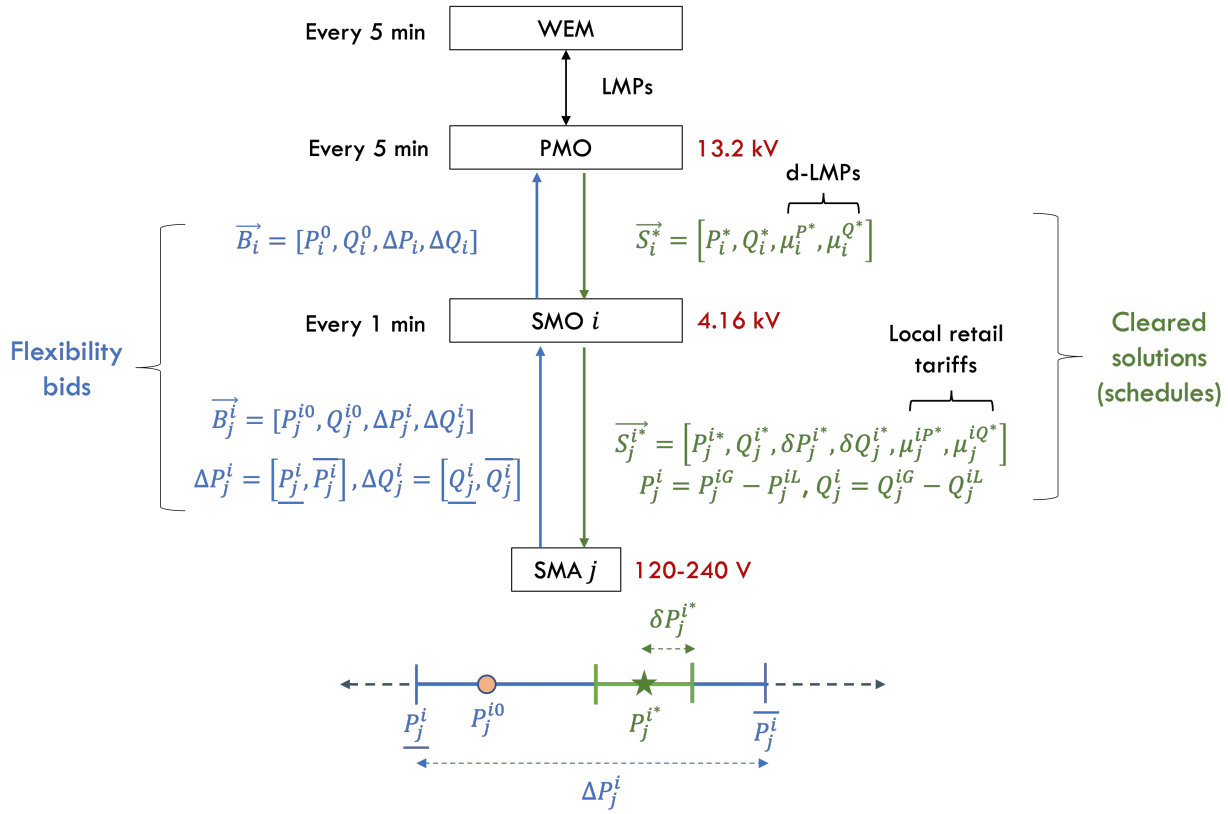

Figure 5.4: Overall inputs and outputs in the LEM.

5.5.1 SM bidding

During the bidding phase, each SMA $j \in \mathcal{N}_i$ submits a bid $\mathcal{B}_j^{iS}$ defined as

$$\mathcal{B}_j^{iS} = \{P_j^{i0}, Q_j^{i0}, \underline{P}_j^i, \underline{Q}_j^i, \overline{P}_j^i, \overline{Q}_j^i, \beta_j^{iP}, \beta_j^{iQ}\}.$$

P_j^{i0} and Q_j^{i0} denote the baseline active and reactive injections of SMA j, along with the upward $(\overline{P}_j^i, \overline{Q}_j^i)$ and downward flexibility $(\underline{P}_j^i, \underline{Q}_j^i)$. β_j^{iP} and β_j^{iQ} denote the disutility parameters associated with providing active and reactive power flexibility, respectively. It should be noted that Bid B_j^{iS} requires SMA j to have a realistic estimate of its energy profile for the next 1 minute. Since it is not always trivial to predict future power availability, agents deploy a decentralized federated learning (FL)-based framework [126] to determine their bids. Using FL helps ensure that the privacy of the participating agents is preserved and the computational aspects of the prediction algorithm scale well as the number of agents increases. Further details on the FL implementation can be found in **??**. Fig. 5.4 summarizes details of the overall LEM.

5.5.2 SM clearing

Once the SMO i has received bids from the participating SMAs, it clears the market with active and reactive power injection setpoints (P_j^{i*}, Q_j^{i*}) and the corresponding retail tariffs $(\mu_j^{iP*}, \mu_j^{iQ*})$. In addition, the SMO also solves for the optimal flexibility ranges $(\delta P_j^{i*}, \delta Q_j^{i*})$ for $j \in \mathcal{N}_i$. The SMO clears the markets with the following objectives: (O1) maximization of aggregate resilience f_i^1, (O2) minimization of the net cost to the SMO, f_i^2, (O3) maximization of total flexibility f_i^3 that the SMO can extract from all its SMAs and (O4) minimization of the disutility of the SMAs f_i^4, arising from flexibility provision. This gives rise to a multiobjective constrained optimization problem:

$$\min_{\mathbf{y}_i^S} f_i^S = \{f_i^1, f_i^2, f_i^3, f_i^4\}^\top \tag{5.2a}$$

$$\text{s.t.} \underline{P}_j^i + \delta P_j^i \leq P_j^i \leq \overline{P}_j^i - \delta P_j^i \; \forall j \in \mathcal{N}_i, \; \forall \text{ constraints} \tag{5.2b}$$

$$\underline{Q}_j^i + \delta Q_j^i \leq Q_j^i \leq \overline{Q}_j^i - \delta Q_j^i \tag{5.2c}$$

$$\delta P_j^i, \; \delta Q_j^i \geq 0, 0 \leq \mu_j^{iP} \leq \overline{\mu}^{iP}, 0 \leq \mu_j^{iP} \leq \overline{\mu}^{iQ} \tag{5.2d}$$

$$\sum_{t_p} \sum_{t_s} \sum_{j \in \mathcal{N}_i} \mu_j^{iP}(t) P_j^i(t) \Delta t_s \leq \sum_{t_p} \mu^{iP^*}(\hat{t}_p) P_i^*(\hat{t}_p) \Delta t_p \tag{5.2e}$$

$$\sum_{t_p} \sum_{t_s} \sum_{j \in \mathcal{N}_i} \mu_j^{iQ}(t) Q_j^i(t) \Delta t_s \leq \sum_{t_p} \mu^{iQ^*}(\hat{t}_p) Q_i^*(\hat{t}_p) \Delta t_p \tag{5.2f}$$

$$\sum_{j \in \mathcal{N}_i} P_j^i(t_s) = P^{i*}(\hat{t}_p), \quad \sum_{j \in \mathcal{N}_i} Q_j^i(t_s) = Q^{i*}(\hat{t}_p) \tag{5.2g}$$

The constraints include capacity limits and operational bounds on SMA injections (including flexibilities), budget balance constraints, price ceilings, and lossless power balance. Note that here we do not account for all the power physics; these will be considered in the PM. The

decision variables consist of the P and Q injection setpoints as well as retail tariffs for each SMA i.e. $\mathbf{y}_i^S = \{\mathbf{y}_j^{iS}\}\ \forall j \in \mathcal{N}_i$ where $\mathbf{y}_j^{iS} = [P_j^i, Q_j^i, \delta P_j^i, \delta Q_j^i, \mu_j^{iP}, \mu_j^{iQ}]$. We note from the choice of f^1 that the solution of Eq. (5.2) requires the resilience scores RS_j^i. This is assumed to be communicated by the secondary resilience manager (SRM) to the SMA. The details of the SRM are addressed in the next section.

In general, the optimization problem in Eq. (5.2) has multiple solutions known as Pareto points, with each solution prioritizing different objectives. However, since the objective functions have different units, instead of finding the Pareto solutions, we use a hierarchical approach proposed in [116] where the SMO optimizes one objective at a time in descending order of importance. While optimizing the subsequent objective functions, additional constraints on the degradation of prior objectives are added to the optimization problem (see [73] for details). The cleared market schedules $\mathbf{y}_i^{S*}$ are sent by the SMO to their corresponding SMAs, as well as to their SRM.

5.5.3 Objective functions for optimization in the Secondary Market

The four objective functions considered in the SM clearing are defined as:

O1. Maximization of aggregate resilience, f_i^1, given by the following, where RS_j^i denotes the resilience score of SMA j under SMO i

$$f_i^1 = -\sum_{j=1}^{n} RS_j^i((P_j^i - P_j^{i0})^2 + (Q_j^i - Q_j^{i0})^2)$$

O2. Minimization of net cost, f_i^2 to the SMO for running the SM

$$f_i^2 = \sum_{j=1}^{n} \mu_j^{iP} P_j^i + \mu_j^{iQ} Q_j^i$$

O3. Maximization of total flexibility, f_i^3 that the SMO can extract from all its SMAs

$$f_i^3 = -\sum_{j=1}^{n} (\delta P_j^i + \delta Q_j^i)$$

O4. Minimization of disutility of the SMAs, f_i^4 arising from flexibility provision

$$f_i^4 = \sum_{j=1}^{n} \beta_j^{iP}(P_j^i - P_j^{i0})^2 + \beta_j^{iQ}(Q_j^i - Q_j^{i0})^2.$$

5.5.4 Three-phase SM optimization problem

$$\min \sum_{j\in\mathcal{N}_{J,i}} \{f^i_{j,1}, f^i_{j,2}, f^i_{j,3}, f^i_{j,4}\} \tag{5.3a}$$

$$f^i_{1,j} \succ f^i_{2,j} \succ f^i_{3,j} \succ f^i_{4,j}, \quad \Phi = \{a, b, c\} \tag{5.3b}$$

$$f^i_{j,1} = -C^i_j \left(\sum_{\phi\in\Phi} (P^{i,\phi}_j - P^{i0,\phi}_j)^2 + (Q^{i,\phi}_j - Q^{i0,\phi}_j)^2 \right)$$

$$f^i_{j,2} = \sum_{\phi\in\Phi} \mu^{iP,\phi}_j P^{i,\phi}_j + \mu^{iQ,\phi}_j Q^{i,\phi}_j$$

$$f^i_{j,3} = -\sum_{\phi\in\Phi} (\delta P^{i,\phi}_j + \delta Q^{i,\phi}_j)$$

$$f^i_{j,4} = \beta^{iP}_j \sum_{\phi\in\Phi} \left(P^i_j - P^{i0}_j\right)^2 + \beta^{iQ}_j \sum_{\phi\in\Phi} \left(Q^i_j - Q^{i0}_j\right)^2$$

subject to:

$$P^{i,\phi}_j - \delta P^{i,\phi}_j \geq \underline{P}^{i,\phi}_j \; Q^{i,\phi}_j - \delta Q^{i,\phi}_j \geq \underline{Q}^{i,\phi}_j \tag{5.3c}$$

$$P^{i,\phi}_j + \delta P^{i,\phi}_j \leq \overline{P}^{i,\phi}_j, \; Q^{i,\phi}_j + \delta Q^{i,\phi}_j \leq \overline{Q}^{i,\phi}_j \tag{5.3d}$$

$$\delta P^{i,\phi}_j, \; \delta Q^{i,\phi}_j \geq 0, \; 0 \leq \mu^{iP}_j \leq \overline{\mu}^{iP}, 0 \leq \mu^{iP}_j \leq \overline{\mu}^{iQ} \tag{5.3e}$$

$$\sum_{t_s}^{t_s+\Delta t_p} \sum_{j\in\mathcal{N}_{J,i}} \sum_{\phi\in\Phi} \left(\mu^{iP,\phi}_j(t) P^{i,\phi}_j(t) + \mu^{iQ,\phi}_j(t) Q^{i,\phi}_j(t) \right) \Delta t_s$$

$$\leq \sum_{\phi\in\Phi} \left(\mu^{P^*,\phi}_i(\hat{t}_p) P^{\phi^*}_i(\hat{t}_p) + \mu^{Q^*,\phi}_i(\hat{t}_p) Q^{\phi^*}_i(\hat{t}_p) \right) \Delta t_p \tag{5.3f}$$

$$\sum_{j\in\mathcal{N}_{J,i}} P^{i,\phi}_j(t_s) = P^{\phi^*}_i(\hat{t}_p), \quad \sum_{j\in\mathcal{N}_{J,i}} Q^{i,\phi}_j(t_s) = Q^{\phi^*}_i(\hat{t}_p) \tag{5.3g}$$

5.5.5 SM monitoring and resilience scores

The final stage in the SM is monitoring. During the market operation, the responses of each SMA j to the market schedules, in terms of its actual DER injections $\hat{P}^i_j$ and $\hat{Q}^i_j$ are suitably monitored by its corresponding SRM. In addition to the market operators, we also propose the addition of two new entities, which we denote as the primary resilience manager (PRM) and the SRM, both of which provide grid functionalities, with the PRM located at the primary circuit level and the SRM at the secondary level, as shown in Fig. 5.5. With the market clearing providing the first step of awareness in the form of power available at each of the nodes at the secondary and primary level, the PRM and SRM monitor the actual injections, determine corresponding scores of commitment, trustability, and resilience (to be defined below), and communicate them using protected channels to the PRM. Not only do these entities enable a separation between grid-specific decision-making from market-specific decisions, but they also provide a pathway for mitigating the impact of any attacks that can occur through the addition of local resources, as will be shown in the following sections.

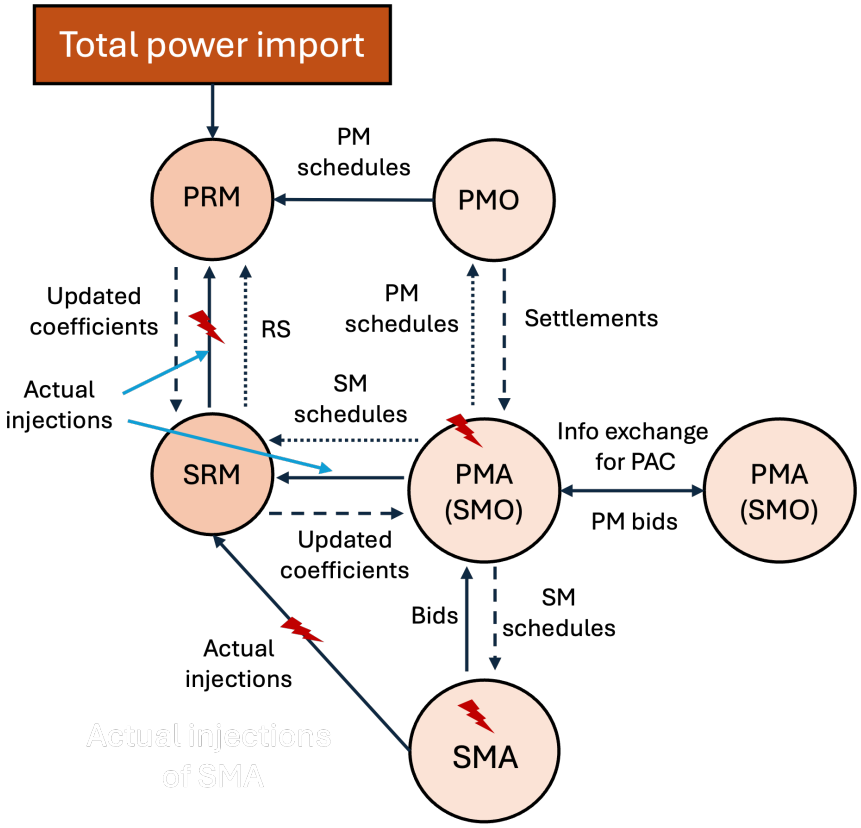

Figure 5.5: Sequence of communication steps and events leading to SA with an LEM. The red arrows indicate the entities and communication links that would be affected by an attack. A more detailed diagram can be found in Fig. 5.6.

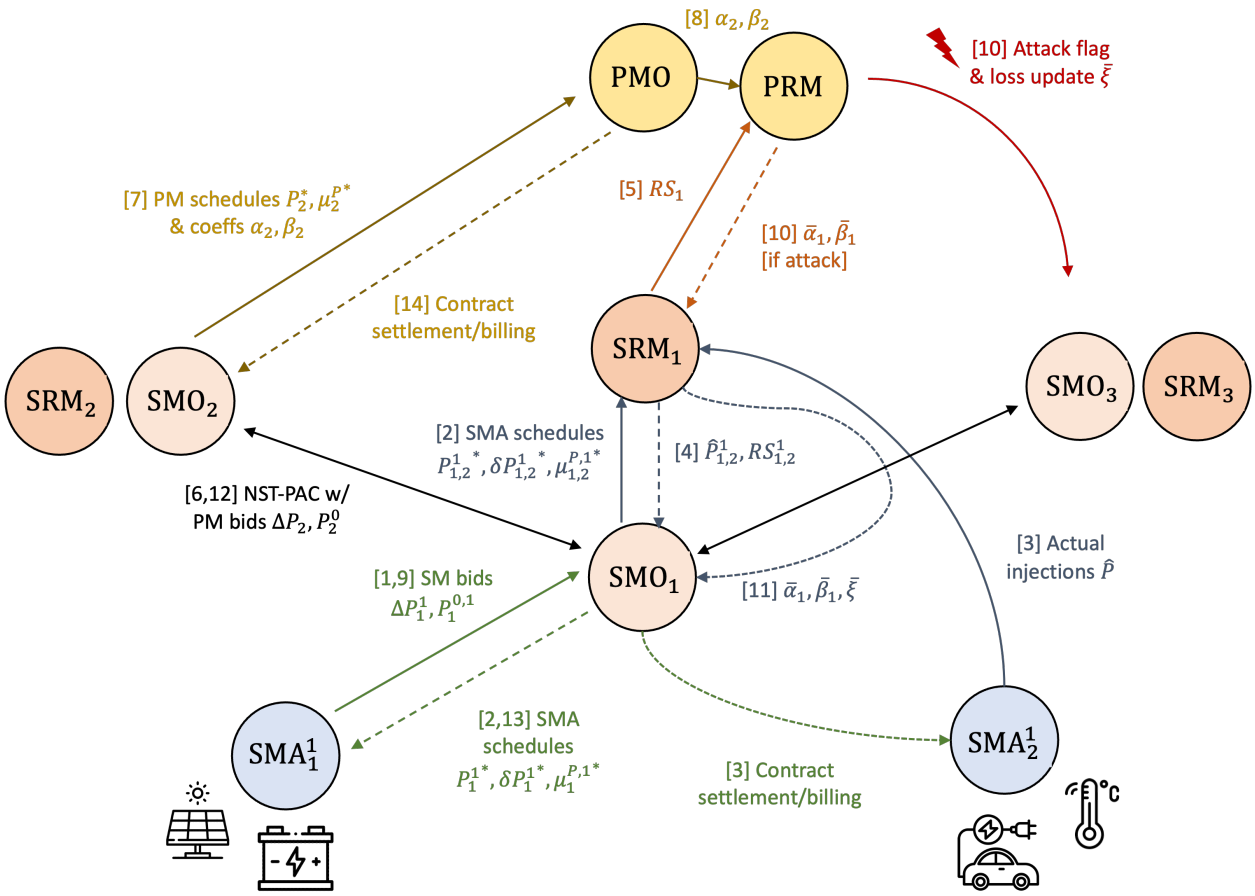

Figure 5.6: Diagram showing a more detailed communication scheme and steps for information exchange between the various market operators, agents, and resilience managers at the secondary and primary levels.

In this monitoring stage, the SRM assigns each SMA an RS that is updated constantly based on its performance in the market and susceptibility to being compromised. The RS is a weighted combination of its commitment score (CS) and trustability score (TS). Formally, for an agent j

$$\text{RS}_j = \alpha \text{CS}_j + (1-\alpha)\text{TS}_j,$$

where $\alpha \in (0,1)$ is a parameter chosen by the SRM. The CS and TS are defined below.

- *Commitment score (CS).* The CS of an agent measures its reliability in executing its cleared schedules and is updated at every SM clearing instance. The first step in updating CS_j for each agent $j \in \mathcal{N}_i$ is the computation of any relative deviation between the cleared schedule and its executed value over the past market period. A moving average is then computed to account for the past performance. Finally, a min-max normalization across all the SMAs is performed to keep $CS_j \in [0,1]$ for all j (see Section 5.6.5 for further details).

- *IoT Trustability score (TS).* The TS captures the possibility of the agents (or the devices underneath them) being compromised. TSs are computed using an FL-based anomaly detector, and like the CS, past values are used again to compute a weighted moving average. However, unlike the CS, which solely depends on power injections, the TS is a cyber-power metric [127] that also takes into account the associated cyber information, e.g., packet length, arrival time, and communication protocols, etc. (see Section 5.7 for further details).

In summary, the overall operation of the SM allows the generation of the schedules $\mathbf{y}_i^{S*} = [P_j^{i*}, Q_j^{i*}, \delta P_j^{i*}, \delta Q_j^{i*}, \mu_j^{iP*}, \mu_j^{iQ^*}]$ and RS_j^i $\forall$ SMAs j, all of which provide SA_j^i for the SRMs corresponding to all SMAs j at the primary node i. Similar measures of the resilience of large-scale networks to attacks can be found in [128].

5.6 Primary market

PM transactions happen between the PMO and the PMAs. Similar to the SM, the operation of the PM also consists of bidding, clearing, and monitoring.

5.6.1 PM bidding

Here, we describe the link between the PM and SM. As noted previously, the PM is cleared every 5 minutes while the SM operates more frequently at 1-minute intervals. Before each PM clearing, the SMOs (or PMAs) aggregate the schedules and cleared flexibilities of all their SMAs resulting from the most recent prior SM clearing (at the lower level) to submit their flexibility bid to the upper-level PM. All market bidding and clearing for both the SM and PM are based on forecasts (assuming perfect foresight) and for the very next period. The complete bid submitted by each SMO $j \in \mathcal{N}$ into the PM $\mathcal{B}_i^P$ defined as:

$$\mathcal{B}_i^P = \{P_i^0, Q_i^0, \underline{P}_i, \underline{Q}_i, \overline{P}_i, \overline{Q}_i, \alpha_i^P, \alpha_i^Q, \beta_i^P, \beta_i^Q\}. \tag{5.4}$$

$$P_i^0(t_p) = \sum_{j \in \mathcal{N}_i} P_j^{i*}(t_p),\ Q_i^0(t_p) = \sum_{j \in \mathcal{N}_i} Q_j^{i*}(t_p) \tag{5.5}$$
$$\underline{P}_i = \sum_{j \in \mathcal{N}_i} P_j^i - \delta P_j^{i*}, \overline{P}_i = \sum_{j \in \mathcal{N}_i} P_j^{i*} + \delta P_j^{i*}$$
$$\underline{Q}_i = \sum_{j \in \mathcal{N}_i} Q_j^{i*} - \delta Q_j^{i*}, \overline{Q}_i = \sum_{j \in \mathcal{N}_i} Q_j^{i*} + \delta Q_j^{i*}$$

In the above, (i) P_i^0, Q_i^0 denote the nominal baseline active and reactive power injection bids of the SMOs, (ii) $(\underline{P}_i, \underline{Q}_i)$ and $(\overline{P}_i, \overline{Q}_i)$ denote the downward and upward flexibilities (around their nominal values) in active and reactive power, respectively, (iii) α_i^P, α_i^Q are the local net generation costs, and (iv) β_i^P, β_i^Q are the flexibility disutility parameters of SMO i for P and Q, respectively. The SMO computes (iii) and (iv) as weighted averages of all their SMA retail tariffs and SMA disutility parameters, respectively, as follows:

$$\alpha_i^P = \frac{\sum_{j \in \mathcal{N}_i} \mu_j^{iP*} |P_j^{i*}|}{\sum_{j \in \mathcal{N}_i} |P_j^{i*}|}, \beta_i^P = \frac{\sum_{j \in \mathcal{N}_i} \beta_j^{iP*} |P_j^{i*}|}{\sum_{j \in \mathcal{N}_i} |P_j^{i*}|}$$

We note that standalone PMAs, such as a large industrial facility, a community solar farm, or an EV charging station, may also be present. In this case, the PMA would directly bid into the PM on its own instead of aggregating over SMAs.

5.6.2 PM clearing

At each PM clearing instance, an optimal power flow (OPF) problem is solved to optimize the PMO's objective while satisfying all grid physics and network power flow constraints. For simplicity, in this work, we consider the cost functions of all the PMAs (or SMOs) to be quadratic. The objective function utilized is a weighted linear combination of (i) maximization of social welfare, (ii) minimization of total generation costs, and (iii) minimization of electrical line losses (see Section 5.6.4 for details of these functions). The total cost includes paying the locational marginal price (LMP) λ for importing power from the transmission grid at the PCC, as well as the payments to local generator PMAs that provide net positive injections into the PM. We divide by suitable base values to convert all quantities to per unit (between 0 and 1 p.u.). Thus, it is reasonable to combine all the terms into a single objective function using a simple weighted sum.

With the objective function thus defined, the constraints are determined by the choice of the power flow model used to describe the system. Since the original alternating current OPF (ACOPF) is inherently nonconvex and NP-hard, we need to convexify the problem to make it more tractable. In this study, we considered two different approaches for this convexification. The first is a branch flow (BF) model or nonlinear *DistFlow* [129] based on a second-order conic program (SOCP) convex relaxation - this is a simpler implementation that is valid for radial and balanced networks. The second is a linear current injection (CI) model [130] based on a McCormick envelope convex relaxation that is more generally applicable to unbalanced and meshed grids common in distribution systems (in addition to radial, balanced), although this adds some overhead due to certain pre-processing steps needed.

We deployed both of these models for different use cases considered in this paper, as shown in Table 5.1. Further details of the BF and CI approaches are provided in Section 5.6.3 and Section 5.6.3, respectively. The exact set of decision variables $\mathbf{y}_i^P$ for each PMA i differs slightly depending on the OPF model used. Both models solve for the nodal power injections and voltages. However, the BF model only considers branch currents, while the CI model also considers nodal current injections. BF also models all variables as only having a single phase, while the CI models these as three-phase, complex phasor quantities. For simplicity, we have only included the single-phase formulations in the paper thus far. However, these can easily be extended to the complex three-phase representation by simply modifying all variables to 3-dimensional complex vectors instead of scalars. A three-phase extension of the SM optimization is also given in Section 5.5.4. Thus, the decision vectors for the BF model are given by $\mathbf{y}_i^{P,BF} = [P_i^G, Q_i^G, P_i^L, Q_i^L, v_i, I_{ik}]\ \forall i \in \mathcal{N},\ (ik) \in \mathcal{E}$ and for the CI model, $\mathbf{y}_i^P = [P_i^\phi, Q_i^\phi, V_i^{\phi,R}, V_i^{\phi,I}, I_i^{\phi,R}, I_i^{\phi,I}, I_{ik}^{\phi,R}, I_{ik}^{\phi,I}]$, where phases $\phi \in \{a, b, c\}$ are the phases and $\mathcal{E}$ is the set of all network edges or branches.

5.6.3 Power system models

Branch flow model

The branch flow OPF problem is formally stated as follows, where R and X denote the network resistance and reactance matrices, respectively, v and I denote the nodal voltage magnitudes and branch currents respectively, and $\mathcal{E}$ denotes the set of all edges in the network. The primal decision variables here for each PMA i are $\mathbf{y}_i^P = [P_i^G, Q_i^G, P_i^L, Q_i^L, v_i, I_{ki}]\ \forall i \in \mathbf{N},\ (ik) \in \mathcal{E}$.

$$\min_{\mathbf{y}^P}\ f^{S-W}(\mathbf{y}^P) \tag{5.6}$$

subject to:

$$\begin{aligned}
&v_i - v_k = \left(R_{ki}^2 + X_{ki}^2\right)|I_{ki}|^2 - 2\left(R_{ki}P_{ki} + X_{ki}Q_{ki}\right)\\
&P_i^G - P_i^L = -P_{ki} + R_{ki}|I_{ki}|^2 + \sum_{k:(i,k)\in\mathcal{E}} P_{ik}\\
&Q_i^G - Q_i^L = -Q_{ki} + X_{ki}|I_{ki}|^2 + \sum_{k:(ik)\in\mathcal{E}} Q_{ik}\\
&P_{ki}^2 + Q_{ki}^2 \le \overline{S}_{ki}^2,\ P_{ki}^2 + Q_{ki}^2 \le v_i|I_{ki}|^2,\ \underline{v}_i \le v_i \le \overline{v}_i\\
&\underline{P}_i^G \le P_i^G \le \overline{P}_i^G,\ \underline{P}_i^L \le P_i^L \le \overline{P}_i^L\\
&\underline{Q}_i^G \le Q_i^G \le \overline{Q}_i^G,\ \underline{Q}_i^L \le Q_i^L \le \overline{Q}_i^L
\end{aligned} \tag{5.7}$$

Current injection model

The primal decision variables for each SMO i obtained by solving the optimization problem $\mathbf{y}_i^P = [P_i^\phi, Q_i^\phi, V_i^{\phi,R}, V_i^{\phi,I}, I_i^{\phi,R}, I_i^{\phi,I}, I_{ik}^{\phi,R}, I_{ik}^{\phi,I}]$ consists of (i) active ($P_i^{\phi*}$) and reactive ($Q_i^{\phi*}$) power setpoints (ii) real and imaginary components of nodal voltages ($V_i^{\phi,R*}, V_i^{\phi,I*}$) and current injections ($I_i^{\phi,R*}, I_i^{\phi,I*}$). Note that these are solved for each non-zero phase $\phi \in \mathcal{P} = \{a, b, c\}$.

The CI-OPF problem formulation is given by:

$$\min_x f^{obj}(x) \tag{5.8a}$$

$$I^R = \mathrm{Re}(\mathrm{Y}V),\ I^I = \mathrm{Im}(\mathrm{YV}) \tag{5.8b}$$

$$P_i^\phi = V_i^{\phi,R} I_i^{\phi,R} + V_i^{\phi,I} I_i^{\phi,I} \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{5.8c}$$

$$Q_i^\phi = -V_i^{\phi,R} I_i^{\phi,I} + V_i^{\phi,I} I_i^{\phi,R} \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{5.8d}$$

$$(I_{ik}^{\phi,R})^2 + (I_{ik}^{\phi,I})^2 \leq \overline{I_{ik}^{\phi}}^2 \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P}, (ik) \in \mathcal{E} \tag{5.8e}$$

$$\underline{V_i^{\phi}}^2 \leq (V_i^{\phi,R})^2 + (V_i^{\phi,I})^2 \leq \overline{V_i^{\phi}}^2 \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{5.8f}$$

$$\underline{P_i^\phi} \leq P_i^\phi \leq \overline{P_i^\phi},\ \underline{Q_i^\phi} \leq Q_i^\phi \leq \overline{Q_i^\phi} \tag{5.8g}$$

where Y is the 3-phase bus admittance matrix for the network, and V and I are matrices of nodal voltages and currents, respectively. Problem Eq. (5.8) is nonconvex due to bilinear constraints Eqs. (5.8c) and (5.8d), and the ring constraint Eq. (5.8f) on voltage magnitudes. We obtain a convex relaxation by using McCormick envelopes (MCE), which represent the convex hull of a bilinear product $w = xy$ by using upper and lower limits on x, y. Thus, we replace the bilinear equality with a series of linear inequalities, denoted as $\mathrm{MCE}(w) = \{w = xy : x \in [\underline{x}, \overline{x}], y \in [\underline{y}, \overline{y}]\}$:

$$MCE(w, \underline{x}, \overline{x}, \underline{y}, \overline{y}) = \begin{cases} w \geq \underline{x}y + x\underline{y} - \underline{x}\underline{y} \\ w \geq \overline{x}y + x\overline{y} - \overline{x}\overline{y} \\ w \leq \underline{x}y + x\overline{y} - \overline{x}\underline{y} \\ w \leq \overline{x}y + x\underline{y} - \underline{x}\overline{y} \end{cases} \tag{5.9}$$

We introduce auxiliary variables for each of the four bilinear terms $\{a_i^\phi, b_i^\phi, c_i^\phi, d_i^\phi\} = \{V_i^{\phi,R} I_i^{\phi,R}, V_i^{\phi,I} I_i^{\phi,I}, V_i^{\phi,R} I_i^{\phi,I}, V_i^{\phi,I} I_i^{\phi,R}\}$ allowing us to convert constraints Eqs. (5.8c) and (5.8d) to linear constraints with MCE constraints on each of the auxiliary variables. We also need additional constraints on the nodal current injections and nodal voltages in order to define the MCE constraints. These voltage and current bounds can be determined by applying a suitable preprocessing method using the nodal P and Q limits from the SMO bids [130]. The resulting bounds will also implicitly satisfy constraints Eq. (5.8e) and Eq. (5.8f). Thus, we can replace constraints Eqs. (5.8c) to (5.8f) with the following set of constraints in order to obtain the relaxed CI-OPF problem, which reduces to a linear program that can be solved easily. However, we do incur the overhead of computing the tightest possible V and I bounds to obtain a good convex relaxation, which in turn ensures

that the relaxed solutions are feasible for the original problem.

$$P_i^{\phi} = a_i^{\phi} + b_i^{\phi}, \;\; Q_i^{\phi} = -c_i^{\phi} + d_i^{\phi} \quad \forall i \in \mathcal{N}, \phi \in \mathcal{P} \tag{5.10a}$$

$$\underline{I_i^{\phi,R}} \leq I_i^{\phi,R} \leq \overline{I_i^{\phi,R}}, \; \underline{I_i^{\phi,I}} \leq I_i^{\phi,I} \leq \overline{i^{\phi,I}} \tag{5.10b}$$

$$\underline{V_i^{\phi,R}} \leq V_i^{\phi,R} \leq \overline{V_i^{\phi,R}}, \; \underline{V_i^{\phi,I}} \leq V_i^{\phi,I} \leq \overline{V_i^{\phi,I}} \tag{5.10c}$$

$$a_i^{\phi} \in MCE(V_i^{\phi,R} I_i^{\phi,R}, \underline{V_i^{\phi,R}}, \overline{V_i^{\phi,R}}, \underline{I_i^{\phi,R}}, \overline{I_i^{\phi,R}}) \tag{5.10d}$$

$$b_i^{\phi} \in MCE(V_i^{\phi,I} I_i^{\phi,I}, \underline{V_i^{\phi,I}}, \overline{V_i^{\phi,I}}, \underline{I_i^{\phi,I}}, \overline{I_i^{\phi,I}}) \tag{5.10e}$$

$$c_i^{\phi} \in MCE(V_i^{\phi,R} I_i^{\phi,I}, \underline{V_i^{\phi,R}}, \overline{V_i^{\phi,R}}, \underline{I_i^{\phi,I}}, \overline{I_i^{\phi,I}}) \tag{5.10f}$$

$$d_i^{\phi} \in MCE(V_i^{\phi,I} I_i^{\phi,R}, \underline{V_i^{\phi,I}}, \overline{V_i^{\phi,I}}, \underline{I_i^{\phi,R}}, \overline{I_i^{\phi,R}}) \tag{5.10g}$$

5.6.4 Objective functions for optimization in the Primary Market

We define the following functions

$$\begin{aligned} f^P(\mathbf{y}^P) = \sum_{i \in \mathcal{N}} f_i^P(\mathbf{y}_i^P) = \sum_{i \in \mathcal{N}} \Big[& f_i^{\text{Load-Disutil}}(\mathbf{y}_i^P) \\ & + f_i^{\text{Gen-Cost}}(\mathbf{y}_i^P) \Big] + \xi \Big[\sum_{(ki) \in \mathcal{E}} f_{ki}^{\text{Loss}}(\mathbf{y}_i^P) \Big] \end{aligned} \tag{5.11}$$

$$f_i^{\text{Load-Disutil}}(\mathbf{y}_i^P) = \beta_i^P (P_i^L - P_i^{L0})^2 + \beta_i^Q (Q_i^L - Q_i^{L0})^2 \tag{5.12}$$

$$f_i^{\text{Gen-Cost}}(\mathbf{y}_i^P) = \begin{cases} \alpha_i^P (P_i^G)^2 + \alpha_i^Q (Q_i^G)^2, \\ \lambda_i^P P_i^G + \lambda_i^Q Q_i^G, \text{if } i \qquad \text{is PCC} \end{cases} \tag{5.13}$$

$$f_{ki}^{\text{Loss}}(\mathbf{y}_i^P) = R_{ki} |I_{ki}|^2 \tag{5.14}$$

The objective function used in Eq. (5.11) used is a weighted linear combination of (i) maximizing social welfare in Eq. (5.12), (ii) minimizing total generation costs in Eq. (5.13) and (iii) minimizing electrical line losses in Eq. (5.14). The total cost includes paying the locational marginal price (LMP) λ for importing power from the transmission grid at the point of common coupling (PCC), as well as the payments to local generator PMAs that provide net positive injections into the PM. We divide by suitable base values to convert all quantities to per unit (between 0 and 1 p.u.). Thus, it is reasonable to combine all the terms into a single objective function using a simple weighted sum. The hyperparameter ξ controls the tradeoff between penalizing line losses versus optimizing for other objectives. The coefficients α_i, β_i, are communicated by each PMA i as part of their bids, while ξ is a global hyperparameter common to all PMAs and determined by the PMO. Here, R denotes the network resistance matrix and $\mathcal{E}$ denotes the set of all edges in the network.

5.6.5 Computation of commitment scores

We describe here the details of computing the commitment reliability score, mentioned in Section 5.5.5. From the SM clearing, the SMAs j are directed by their SMO i to keep their

net injections within the intervals $[P_j^{i*} - \delta P_j^{i*}, P_j^{i*} + \delta P_j^{i*}]$. We first compute the deviations (if any) in their actual responses $\hat{P}_j^i$ from this range, where $[\![\cdot]\!]$ denotes the indicator function:

$$\begin{aligned} e_j^{iP}(t_s) = [\![\hat{P}_j^i > \overline{P}_j^{i*}]\!](\hat{P}_j^i - \overline{P}_j^{i*}) + [\![\hat{P}_j^i < \underline{P}_j^{i*}]\!](\underline{P}_j^{i*} - \hat{P}_j^i) \\ + [\![\underline{P}_j^{i*} \leq \hat{P}_j^i \leq \overline{P}_j^{i*}]\!] \max(\hat{P}_j^i - \overline{P}_j^{i*}, \underline{P}_j^{i*} - \hat{P}_j^i) \end{aligned} \tag{5.15}$$

We then obtain relative deviations by comparing these with the magnitudes of their corresponding baseline setpoints:

$$\overline{e_j^{iP}}(t_s) = \frac{e_j^{iP}(t_s)}{|P_j^{i*}(t_s)|}, \quad \overline{e_j^{iQ}}(t_s) = \frac{e_j^{iQ}(t_s)}{|Q_j^{i*}(t_s)|} \tag{5.16}$$

These are then normalized to unit vectors to compare the deviations among all SMAs overseen by the SMO. This allows the SMO to assess their relative performance across all its SMAs.

$$\widetilde{\mathbf{e^{iP}}}(t_s) = \frac{\overline{\mathbf{e^{iP}}}(t_s)}{\|\overline{\mathbf{e^{iP}}}(t_s)\|}, \ \widetilde{\mathbf{e^{iQ}}}(t_s) = \frac{\overline{\mathbf{e^{iQ}}}(t_s)}{\|\overline{\mathbf{e^{iQ}}}(t_s)\|} \tag{5.17}$$

The scores are then updated, with the score being increased when the SMAs follow their contracts and decreased otherwise:

$$C_j^i(t_s) = \begin{cases} 1 & \text{if } t_s = 0 \\ C_j^i(t_s - 1) - \frac{\widetilde{e_j^{iP}}(t_s) + \widetilde{e_j^{iQ}}(t_s)}{2} & \text{if } t_s > 0 \end{cases} \tag{5.18}$$

Finally, we perform min-max normalization across all the SMAs' scores to ensure that $0 \leq C_j \leq 1 \ \forall$ SMAs j.

$$\overline{C}_j^i = \frac{C_j^i - \max_j C_j^i}{\max_j C_j^i - \min_j C_j^i}$$

5.7 Trustability scores and resilience metrics

5.7.1 Computation of IoT trustability scores

The IoT trustability score (TS) is computed utilizing the federated self-learning concept [127]. Anomalies in IoT data are the key factor in the formulation of the IoT TS. Another contributing factor is the IoT device's market commitment history. More details of the features are shown in Table 5.2.

To detect anomalies, we learn the IoT data's expected behavior and prediction for the short time steps. Predicted data is compared with measured data for anomaly detection. For prediction, we are using an autoencoder neural network for federated unsupervised learning. There will be one autoencoder model for each IoT device to train on its physical data, and one more autoencoder model to train only on IoT network packet data. For each type of data, there is a tolerance value T_{err} for the relative error (RE). If any data point (DP) crosses T_{err},

Table 5.2: Features considered for each type of data.

Data Source	Features
IoTs network packet	Source/Destination IP, Source/Destination port, Packet length, Protocols, Intra-packet arrival time
HVAC	Timestamp, Load, Indoor temperature, outdoor temperature, Temperature setpoint, Indoor area, Building thermal insulation
PV	Timestamp, Power generation, Rating, Solar irradiance
Battery	Timestamp, Charging/Discharging rate, SoC, KW capacity
EV	Timestamp, Charging rate, SoC

then that is flagged as an anomalous data point (ADP). So, for any reporting time period Δt, non-anomaly ratio (NAR) is calculated using,

$$NAR = 1 - \frac{\text{Total ADP number over } \Delta t}{\text{Total DP number over } \Delta t} \tag{5.19}$$

Next, the cumulative non-anomaly ratio ($CNAR$) is computed, where T is the fixed total time period and is always divisible by Δt.

$$CNAR_t = \sum_{j=1}^{\frac{T}{\Delta t}} \frac{T}{j\Delta t} NAR_{t-j\Delta t} \tag{5.20}$$

IoT Trustability Score (TS) for time t and building/house i is calculated:

$$TS_{t,i} = w_t \times NAR_t + w_{t-} \times \frac{CNAR_t}{CNAR_{max}} \tag{5.21}$$

where

$$(I)\ w_t \geq w_{t-} \qquad (II)\ w_t + w_{t-} = 1 \tag{5.22}$$

Here, $CNAR_{max}$ is calculated using (5.20) with the maximum NAR being $NAR = 1$ for the whole time period T. Finally, to get the overall TS_t of any observation node with IoTs at time t, we average the $TS_{t,i}$ of all the clients i of that observation node to calculate TS_t:

$$TS = \frac{\sum_{i=1}^{M} TS_{t,i}}{M} \tag{5.23}$$

where M is the total number of clients or buildings/houses at that observation node.

5.7.2 Secondary Transformer and Primary Node Resiliency Metric (STNR and PNR)

Secondary transformer node resiliency (STNR) is computed using multiple resiliency factors and TS.

$$STNR_j = \prod_{i=1}^{n_c} F_i^{W_i} \tag{5.24}$$

where n_c is the total number of factors for the category of the secondary level node, F_i is the value for each factor, and W_i is the normalized weight for each factor. These factors f influencing resiliency are determined and assigned weights to aggregate into the PNR score. Factors that can be determined directly from the secondary level configuration are described in Fig. 5.7. All the device and communication vulnerabilities present at the secondary (DCVS) level of a primary node are identified using the National Vulnerability Database (NVD) [131]. Then the DCVS factor is calculated as,

$$DCVS = \frac{1}{\sum_{i=1}^{N_s} CVSS_i} \tag{5.25}$$

where N_s is the number of total vulnerabilities present at the secondary level. Here, the common vulnerability scoring system (CVSS) is one of several methods to measure the impact of vulnerabilities in devices known as Common Vulnerabilities and Exposures (CVEs). It is an open set of standards used to assess the vulnerability of software and assign severity along a scale of 0-10. The National Institute of Standards and Technology (NIST) analyzes all identified vulnerabilities and lists these in the NVD. In case of the absence of any vulnerability, DCVS will be equal to 1.

Weight assignment and aggregation are managed by fuzzy multiple-criteria decision-making (MCDM), specifically the fuzzy analytic hierarchy process (Fuzzy AHP). A weighted average of the STNR results in the primary node resiliency (PNR):

$$PNR_k = \frac{\sum_{j=1}^{n}(STNR_j \times W_j)}{\sum_{j=1}^{n} W_j} \tag{5.26}$$

where W_i is the weighted coefficient for the i^{th} secondary feeder node.

5.7.3 Distribution System Resiliency (DSR)

Let $F = (f_{ij}) \in R_+^{m \times n}$ be the factors value matrix, where f_{ij} is value of factor i of primary node j. The higher the value of f_{ij}, the more the node will contribute to the resiliency metric in regard to that factor. Now, following the data envelopment analysis (DEA) method, each node p can choose a set of weights $w^p = (w_1^p, ...w_m^p)$, where $\sum_{i=1}^{m} w_i^p = 1$. Now the relative contribution (RC) of the node p to the total contribution of all the nodes towards DSR, as measured by node p's weight selection, can be evaluated as,

$$RC^p = \frac{\sum_{i=1}^{m} w_i^p f_{ip}}{\sum_{i=1}^{m} w_i^p \sum_{j=1}^{n} (f_{ij})} \tag{5.27}$$

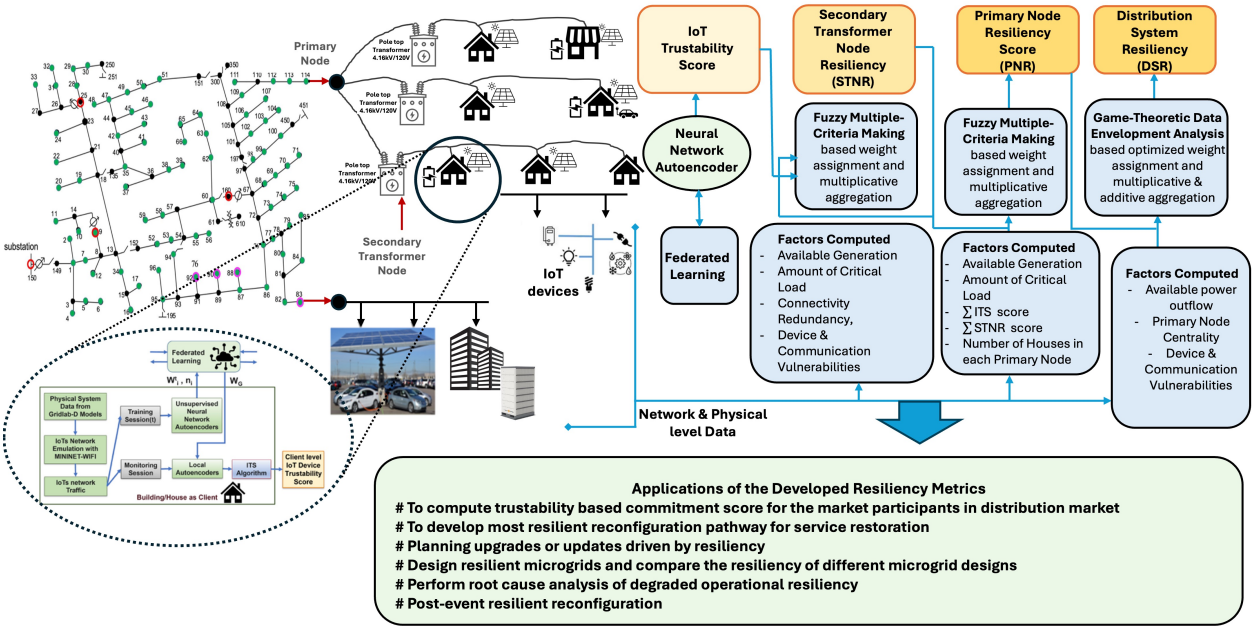

Figure 5.7: Overview of the developed resilience score for the distribution system with IoTs

Now, each node wants to maximize this ratio in Eq. (5.27) to have the best set of weights so that they can contribute to the maximum possible value in DSR. Using the weight vector for each node, a combination of multiplicative and additive methods are used to get the DSR.

$$DSR = \sum_{j=1}^{n} \left(\prod_{i=1}^{m} (f_{ij})^{w_i^j} \right) \tag{5.28}$$

Details related to the computation of DSR are shown in Fig. 5.7.

5.7.4 Distributed optimization for PM clearing

Since the number of nodes (and hence the number of PMAs) in a primary feeder could be arbitrarily large, rather than using traditional centralized optimization solvers, we employ a distributed proximal atomic coordination (PAC) algorithm [132] to solve the OPF using peer-to-peer communication between the agents. This also helps preserve data privacy since each PMA only needs to exchange limited information with its immediate neighbors. A distributed approach also enables the PMAs to clear the market independently of the PMO, alleviates the communication burden, and reduces latencies since PMAs do not need to send all their data to a centralized entity, thus allowing for scalability. This is achieved by a process called atomization, wherein the overall global optimization problem is decomposed into several local optimization problems called atoms for each PMA. The constraints can also similarly be decoupled. However, certain network constraints also depend on other PMAs' variables. To deal with this, we include additional coupling or consensus constraints to ensure consistency. We also used an enhanced variant of PAC known as NST-PAC that employs Nesterov (NST) acceleration and has enhanced privacy features by further masking

the variables exchanged between atoms (i.e., the PMAs) [133]. After a sufficient number of iterations, both the PAC and NST-PAC algorithms provably converge to globally optimal and feasible solutions $\mathbf{y}^{P*} = \{\mathbf{y}_i^{P*}\}$ for each of the PMAs. These cleared market schedules are communicated by the PMAs to their respective SRMs as well as to the PMO.

For a given global optimization (primal) problem with equality and inequality constraints for K number of nodes (or agents):

$$\min_x \sum_{i=1}^{K} f_i(x) \text{ s.t. } Gx = b, \quad Hx \leq d \tag{5.29}$$

We can decompose this into $\mathcal{S} = \{S_1, S_2, \dots S_K\}$ coupled optimization problems, known as atoms (representing each SMO i). We separate the vector of all decision variables x into two sets: $\mathcal{L} = \{L_i, \forall i \in [K]\}$ and $\mathcal{O} = \{O_i, \forall i \in [K]\}$ which is a partition of decision variables into those that are owned and copied by atom i, respectively. We can similarly also decompose the constraints into sets owned by each atom $\mathcal{C} = \{C_i, \forall i \in [K]\}$. These variable copies across multiple atoms can then be used to satisfy coupled constraints and global objectives. Note that for a number K, $[K] = \{1, 2, \dots K\}$.

The decomposed (or atomized) optimization problem is shown in Eq. (5.30), where a_j and $f_j(a_j)$ are the primal decision variables (both owned and copies) and individual objective functions corresponding to each SMO atom, respectively. G_j and H_j are the atomic constraint submatrices of G and H, while b_j and d_j are subvectors of b and d of the right hand side constraint vectors b and d, respectively. B is the directed graph incidence matrix defining the owned and copied atomic variables. This incidence matrix allows us to fully parallelize the distributed optimization by defining coordination or consensus constraints, which enforce that all the copied variables for each atom j must equal the values of their corresponding owned values in every other atom $i \neq j$. B_j and B^j denote the incoming and outgoing edges for atom j respectively. Here

$$\begin{aligned} &\min_{a_j} \sum_{j \in K} f_j\left(a_j\right) \\ &\text{s.t. } G_j a_j = b_j,\ H_j a_j \leq d_j,\ B_j a = 0\ \forall j \in [K] \\ &B_{im} \triangleq \begin{cases} -1, & \text{if } i \text{ is 'owned" and } m \text{ a related "copy"} \\ 1, & \text{if } m \text{ is "owned" and } i \text{ a related "copy"} \\ 0, & \text{otherwise} \end{cases} \end{aligned} \tag{5.30}$$

The augmented Lagrangian is first atomized or decomposed for each node or SMO, introducing dual variables η and ν corresponding to primal equality and coordination constraints, respectively. Note that the inequality constraints are handled directly during the primal

minimization step by appropriately defining the feasible set.

$$\begin{aligned}\mathcal{L}(a,\eta,\nu) &= \sum_{j\in K}\left[f_j\left(a_j\right)+\eta_j^T\left(G_j a_j-b_j\right)+\nu_j^T B_j a\right]\\ &= \sum_{j\in K}\left[f_j\left(a_j\right)+\eta_j^T\left(G_j a_j-b_j\right)+\nu^T B^j a_j\right]\\ &\triangleq \sum_{j\in K}\mathcal{L}_j\left(a_j,\eta_j,\nu\right)\end{aligned} \tag{5.31}$$

PAC algorithm

We can then apply the prox-linear approach of [134] to Eq. (5.31) and obtain the proximal atomic coordination (PAC) algorithm [132, 135]:

$$\begin{aligned}a_j[\tau+1] &= \underset{a_j\in\mathbb{R}^{|T_j|}}{\operatorname{argmin}}\left\{\begin{array}{l}\mathcal{L}_j\left(a_j,\bar{\mu}_j[\tau],\bar{\nu}[\tau]\right)\\ +\frac{1}{2\rho}\left\|a_j-a_j[\tau]\right\|_2^2\end{array}\right\}\\ \mu_j[\tau+1] &= \mu_j[\tau]+\rho\gamma_j\tilde{G}_j a_j[\tau+1]\\ \bar{\mu}_j[\tau+1] &= \mu_j[\tau+1]+\rho\hat{\gamma}_j[\tau+1]\tilde{G}_j a_j[\tau+1]\end{aligned}$$

Communicate a_j for all $j\in[K]$ with neighbors

$$\begin{aligned}\nu_j[\tau+1] &= \nu_j[\tau]+\rho\gamma_j[B]^{O_j}a[\tau+1]\\ \bar{\nu}_j[\tau+1] &= \nu_j[\tau+1]+\rho\hat{\gamma}_j[\tau+1][B]^{O_j}a[\tau+1]\end{aligned}$$

Communicate $\bar{\nu}_j$ for all $j\in[K]$ with neighbors.

The primal and dual variables are initialized as follows, $\forall j\in[K]$:

$$\begin{aligned}a_j[0] &\in \mathbb{R}^{||T_j|}\\ \mu_j[0] &= \rho\gamma_j\tilde{G}_j a_j[0]\\ \bar{\mu}_j[0] &= \mu_j[0]+\rho\hat{\gamma}_j[0]\tilde{G}_j a_j[0]\\ \nu_j[0] &= \rho\gamma_j[B]^{O_j}a[0]\\ \bar{\nu}_j[0] &= \nu_j[0]+\rho\hat{\gamma}_j[0][B]^{O_j}a[0]\end{aligned}$$

NST-PAC algorithm

This work employs an enhanced, accelerated version called NST-PAC developed in [133]. It is a primal-dual method incorporating both $L2$ and proximal regularization terms. The convergence speed is increased by using time-varying gains and Nesterov-accelerated gradient updates for both the primal and dual variables. The iterative NST-PAC algorithm consists

of the following steps at each iteration τ:

$$
\begin{aligned}
a_j[\tau+1] = \underset{a_j}{\operatorname{argmin}} \big\{ & \mathcal{L}_j\left(a_j, \hat{\eta}_j[\tau], \hat{\nu}[\tau]\right) \\
& + \frac{\rho_j \gamma_j}{2} \left\| G_j a_j - b_j \right\|_2^2 + \frac{\rho_j \gamma_j}{2} \left\| B_j a_j \right\|_2^2 \\
& + \frac{1}{2\rho_j} \left\| a_j - a_j[\tau] \right\|_2^2 \big\} \\
\hat{a}_j[\tau+1] = & \, a_j[\tau+1] + \alpha_j[\tau+1]\left(a_j[\tau+1] - a_j[\tau]\right) \\
\eta_j[\tau+1] = & \, \hat{\eta}_j[\tau] + \rho_j \gamma_j \left(G_j \hat{a}_j[\tau+1] - b_j\right) \\
\hat{\eta}_j[\tau+1] = & \, \eta_j[\tau+1] + \phi_j[\tau+1]\left(\eta_j[\tau+1] - \eta_j[\tau]\right) \\
\text{Communicate } & \hat{a}_j \text{ for all } j \in [K] \text{ with neighbors} \\
\nu_j[\tau+1] = & \, \hat{\nu}_j[\tau] + \rho_j \gamma_j B_j \hat{a}_j[\tau+1] \\
\hat{\nu}_j[\tau+1] = & \, \nu_j[\tau+1] + \theta_j[\tau+1]\left(\nu_j[\tau+1] - \nu_j[\tau]\right) \\
\text{Communicate } & \hat{\nu}_j \text{ for all } j \in [K] \text{ with neighbors}
\end{aligned}
\tag{5.32}
$$

The algorithm further protects privacy by masking both the primal and dual variables. Masking is implemented by using iteration-varying and atom-specific parameters $\alpha_j[\tau]$, $\phi_j[\tau]$ and $\theta_j[\tau]$. Masking the dual variables (or shadow prices), in particular, is desirable since these may reveal sensitive data related to costs, operating constraints, or other preferences of SMOs. Instead, masked variables $\hat{a}$ and $\hat{\nu}$ are exchanged between atoms. By iteratively solving the local, decomposed optimization problems across all SMOs, NST-PAC (and PAC) provably converge to the globally optimal ACOPF (relaxed) solutions for the whole primary feeder [132, 133].

5.7.5 PM monitoring and resilience scores

During the actual market operation, the injections $\hat{P}_i$ and $\hat{Q}_i$ from the DERs at PMA j are monitored by their SRM. These could be either from standalone PMAs or aggregated information from all the SMAs at a given PMA. The SRM also assembles resilience scores RS_i for each PMA i. This is done through aggregation (via a weighted average) of $\mathrm{RS}_j^i \; \forall j \in \mathcal{N}_i$. The RSs for standalone PMAs can also be directly computed at the SRM using their monitored injections. $\mathbf{y}_i^P$ and RS_i thus provide complete SA at each PMA node i. All SRMs send this information to the PRM so that the PRM has complete SA of all PMAs. This SA can then be used to redispatch the ICAs in both the PM and SM to mitigate the impact of various attacks. Further details on the mitigation strategy can be found in Section 5.10.

5.8 Reconfiguration paths

The final tool that we use in our proposed EUREICA framework is the determination of reconfiguration in the wake of islanding, which can occur if an attack or natural disaster causes an entire section of the grid to be disconnected from the main grid. Reconfiguration can also In such cases, an algorithm that determines a self-sustaining operation of the islanded

system, which is enabled by reconfiguration paths with suitable switch settings, is essential. We propose a reconfiguration algorithm (see Section 5.11 for details) that considers power flow feasibility, available distributed generators (DGs), critical load, as well as RS information, to determine switching actions to restore specific sections of the distribution feeder. The reconfiguration paths will be determined based on the available amount of generation and the amount of critical load to be supplied, which is obtained through the SA provided by the EUREICA framework. In addition, the TSs are used at the secondary feeder level to intelligently disconnect non-critical loads, thus enabling the maximum restoration of critical loads. Once the feasible paths are determined for the optimal selection of loads, the RSs for all feasible paths are computed, and the most resilient path is implemented in the system.

5.9 Attack scenarios

We now apply our EUREICA framework to enhance the resilience of the distribution grid to cyber-physical attacks under the 'black sky' scenario. This is the primary goal of this project. We use our coordination mechanism to validate the mitigation of attacks of different levels of severity, with attack magnitudes that range from 5 to 40% of the total peak load. Both grid-connected and islanded cases are studied. In all cases, we show that grid resilience can be obtained through a combination of locally available flexible assets and reconfiguration of the grid topology. In addition to numerical simulations, we report results from the validation partners as well. In the following section, we present use cases that illustrate how SA can be leveraged to ensure grid resilience in a distribution grid with a high penetration of DERs. We consider four different attack scenarios, all of which are motivated by the two large-scale attacks in [106, 124] on power grids. Disruptive attacks are assumed to occur in the form of (a) a sudden loss of generation, and/or (b) a sudden increase in load, at multiple vulnerable locations. All use cases are simulated using an IEEE 123-node test feeder; extensions to more realistic and larger networks [136] can be implemented similarly.

5.9.1 Attack 1

In this attack, it is assumed that a small percentage of generation or load resources at either the primary or secondary feeder level are compromised. In particular, it is assumed that these units are offline due to either an outage, a natural calamity, or a malicious cyber-attacker using elevated privileges to disconnect the units. In addition to the generation shortfall, it is assumed that the communication link between the market operators (PMO/SMO) and the resilience managers (PRM/SRM) is also affected by a denial of service (DoS) attack, which compromises the availability of a resource (see [137] for an attack which occurred on an sPower installation in Utah). Attack 1 draws inspiration from [106], where a malicious attacker used (i) elevated and unauthorized access to disconnect several resources, and (ii) severed communication links, to hamper operator visibility and response. While these attacks occurred at the transmission level, it is feasible that a similar impact can be had by targeting distribution grid entities, especially with the larger attack surface provided by grid-edge devices. Independently, it is possible that IoT load devices such as heating, ventilation, and air conditioning (HVAC) devices, WHs, EV chargers, or refrigerators may be attacked as

well as noted in [112]. Elements of both of these types of attacks are explored here in two different cases, 1a and 1b.

Case 1a

In this case, the grid is assumed to be subjected to a sudden increase in load at the primary feeder level (SMO or PMA) due to malicious agents. There are several large loads connected to the primary feeder, such as commercial buildings or industries, and a malicious agent can manipulate the loads in these entities to affect the grid. Typically, the grid would rely on the margin provided by grid inertia to mitigate the effect of a sudden load increase. However, in a case where the grid's resources are stretched, such as a cold snap or similar natural hazards, it is imperative that the grid-edge IoT resources be tapped to mitigate this condition. Examples of this scenario are already seen in operations, such as requests from grid operators in Alaska, Texas, and others in response to cold snaps. The operators requested customers to reduce their power consumption to support large critical loads such as chillers in hospitals. Furthermore, increased DER penetration will also lead to a loss of inertia, currently provided largely by large coal and gas plants. Case 1(b) details the performance of the proposed framework from this generation shortfall, even when the PRM does not have complete observability in the system.

Case 1b

Here, several generating resources are assumed to be unavailable at the primary feeder level (i.e., SMO or PMA). There are several scenarios that motivate this case, for example, in the case of several cloudy days in a row (affecting wind power) or unforeseen maintenance on generating units, the grid operates at a lower margin than under normal conditions. There is also the case of a malicious actor disconnecting generation resources. Here, the grid experiences a generation shortfall, and in combination with the DoS attack, the system operator (PRM) loses observability.

Case 1c

In this case, the grid is subjected to a sudden increase in load and/or corruption of distributed generation from the IoT devices, in a coordinated fashion, directly at the secondary feeder. DER IoT devices will soon be operated via cloud-based service mechanisms that allow them to be controlled remotely. Thus, a sufficiently motivated malicious actor could gain control of a large number of these to suddenly reduce generation or increase load in a coordinated fashion. We simulated a case in which a large number of DGs (such as solar PV smart inverters) are attacked at the SMA level.

5.9.2 Attack 2

A larger-scale attack is assumed to occur at the distribution grid level in the form of several DGs being corrupted, causing them to go offline. The scale of this attack is assumed to be such that the impact is felt even in the transmission grid. We will explore how SA by the PRM and SRM helps mitigate this impact. Similar to attack 1.0, this use case combines

elements of both [112] and [106]. The similarity to the latter is that the corruption is inserted in the form of outages of large DGs, while that to the former is that it introduces oscillations at the transmission level. For this purpose, we will utilize the well-known Kundur 2-area test system used to understand the transient and dynamic transmission-level impacts [138]. In particular, we will assume that there is an outage in one of the two areas (Area 2) that is load-rich, which introduces additional stress on the tie-line connecting the 2 areas (see Fig. 5.36 for a diagram of the 2-area system).

5.9.3 Attack 3

The substation transformer is located at node 150, which is connected to the main transmission grid under normal operating conditions. However, under this attack, the distribution grid is islanded from the main grid at node 150. This could be due to a multitude of factors, such as wildlife tripping the transmission line from the substation to the distribution system, or a cyber-attack (i.e., integrity or disruption attack) that trips the circuit breaker from the main grid. With the increased SA introduced through our framework, we will demonstrate that the distribution system loads can be picked up in a coordinated fashion.

5.10 Mitigation using market operators and resilience managers

Our market framework, consisting of the SM and PM, provides situational awareness (SA) in the form of available power injections at various nodes at the primary and secondary levels. Once the market is cleared, during execution, the actual injections from the SMA and PMA are monitored by the SRM and PRM, respectively (see Fig. 5.5). These injections are then utilized by these managers to compute commitment scores, trustability scores, and resilience scores (RS), as shown in Section 5.6.5 and Section 5.7.1. In what follows, we will show how the SA from the market operators and the RS from the resilience managers can be utilized to mitigate all the attacks.

As a result of continuous monitoring, any unexpected deviation from the agents' nominal performance in the form of a change in the net injection at the PCC raises a flag. Any such flag makes the operators shift from the nominal operating mode to the resilience mode. Minimal visibility regarding actual injections from all PMAs is assumed to be available. Rather, we assume that each SRM only locally observes the actual injection from the corresponding SMA, and each SRM communicates that information to the PRM. More importantly, the attack scenarios considered also assume that this important communication to the PRM from all SRMs is completely sabotaged (as was the case in the Ukraine 2015-16 attacks). Despite this loss of communication, the PRM is able to step in and mitigate the attack, as the flag raised is independent of this communication loss and is due to a physical impact of the agents' deviation from nominal performance. Subsequently, the PMO redispatches trustworthy PMAs so as to bring the power import from the bulk grid down to pre-attack levels. The new setpoints for the PMAs/SMOs are, in turn, suitably disaggregated to compute new setpoints for the SMAs through a re-dispatch by the SM. Before proceeding to the results, we propose

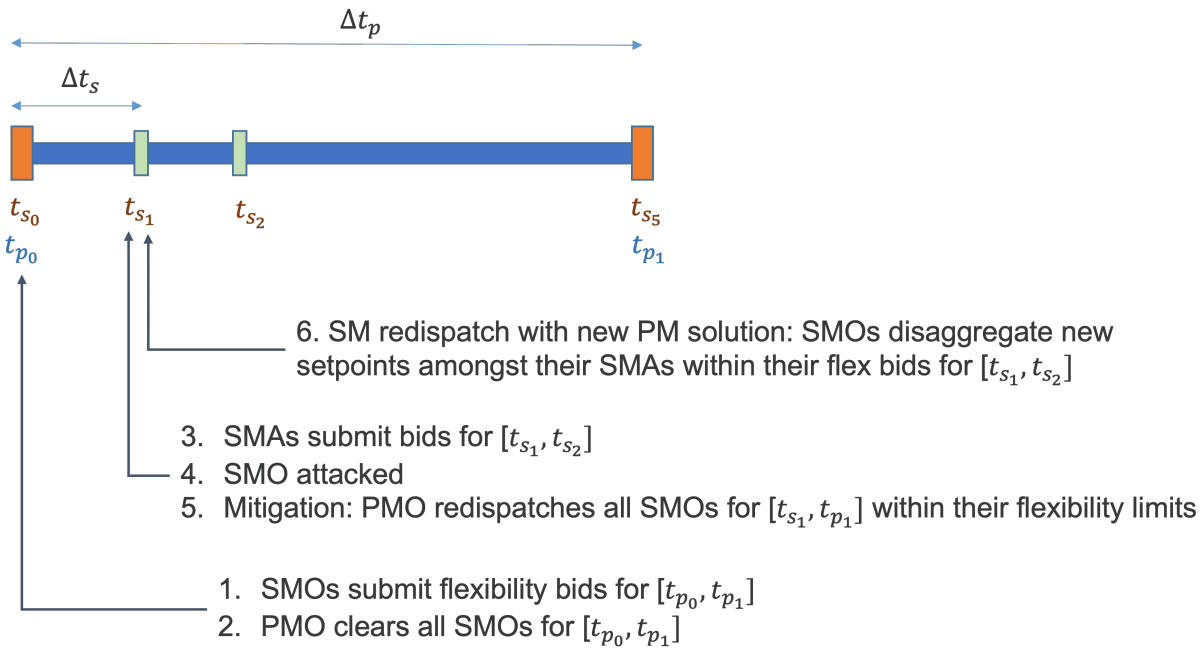

Figure 5.8: Timeline of attack detection and mitigation.

a specific mitigation strategy that leverages the SA provided by our approach.

5.10.1 Algorithm (A) for redispatch by the PMO in a balanced network

We first consider a balanced, equivalent single-phase network using the BF model. The starting point for the overall mitigation sequence is the awareness that an attack has occurred. This is realized by the PRM in the form of a change in the net load from P_{PCC} to $\overline{P}_{PCC}$, which denotes the net load from the entire primary feeder at the substation before and after the attack, respectively. This can be detected by the PRM at the substation or point of common coupling (PCC) since this is the power it imports from the main transmission grid. We propose a redispatch algorithm that the corresponding SMOs can carry out based on the ratio between these two values. To describe this redispatch algorithm, we begin with the cost function in Eq. (5.11) for the PM ACOPF problem. For ease of exposition, this can be rewritten in a simplified manner as:

$$\sum_{i=1}^{n} \left(\frac{1}{2} \alpha_i P_i^{G^2} + \beta_i \left(P_i^L - P_i^{L0} \right)^2 \right) + \xi \cdot losses \tag{5.33}$$

$$\overline{\alpha}_i = \Delta_\alpha \alpha_i,\ \overline{\beta}_i = \Delta_\beta \beta_i,\ \overline{\xi} = \Delta_\xi \xi;\ \alpha, \beta, \xi, \Delta > 0 \tag{5.34}$$

$$\Delta_\alpha = \Delta_\beta = \frac{|P_{PCC}|}{|\overline{P}_{PCC}|},\ \Delta_\xi = \frac{|\overline{P}_{PCC}|}{|P_{PCC}|} \tag{5.35}$$

We note that a change in the power import from the main grid causes $\Delta_\alpha, \Delta_\beta, \Delta_\xi$ to deviate from unity. Suppose that several distributed local generator SMOs are attacked, as in

Attack 1(a). This would increase net feeder load, i.e. $|\overline{P}_{PCC}| > |P_{PCC}|$ (note that both $P_{PCC}, \overline{P}_{PCC} < 0$ since net loads are negative injections), thus causing $\Delta_\alpha < 1$. Applying this cost coefficient update would lower the cost coefficients from α_i to α_i'. This results in dispatching more local generation from remaining online SMOs instead of importing power from the bulk grid. As the SMOs also have information about the flexibility in their corresponding SMAs in the form of $\delta P^*, \delta Q^*$ (see Section 5.5.2), the overall hierarchical PM-SM market structure automatically provides the solutions of the new dispatch. Similarly, a value of $\Delta_\beta < 1$ reduces the disutility coefficients to encourage more demand response via load shifting and/or curtailment, by utilizing the downward flexibility provided by the SMOs bidding into the PM, and subsequently also by the SMAs bidding into the SM. In contrast to these two values, when the net import from the main grid increases, then $\Delta_\xi > 1$ penalizes electrical line losses more heavily in the objective function. As a result, the redispatch discourages imports from the transmission grid in favor of dispatching more local DERs. This is because distribution grids are more lossy (have higher resistance to reactance ratios), and hence prioritizing the loss minimization makes it more efficient to utilize local generation closer to the loads being served.

After deriving the multiplicative coefficient update factors $\Delta_\alpha, \Delta_\beta, \Delta_\xi$, the PRM can broadcast these common values to all the SRMs simultaneously, who in turn send them to their corresponding SMOs. The SMOs update each of their objective function coefficients using these factors and then perform distributed optimization to redispatch the PM, resulting in new P and Q setpoints for SMOs, along with new nodal distribution LMPs (d-LMPs). This is followed by each SMO also re-dispatching their SM, in order to disaggregate the new setpoints among their SMAs. A timeline of the key events is shown in Fig. 5.8.

5.10.2 Algorithm (B) for redispatch in an unbalanced, 3-phase network

For the unbalanced 3-phase case, we use a modified algorithm for the coefficient update. The update rule here is more sophisticated since, in this case, the variables are now 3-phase vectors rather than scalars.

$$\Delta = \mathbf{P}_{PCC} - \overline{\mathbf{P}}_{PCC} \tag{5.36}$$

$$Z_i(\delta_i) = 1 + \frac{RS_i \Delta^\top \delta_i}{\mu \sum_i RS_i} \Longrightarrow \gamma_{i\delta} = \frac{1}{Z_i(\delta_i)} \tag{5.37}$$

$$\overline{\boldsymbol{\alpha}}_i = \gamma_{i\alpha}\boldsymbol{\alpha}_i, \quad \overline{\boldsymbol{\beta}}_i = \gamma_{i\beta}\boldsymbol{\beta}_i, \quad \overline{\boldsymbol{\xi}} = \left(\frac{\sum_i \gamma_{i\alpha} + \gamma_{i\beta}}{2n}\right)^{-1} \boldsymbol{\xi} \tag{5.38}$$

Note here that $\mathbf{P}_{PCC}, \overline{\mathbf{P}}_{PCC}$ are the 3-phase power imports from the tie line before and after the attack. $\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i$ are 3×1 vectors representing cost and disutility coefficient for each phase at SMO node i, and $\boldsymbol{\xi}$ is a 3-phase hyperparameter that penalizes line losses in the objective function. A DG attack that increases net load would result in $\gamma_{i\alpha}, \gamma_{i\beta} < 1$ and $\overline{\boldsymbol{\xi}} < \boldsymbol{\xi}$. Thus, these coefficient updates work using a similar intuition to algorithm (A) in that they favors local DER generation and load flexibility over transmission imports. A key difference here is that the PRM also takes into account the RS of each SMO during the redispatch so that it

relies more heavily on resilient SMOs for attack mitigation. The PRM updates the coefficients α_i, β_i and ξ to α_i', β_i' and ξ', and sends the new coefficient values to all SRMs. The SRMs send these new objective functions to the corresponding SMOs, and the rest of the mitigation procedure follows in the same manner as in the previous section.

5.11 Resilience-drive reconfiguration algorithm for attack mitigation

All possible shortest paths are computed between each generation source and critical load pairs present within the system using the graph network. If the generation is not enough to supply the total critical load, then the algorithm searches for the next available generation. This will continue until the critical load demand is met. As the generation sources are assigned to critical loads, if any source's capacity is more than the assigned load, the source's partial remaining capacity will be utilized for other loads. Once we have all the feasible paths for reconfiguration, we will compute the resiliency metric for each path (see Section 5.7.2 for how to compute the resiliency metrics and PNR), which will support the operator in finding the most resilient path to restore. The reconfiguration paths will be determined based on the stress levels of the grid and the corresponding degree of the failed SMA node, the tolerance bands and flexibility of the ICA, and the security levels and privacy needs of the SMA. This process is outlined in Fig. 5.9.

5.12 Validation platforms

In this section, we briefly describe the various software and hardware tools used by our industry and national lab partners to validate our algorithms and results.

5.12.1 PNNL

Since the proposed framework is intended to be deployed in a large-scale distribution system with a high penetration of renewables, we choose our first validation platform to be driven by HELICS (Hierarchical Engine for Large-scale Infrastructure Co-simulation), an open-source cyber-physical energy co-simulation framework for energy systems [139]. As the core engine of the platform, HELICS provides time management and data exchanges between the simulators, also known as federates. Moreover, through standard procedures and application programming interfaces (APIs), data exchange between federates is performed either as values or messages. Within this platform, GridLAB-D is the distribution system simulator used to simulate all of the use cases we discussed above. GridLAB-D is versatile, with the capability of simulating large 3-phase unbalanced distribution systems, with agent-based and information-based modeling tools and extensive data collection tools for end-use technologies and interface APIs for co-simulation [140]. Thus, it allows us to modify the standard IEEE 123-node system to include (1) residential (either single or multi-family) and commercial buildings with or without heating, ventilation, and air conditioning (HVAC) systems, (2) edge devices with IoT connectivity including HVAC type appliances, small electronics and

Algorithm 1 Compute reconfiguration pathway

Given total generation G, generation units $\mathcal{G}_i = \{G_1, G_2, \ldots, G_i\}$, total load L, load nodes $\mathcal{L}_i = \{L_1, L_2, \ldots, L_i\}$, graph network, switch settings, $\mathcal{S}_i = \{S_1, \ldots, S_n\}$
while $G < L$, **do**
  for Load L_i in $\mathcal{L}_i$ **do**
    Find shortest path $path_i$ to generation units $\mathcal{G}_i$
    Sort paths based on electrical distance in the ascending order
  end for
  Find leaf nodes $\mathcal{L}_i^n$ in $\mathcal{L}_i$
  for all $\mathcal{L}_i^n$ **do**
    Find next node in $path_i$ j to $\mathcal{G}_i$
    $TotalLoss = TotalLoss + PowerLossInPath$
    Update $\mathcal{L}_i^n$ to j
    if Leaf nodes $\mathcal{L}_i^n$ is empty **then**
      break
    end if
  end for
end while
for all *paths* **do**
  If there is a $\mathcal{S}_i$ in path, then assign $\mathcal{S}_i == 1$
  if *path* is a tree **then**
    Break
  else
    Adjust $\mathcal{S}_i$ for *path* to return to tree structure
  end if
end for
Update $\mathcal{G}_{i+1} = \{G_1, G_2, \ldots, G_{i+1}\}$
Update $\mathcal{S}_i = \{S_1^0, \ldots, S_n^0\}$
Compute resiliency metric at the primary level PNR **for the switch setting** $\mathcal{S}_N$.
Implement the *path* **with the highest** PNR.

Figure 5.9: Resilience-based reconfiguration algorithm.

Table 5.3: GridLAB-D IEEE 123-node test feeder features with IoT-enhanced model.

		Number	Capacity
Standard IEEE 123-node test feeder	Spot loads	85	3,985.7 kVA
EUREICA IEEE 123-node test feeder	Houses - Demand response (HVACs in all, WHs in 348)	1,008	variable (4 KW avg/house) (20% to 30% critical)
	Distributed generators (DGs)	380	1,745.8 kVA ($\approx$ 44% system penetration)
	PVs	207	880.84 kVA
	BESSs	173	865 kVA
	Community PVs	12	96 kVA

lighting, (3) distributed generators (DGs) such as photovoltaic (PV) panels, battery energy storage systems (BESSs), and diesel generators. The distribution model is also assumed to have all its loads connected via smart meters with load-shedding capability to regulate energy demand according to a distribution system operator's command. Specifically, the IoT-enhanced version of the IEEE 123-node system consists of the assets detailed in Table 5.3.

5.12.2 NREL

The objective of validating the EUREICA framework at NREL is to evaluate the feasibility of implementing the framework in *real-time*. Since electrons flow in the grid in real-time, it is critical that the operations proposed should also function in real-time and be in compliance with operational requirements. The Advanced Research on Integrated Energy Systems (ARIES) at NREL is a cutting-edge virtual emulation platform that encompasses actual Distributed Energy Resource (DER) hardware systems, such as wind turbines, photovoltaic (PV) arrays with controllers, batteries, and storage systems [141]. The digital real-time simulation (DRTS) cluster of ARIES is used for validating the performance of the EUREICA modules. The overall validation platform is shown in Fig. 5.10. It constitutes 5 components - (i) IoT device virtualization (using Raspberry Pis and Typhoon HIL to characterize IoT devices), (ii) Communication emulation (using analog and network connections to emulate communication at the speed of actual communication in the field) (iii) hardware-in-the-loop interface (to provide increased fidelity for components under study) (iv) digital real-time simulation (using RTDS and Typhoon HIL to emulate the power grid in real-time) (v) Time synchronization (to bring together the hardware components using a time server to ensure accuracy of simulation). This validation platform is used to determine the performance of the EUREICA framework for all the modules.

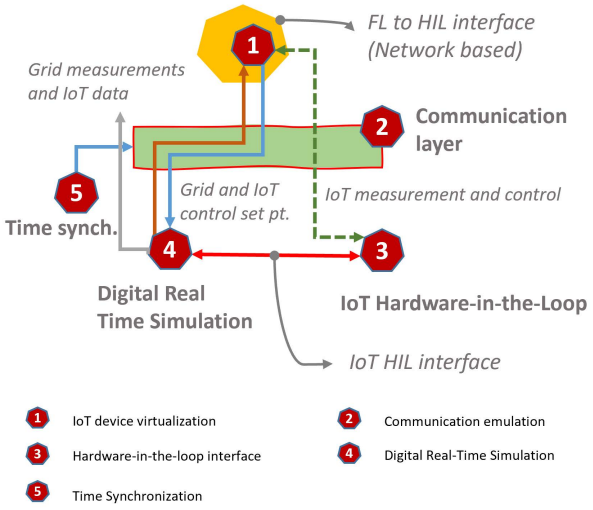

Figure 5.10: ARIES-DRTS Validation Platform at NREL

5.12.3 LTDES

A training simulator-based platform was also used to validate the EUREICA framework. In particular, the General Electric (GE) ADMS DOTS (Advanced Distribution Management System-Distribution Operations Training Simulator), used for training operators and dispatchers, was used as the validation platform. Rather than users waiting to experience challenging events on the job, dispatchers are able to familiarize themselves with advanced application functionality and gain an understanding of how they interact with other subsystems of the ADMS. Into this ADMS-DOTS we introduce DERIM, a Distributed Energy Resource Integration Middleware, an interface that allows integration of various DERs with the ability to communicate as dictated by the EUREICA framework (see Fig. 5.11). This integrated system uses the same software components, programmatic and user interfaces as the real-time ADMS, and creates an effective training and testing environment to operate with the actual network model, data, and functions in a controlled and safe environment. Fig. 5.12 shows an example of what the validation process looks like for attack 1a.

5.13 Results

5.13.1 Numerical simulation setup for markets

All use cases considered are based on an IEEE 123 test feeder (see **??**), which is radial, unbalanced, and multi-phase. The feeder was modeled in the GridLAB-D environment (see Section 5.12.1 for more details) and augmented to have a high penetration of DERs. For all attacks except attack 3, we assumed that the switch settings are in their nominal positions, such that we have one primary feeder, with 85 active nodes with SMOs/PMAs (out of the 123 in total). A PMO was assumed to be at the slack bus (substation), at either 115 or 69kV, with the SMOs at 4.16kV, and each SMA at 120-240V. The flexibility bids for the SMAs and

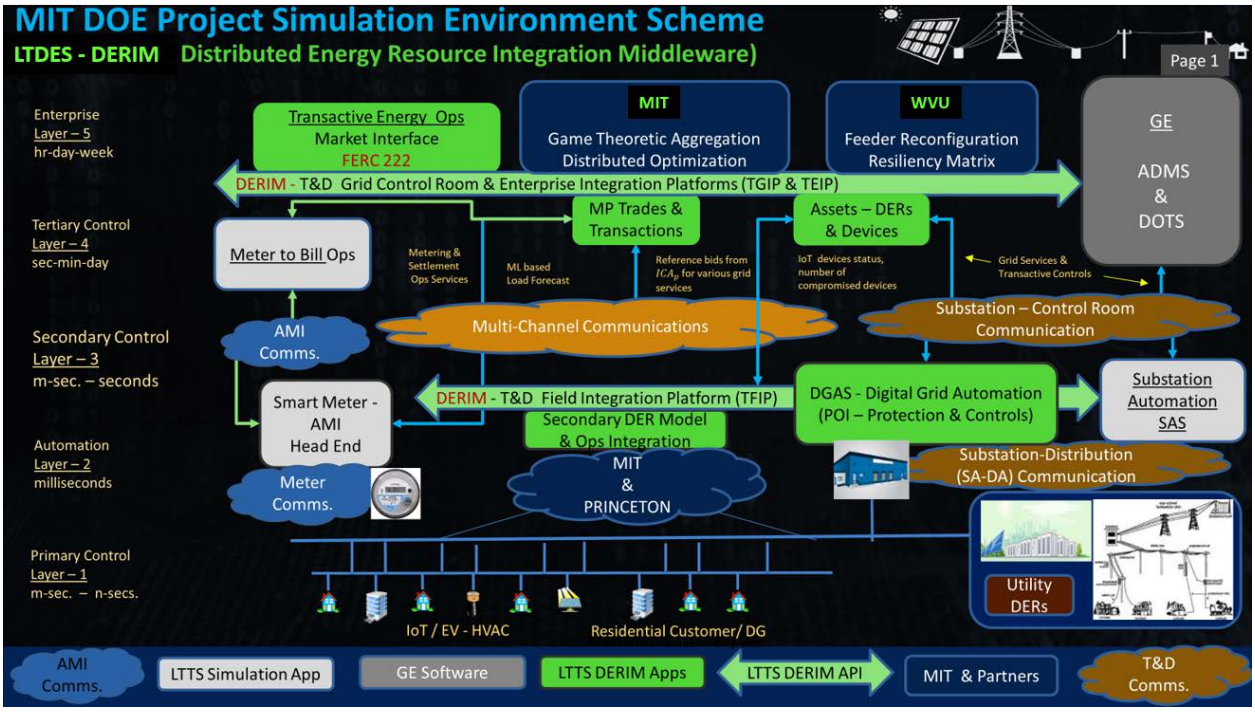

Figure 5.11: DERIM interface with ADMS-DOTS in the LTDES validation platform.

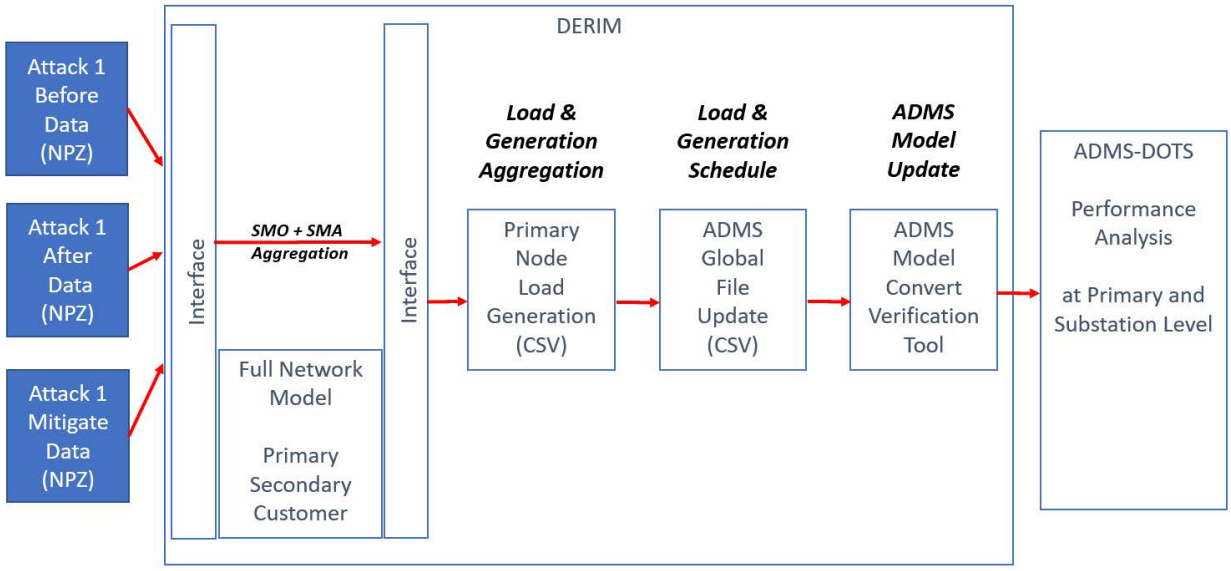

Figure 5.12: Attack 1a validation process workflow in the LTDES validation platform with DERIM and ADMS-DOTS.

the SMOs were randomly generated, allowing each to offer flexibilities of up to ±30% around their baseline power injections [142]. We used 5-minute real-time market LMPs from the California ISO and assumed the Q-LMP to be 10% of the P-LMP. Note that for all attack scenarios except attack 3, we used the CI model to represent the feeder as is.

For attack 2, however, we considered a modified version of the feeder and deployed the BF model instead. Here, we modified the original IEEE-123 feeder to consider a case where we have a few large distributed generators (PV, batteries, diesel generators) concentrated at just five primary feeder (SMO) nodes numbered 25, 40, 67, 81, and 94. This is in contrast to the other attacks, where there was instead a larger number of smaller DERs distributed throughout the network. Another distinguishing factor of this scenario is that the originally unbalanced feeder was converted to an equivalent balanced 3-phase model by (i) assuming all switches to be at their normal positions, (ii) converting single-phase spot loads to 3-phase, (iii) assuming cables to be 3-phase transposed, (iv) converting configurations 1 thru 12 to symmetric matrices and (v) modeling shunt capacitors as 3-phase reactive power generators [135]. Each SMO was assumed to have between 3-5 SMAs, with the number chosen uniformly at random. Since the injection data in the original IEEE-123 model was only available up to the primary feeder node level, we artificially randomly disaggregated the injections at each SMO amongst its SMAs, which could be net loads or generators.

We then performed a co-simulation of both the PM and SM for all attack scenarios. We refer the reader to [73, 116] for the behavior of this market structure for a nominal scenario when there is no attack. In what follows, we only consider the three attack scenarios described above. We also note that our flexibility bids were synthetically created, so the resulting flexible ranges in our simulations may be quite large at times and not realistic in some cases. However, our proposed framework can be generally applied to cases where there is less DER flexibility as well.

With the numerical simulation setup described in Section 5.13.1, we present details of how each of these attacks is mitigated using the proposed EUREICA framework. Note that for all attacks except attack 2.0, we use the mitigation strategy described in Section 5.10.2. For attacks 2a and 2b, we use the algorithm in Section 5.10.1 instead. In addition to market simulations, we also validated our results using high-fidelity software at the Pacific Northwest National Lab (PNNL), LTDES, and the National Renewable Energy Lab (NREL). Technical details for each validation platform can be found in Section 5.12.

5.13.2 Mitigation of attack 1a

We note that in attack 1a, loads are compromised, leading to an increase in the power import from the bulk grid. It is also assumed that the communication from all SRM to the PRM is disrupted, while the communication from the PRM to the SRM remains intact. That is, the PRM loses observability but is still able to communicate the redispatch of the new coefficients to the SRM. We do not consider the case when such observability is not lost, a discussion of which is beyond the scope of this paper. With the redispatch, the PM-SM framework identifies all of the new trustable PMAs (through the SA computations described in Section 5.10), which will provide the injections needed to fully mitigate the attack, and the overall power balance is thus met at all points in the distribution grid.

The steps in mitigation are as follows: 10 SMO nodes are attacked, resulting in a total

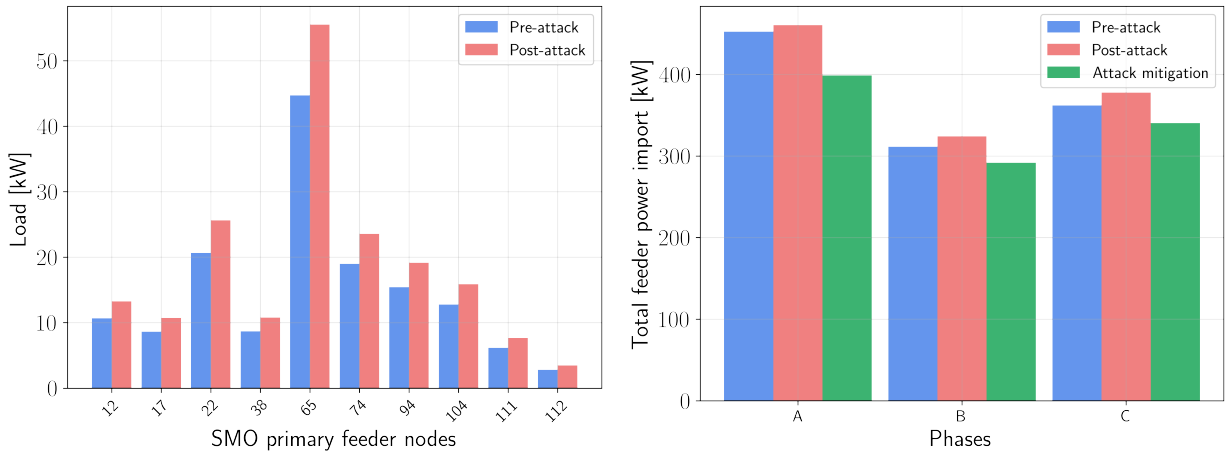

(a) Net load at attacked nodes.

(b) Feeder power import from main grid.

Figure 5.13: Effect of attack 1a and mitigation.

increase in load (generation shortfall) of 36 kW for the entire feeder as seen in Fig. 5.13a. A large number of flexible load nodes across the entire feeder helps with mitigation by curtailment and shifting, as in Fig. 5.14. Flexible load curtailment at individual SMO nodes ranges from a minimum of 0.55 kW to a maximum of 7.8 kW reduction per primary feeder node, using a combination of resources like HVAC, WHs, batteries, and EVs to reduce the net load. There is a 123 kW decrease in power import after mitigation as seen in Fig. 5.13b. The new SMO setpoints from the PM redispatch are then disaggregated amongst their SMAs during the following SM redispatch, with an example for SMO 77 shown in Fig. 5.15.

Attack 1a validation by LTDES

The outputs from the PM-SM market framework were sent to the DERIM interface, using which the effect on the total net load at the substation feeder head could be determined with the DERIM-ADMS-DOTS software platform (see Fig. 5.12 for an overview of the validation process). It is clear from Fig. 5.16 that without the intervention of EUREICA, the impact of the attack is a 37 kW jump in the feeder demand; in contrast with EUREICA, the feeder demand is cut by 94 kW. Moving further ahead from the attack timestep, the feeder net load eventually approaches back to the same value as if there hadn't been an attack.

While Fig. 5.16 zooms in on the period around the attack timestep, Fig. 5.17 shows the total feeder head load over the entire 24 h simulation horizon. We can clearly see the blip at 13:00 PST indicating the impact of the attack. Fig. 5.18 shows the effects of attack and mitigation on the net load at all the SMO primary nodes. This shows that the DERIM-ADMS-DOTS validation produces similar results to our market simulation in Fig. 5.13a and Fig. 5.14. The attack increases the load at the following nodes: 12, 17, 21, 36, 65, 75, 95, 105, 112, and 113. The majority of load curtailment for mitigation is contributed by the larger loads at nodes 1, 16, 48, 76, and 88.

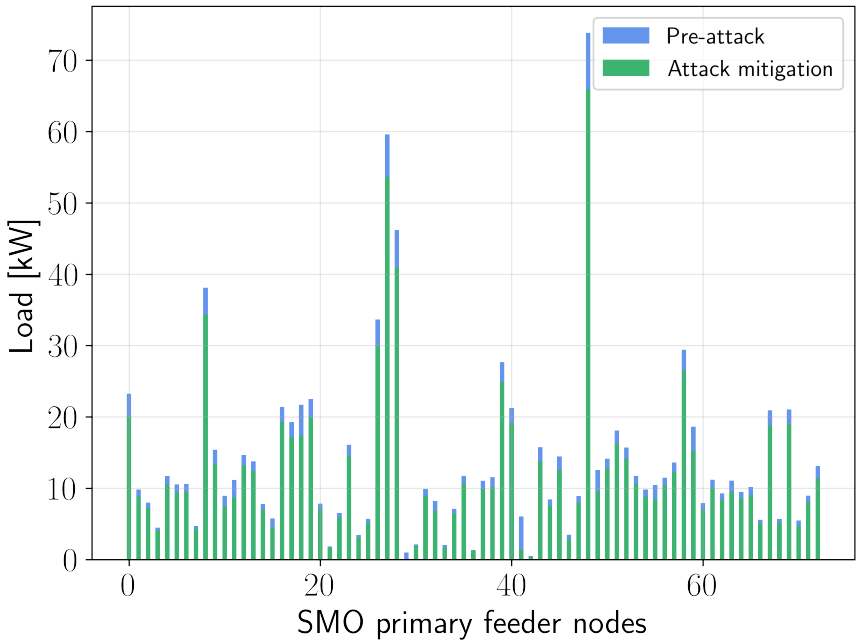

Figure 5.14: Curtailment of flexible loads for attack 1a mitigation.

Figure 5.15: Dis-aggregation of setpoint changes (from the PM) for SMO at node 77 across its 3 SMAs (in the SM) on phase B, after attack 1a mitigation.

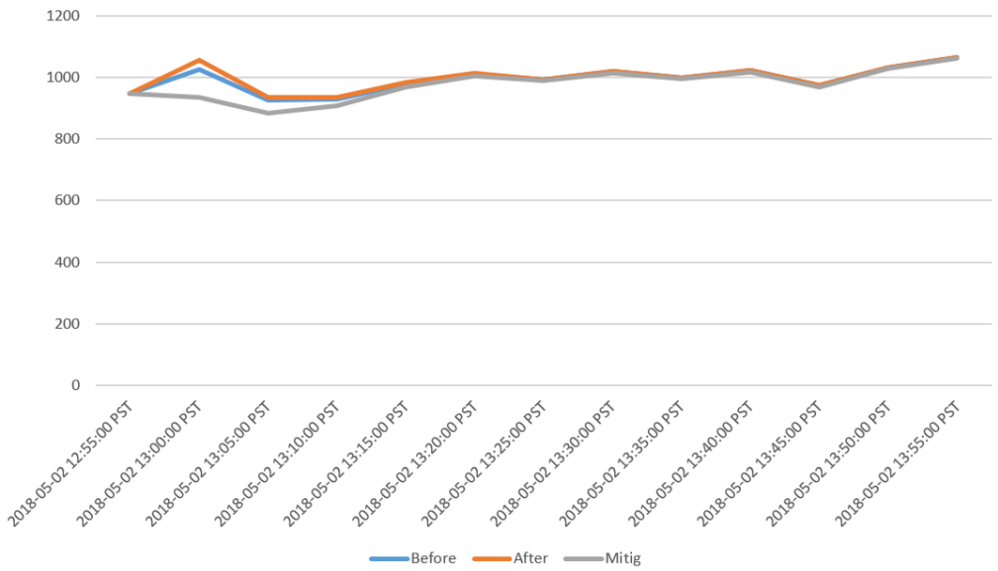

Figure 5.16: LTDES validation of attack 1a in the DERIM-ADMS platform, showing total power import at the substation around the attack time at 13:00.

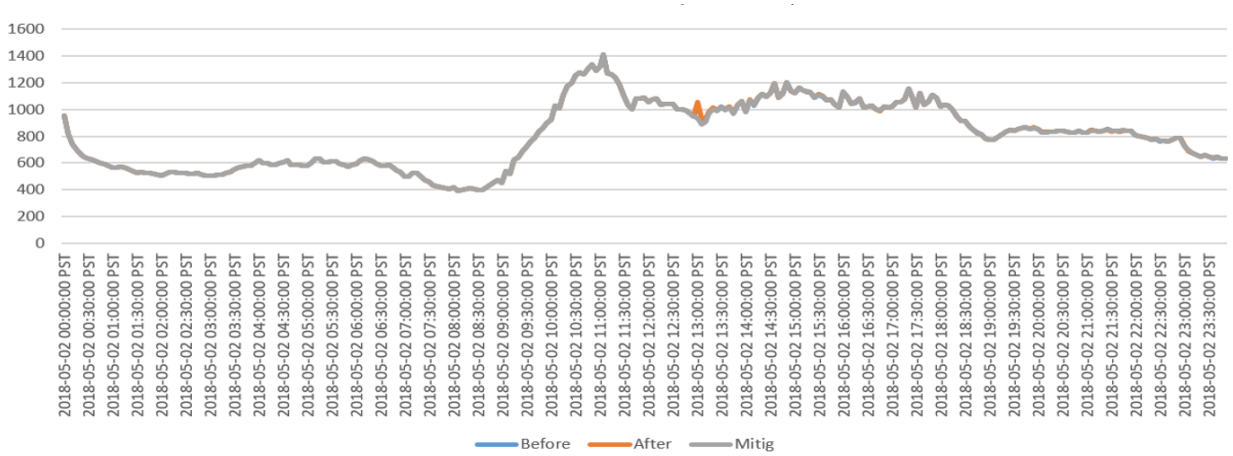

Figure 5.17: Effects of attack 1a on total load at feeder head over 24 h.

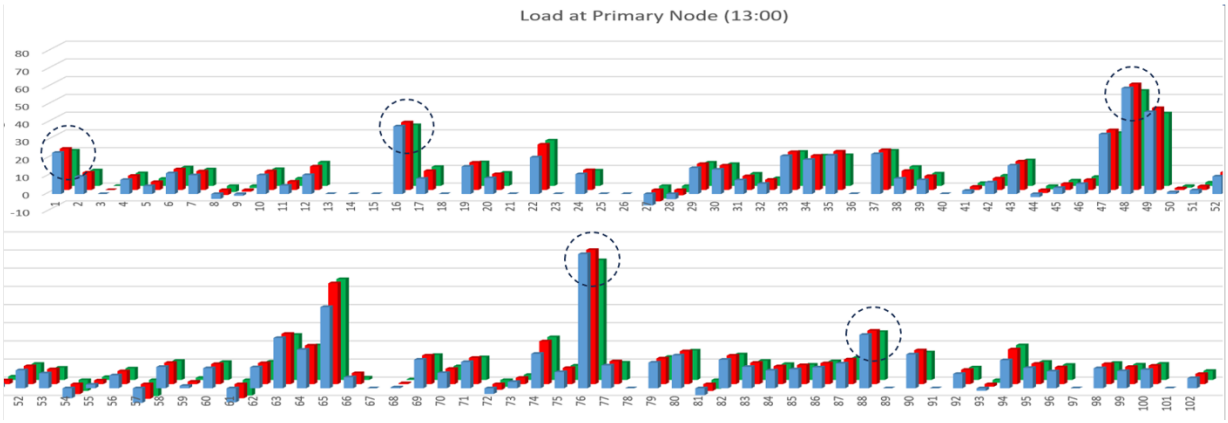

Figure 5.18: Load change at primary nodes during attack 1a. The values (i) without attack, (ii) with attack, and (iii) with attack mitigation are shown in the blue, red, and green bars, respectively. The SMO nodes providing the most flexibility are circled.

Validation of Attack 1a by PNNL using HELICS

This attack artificially increases the load at several devices throughout the network. Fig. 5.19 shows the effects of attack 1a and mitigation on the total feeder load over the course of the 48-hour simulation. This was performed using the HELICS platform and a GridLAB-D model. Firstly, we notice that the application of the LEM during day 2 generally results in curtailment of net load by leveraging DER flexibility, relative to day 1 (when the market is not used). Secondly, upon zooming in on the attack period (around 13:00 PST), we see that the LA attack increases the total system load. However, attack mitigation is quickly able to reduce the system load using flexibility and help the system recover.

Validation of Attack 1a by NREL using ARIES

The market structure is implemented using the same validation platform, with the primary feeders modeled on the RTDS, and secondary feeders, and below on Typhoon HIL and Raspberry Pis. In the implementation, the SMAs receive DER predictions from federated learning. The PMO and SMOs solve for primary and secondary market setpoints at each primary feeder node and secondary feeder, respectively, and then they are distributed to the SMAs and ultimately to the IoT devices. Under nominal conditions (without an attack), the market operates with the objective of voltage regulation and minimization of power import from the main grid.

In the case of attack 1a, the secondary feeder load increases by 63 kW, which may be driven by various factors, such as weather-related load swings or a coordinated cyber attack across IoT devices, such as the MadIoT attack. In this case, the mitigation is provided by using 30 flexible load nodes. Curtailment at the IoT device level ranges from a minimum of 0.2 kW reduction and a maximum of 0.5 kW reduction per primary feeder node. In total, approximately 130 kW of power imported from the main grid decreases after mitigation. Market clearing happens every minute, and the drop in the load is shown in Fig. 5.20. The IoT device response, which is the thermostat in this case, has an instantaneous response,

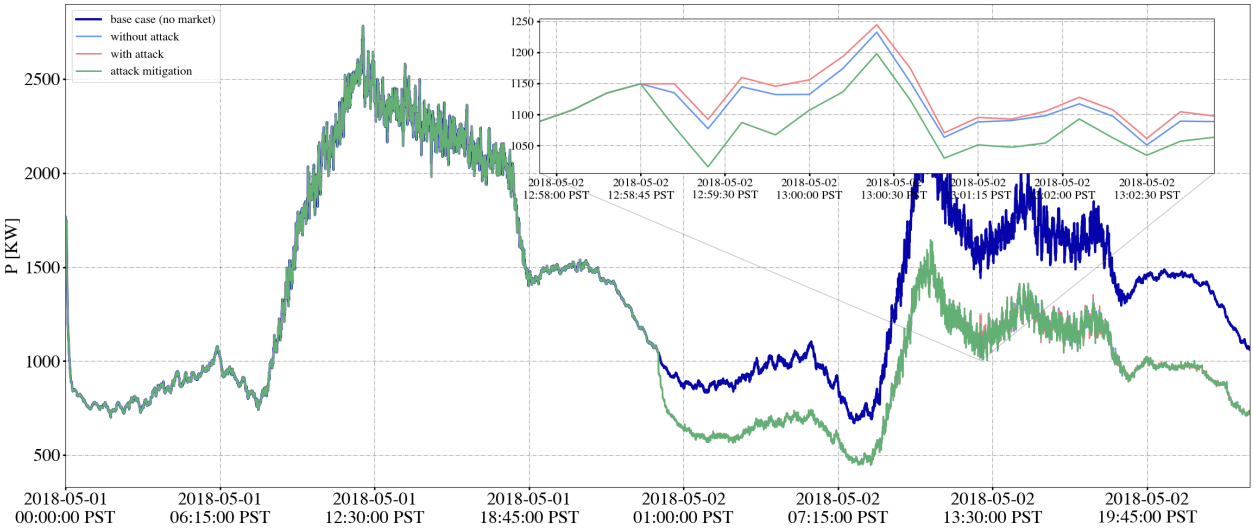

Figure 5.19: Validation of attack 1a mitigation effects of the EUREICA framework using HELICS, showing system load over 48 hours.

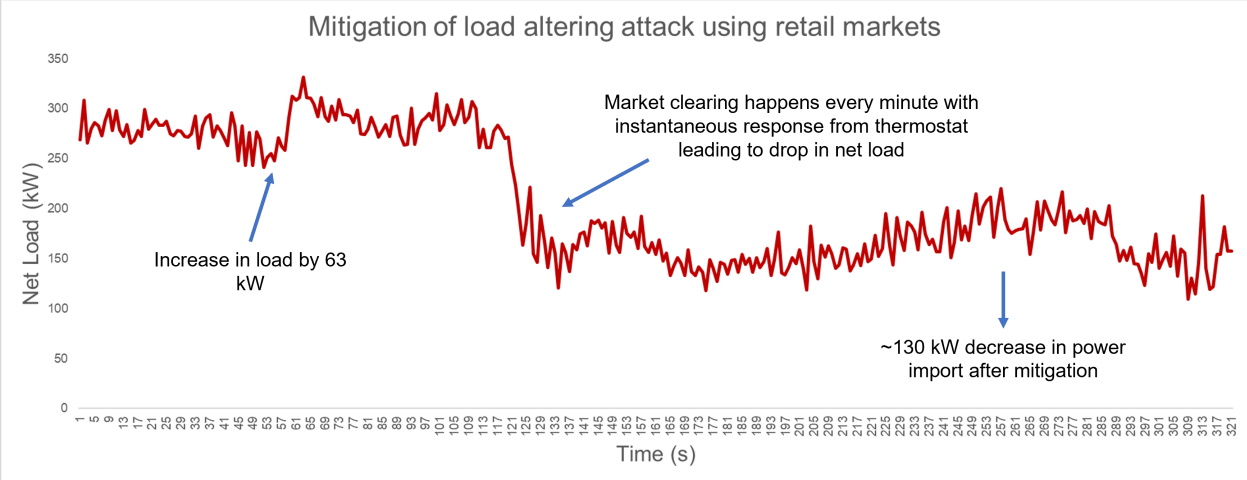

Figure 5.20: Implementation of market services to mitigate load increase in attack 1a.

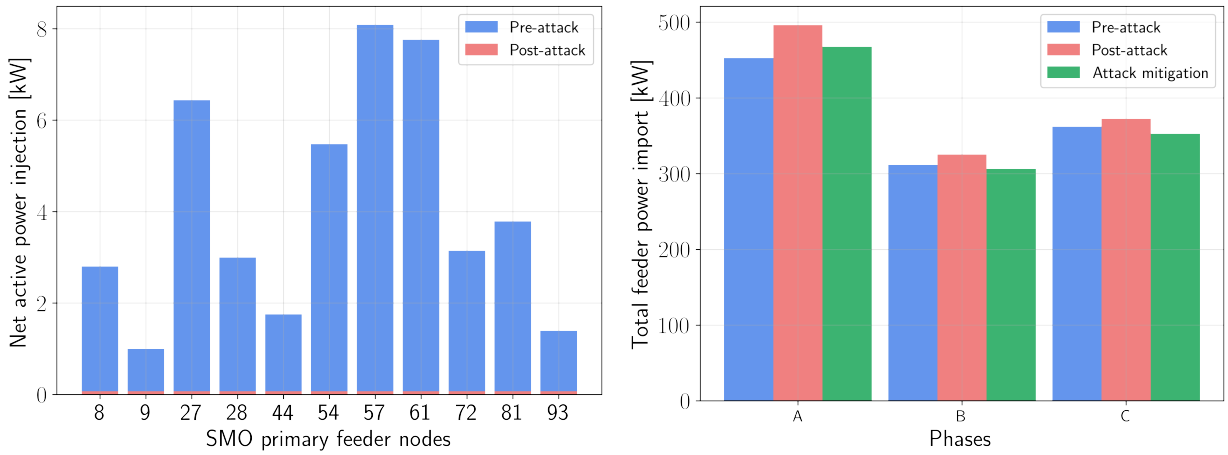

(a) Generation with and without attack 1b.

(b) Feeder power import.

Figure 5.21: Effects of attack 1b on SMO net generation and power import.

with an immediate drop in net load.

5.13.3 Attack 1b mitigation based on resilience

We note that in Attack 1b, there is a loss in net generation, and therefore, the power imported from the bulk grid increases. It is also assumed that the communication from all SRM to the PRM is disrupted, while the communication from PRM to SRM remains intact. That is, the PRM loses observability but is still able to communicate the redispatch of the new coefficients to the SRM. We do not consider the case when such observability is not lost, a discussion of which is beyond the scope of this paper. With the redispatch, the PM-SM framework identifies all of the new trustable PMAs through the SA computations described in Section 3, with the overall power balance met at all points in the distribution grid.

The steps in mitigating this attack are as follows. Due to the attack, 45 kW of net generation was compromised as shown in Fig. 5.21a. The PMO alerts other trustable PMAs/SMOs to redispatch their generation assets in the PM. Trustable PMAs/SMOs will curtail flexible loads to respond and mitigate the attack as in Fig. 5.23. The redispatch is also influenced by the resilience scores of different SMOs over time, shown in Fig. 5.24. SMOs redispatch the SM, which provides correct setpoints to all their SMAs. An example of this in Fig. 5.22 shows how the SMO at node 35 disaggregates its new setpoint amongst its 3 SMAs. As a result of mitigation, the total import from the main grid stays at the same level as shown in Fig. 5.21b.

We now highlight the effects of resilience scores on the mitigation of attack 1b, where a number of DGs are attacked. We consider how the RSs of SMOs and SMAs influence which resources are utilized to mitigate the attack. The RSs of the flexible SMOs are plotted against their absolute and relative levels of net load curtailment in Fig. 5.25a and Fig. 5.25b, respectively. We see that in relative terms, the curtailment is generally distributed evenly to ensure that no single SMO is disproportionately affected. However, if the PMO does need to utilize more flexibility from certain SMOs, it generally calls upon more reliable ones with

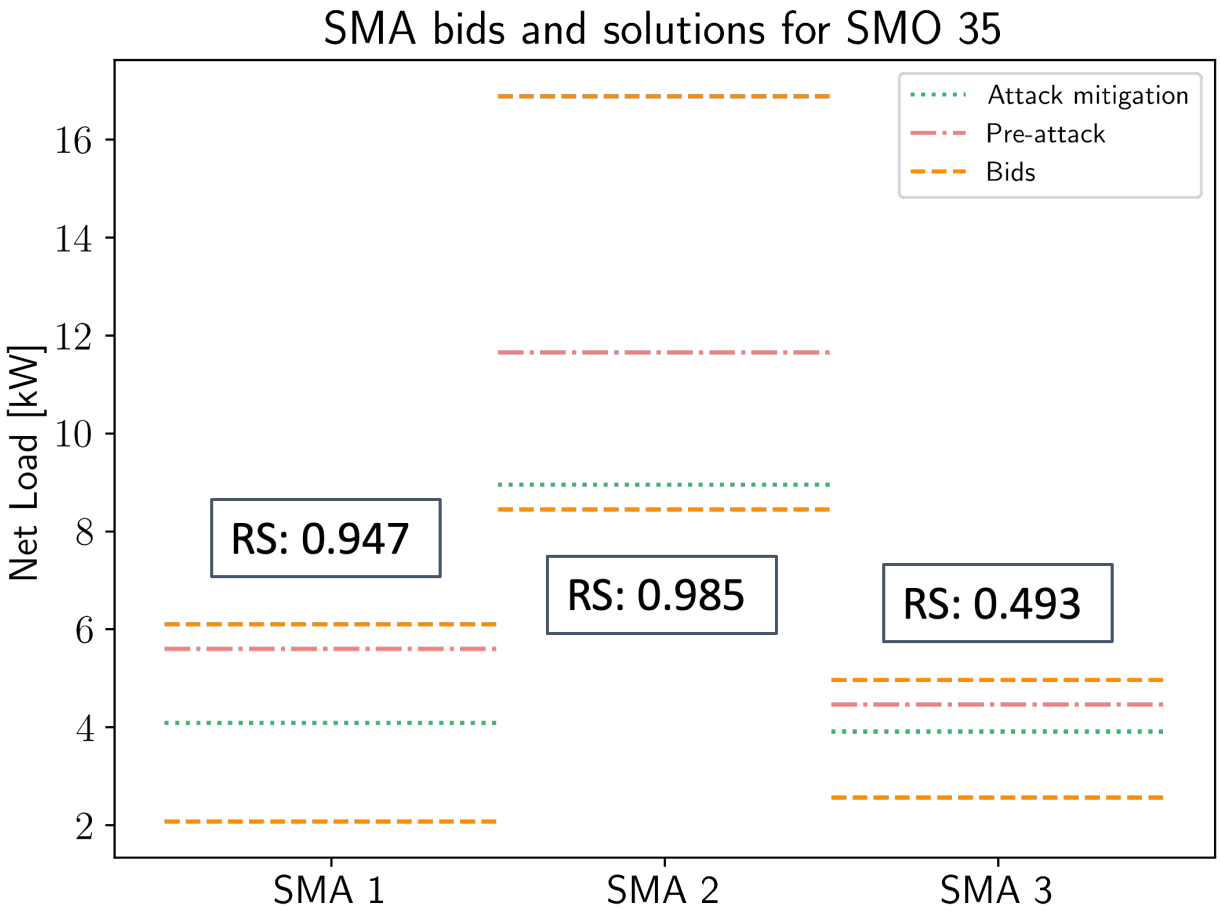

Figure 5.22: Dis-aggregation of changes in the setpoints for SMO (from the PM) at node 35 across its 3 SMAs (in the SM), resulting from attack 1b mitigation, along with each SMA's RS. All 3 SMAs are on phase A.

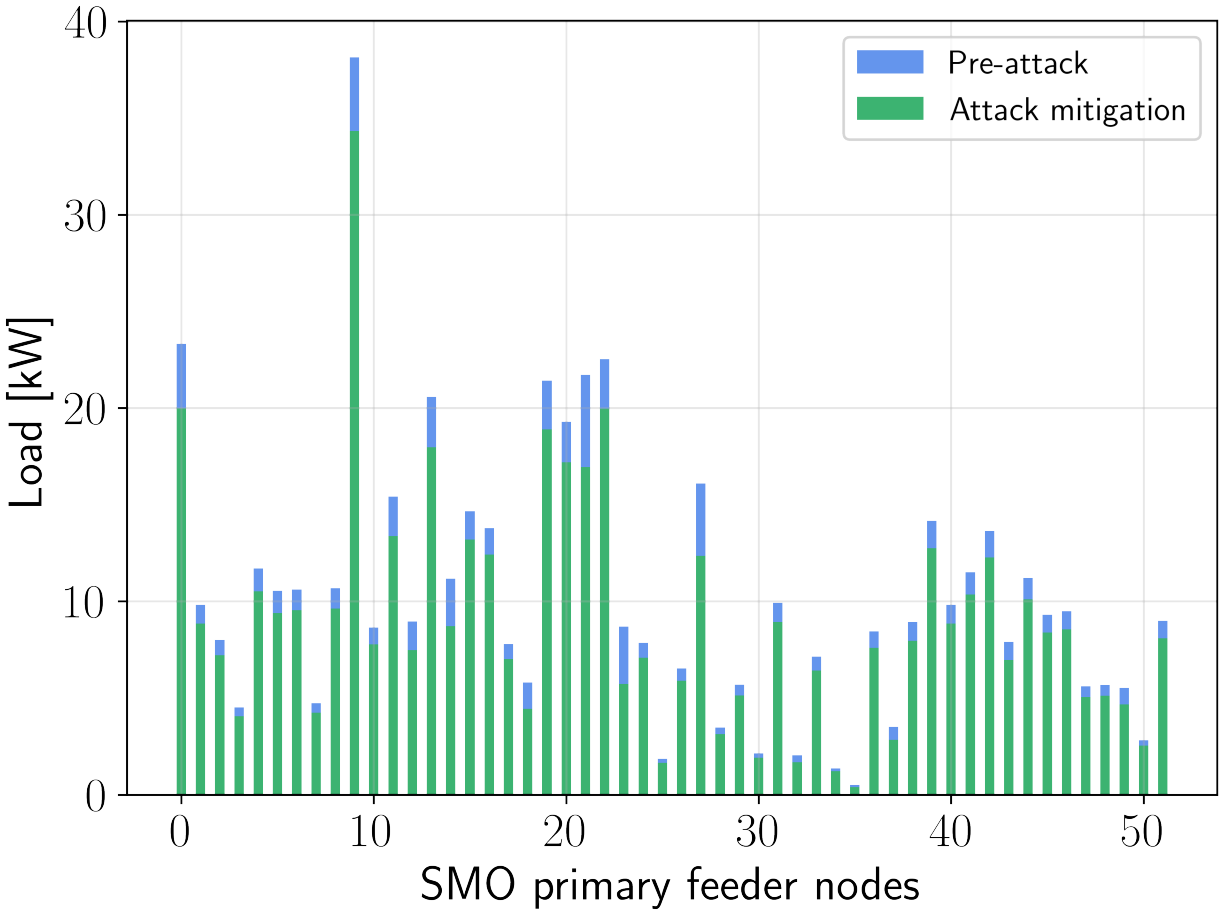

Figure 5.23: Curtailment of flexible loads for attack 1b mitigation.

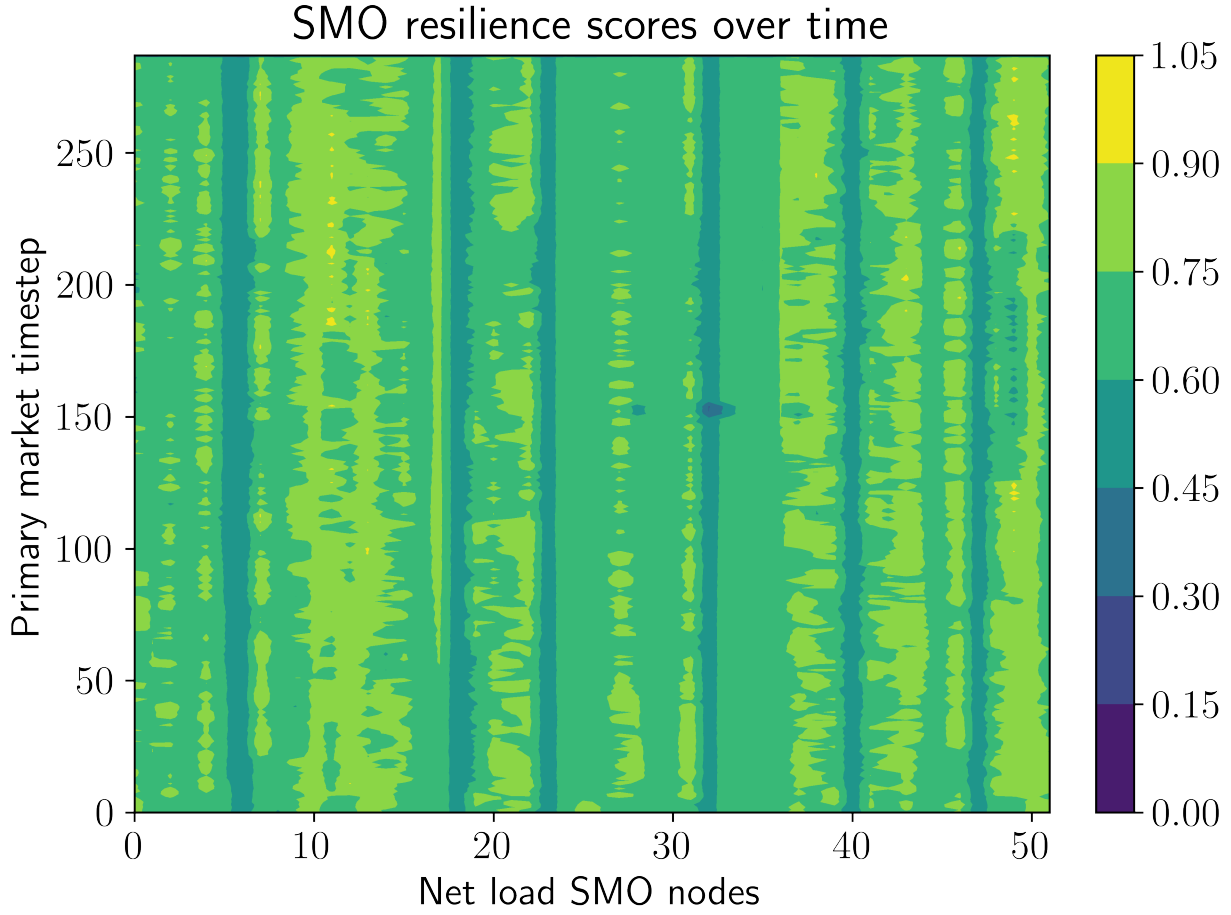

Figure 5.24: Locational-temporal trends of RS across all flexible SMO nodes and over the whole simulation period of 24h.

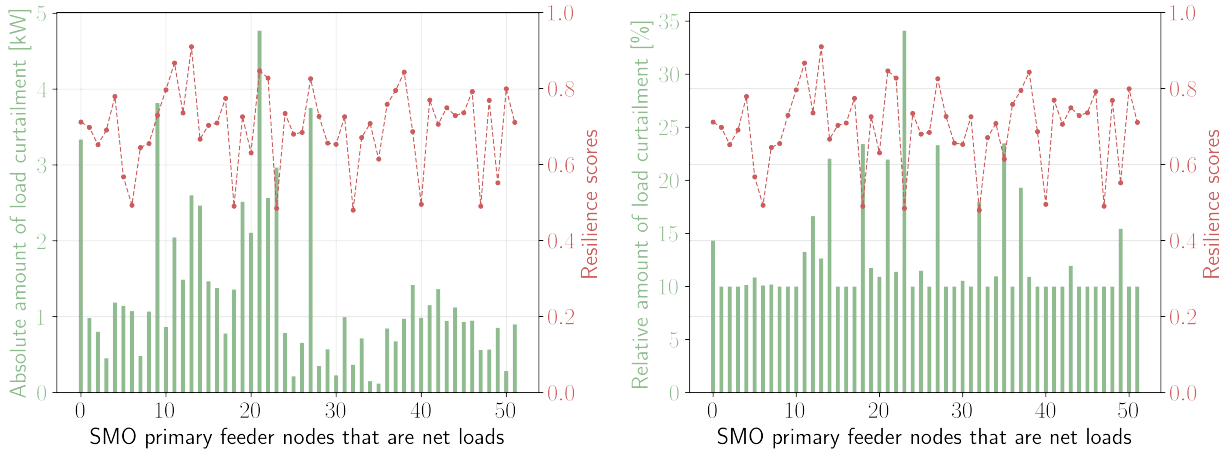

(a) Absolute curtailment and RS.

(b) Relative curtailment and RS.

Figure 5.25: Distribution of absolute and relative amounts of load curtailment across the flexible net load SMOs, along with their corresponding resilience scores.

higher RSs. The absolute amounts of curtailment vary for each SMO based on their baseline load. This also holds while disaggregating SMO setpoints at the SM level, where the SMO allocates greater flexibility to SMAs with higher RSs, as seen in Fig. 5.15.

Validation of Attack 1b by PNNL using HELICS

We utilized the HELICS-based co-simulation platform to simulate this use case in which several of the distributed generation resources are being disconnected, leading to about a 44 kW loss in generation. That is accomplished in the model simulation by taking offline the PVs at the buses as indicated in Fig. 5.21a. However, with the SA enabled through the market module, the SM agents are informed about how much they need to adjust their flexible assets, which results in an approximate 36 kW load and local generation alteration after attack mitigation to counterbalance the distributed generation loss, as seen in Fig. 5.23. The effect of the market integration on the total system load during the second 24-hour period of a 48-hour simulation is depicted in Fig. 5.26. In particular, the window details the attack that happens around 13:00 on the second day and how the total flexible load is manipulated to mitigate the need for increased generation demand from the main grid. We see that the attack mitigation reduces the impact of the attack by lowering the total feeder load and bringing it back down closer to the values if there wasn't an attack. However, we note that even after mitigation, the load is still slightly higher than the 'without attack' case for some periods, but much lower than the 'with attack' case.

Validation of Attack 1b by LTDES using DERIM and ADMS-DOTS

Fig. 5.27 shows the effects of the DG attack 1b on the total system load over the full 24 h simulation, while Fig. 5.28 zooms in on the period around the attack. We see that without the market-based mitigation, the feeder demand would have jumped by 68 kW due to the attack. However, with mitigation, the attack impact is minimal since there's only a 4 kW

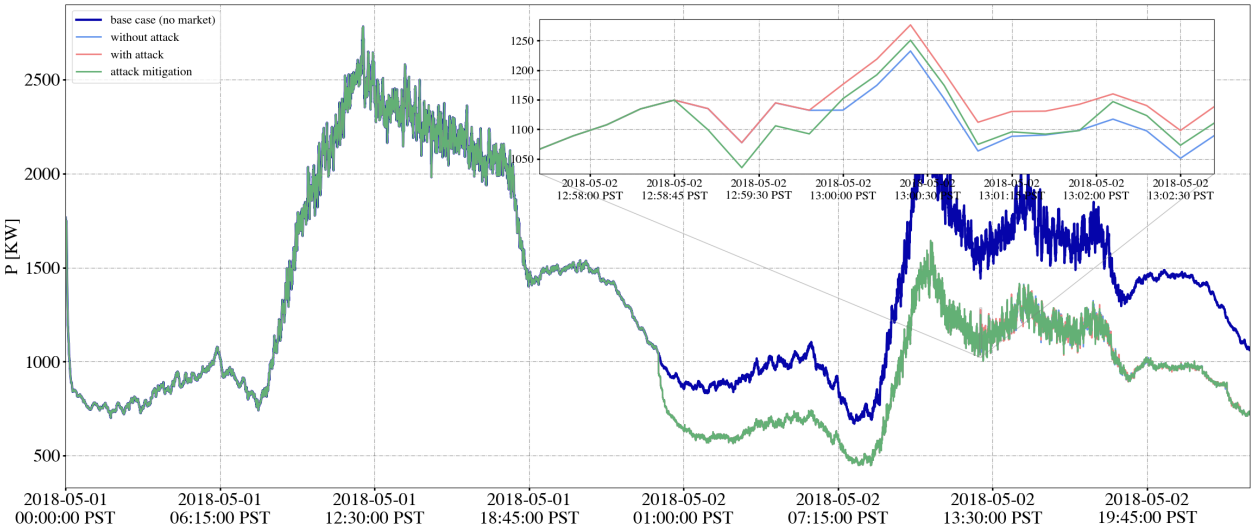

Figure 5.26: Validation of attack 1b mitigation effects of the EUREICA framework using HELICS, showing system load over 48 hours.

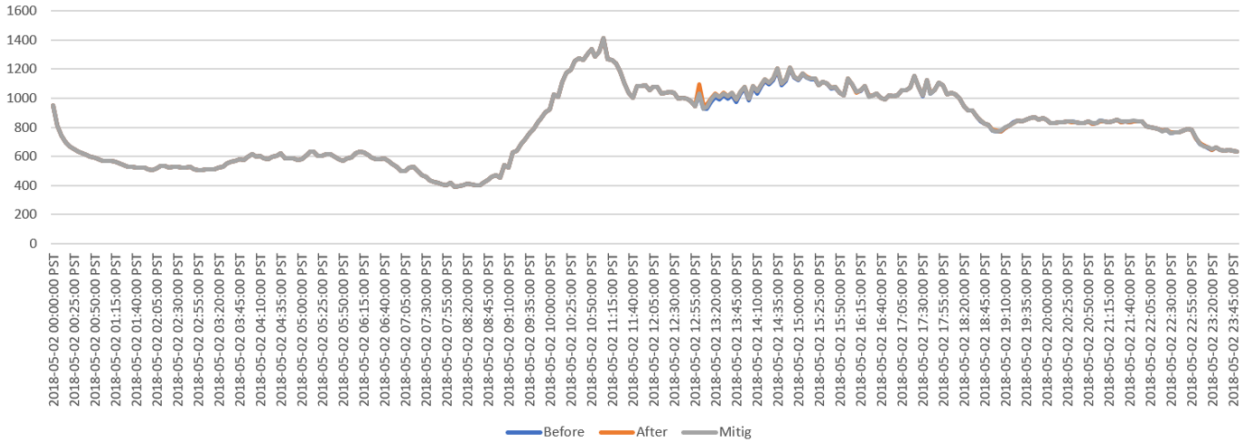

Figure 5.27: Effects of attack 1b on the total load at the feeder head over 24 h.

increase in feeder demand. Fig. 5.29 shows the changes in net injections at all primary nodes during attack 1b. This essentially shows that we leverage flexibility from several primary nodes across the feeder, producing results similar to those shown in Section 5.13.3. The attack causes the following DG nodes to lose power and go offline: 9, 28, 45, 55, 56, 58, 62, 73, 82, and 94. The following flexible load nodes contribute a majority of the curtailment needed to mitigate: 1, 48, 76, and 88.

Fig. 5.30 compares the load setpoints at the SMO level (updated every 5 minutes) for node 76 versus the aggregated setpoints over all the SMAs at this node (cleared every 1 minute). Although these are largely similar, there are some slight differences between the two values. Thus, it may make more sense to utilize the more precise SMA setpoints directly for the ADMS simulation.

We also did some further analysis on the role of the SM and PM in attack mitigation.

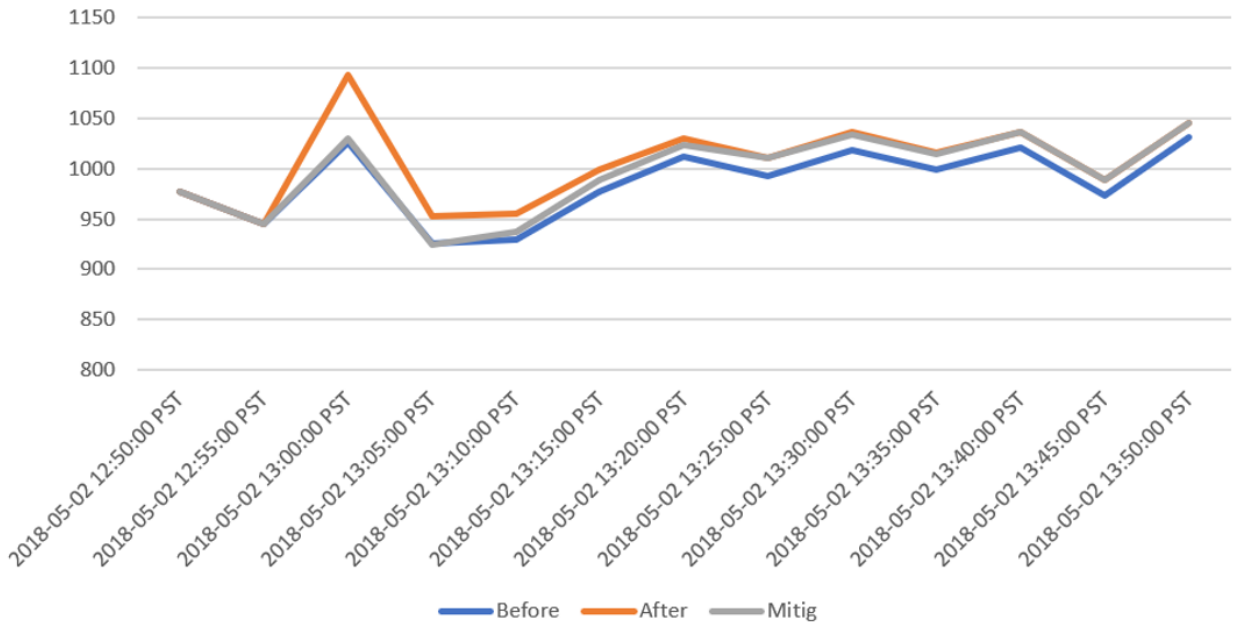

Figure 5.28: LTDES validation of attack 1b in the DERIM-ADMS platform, showing total power import at the substation around the attack time at 13:00.

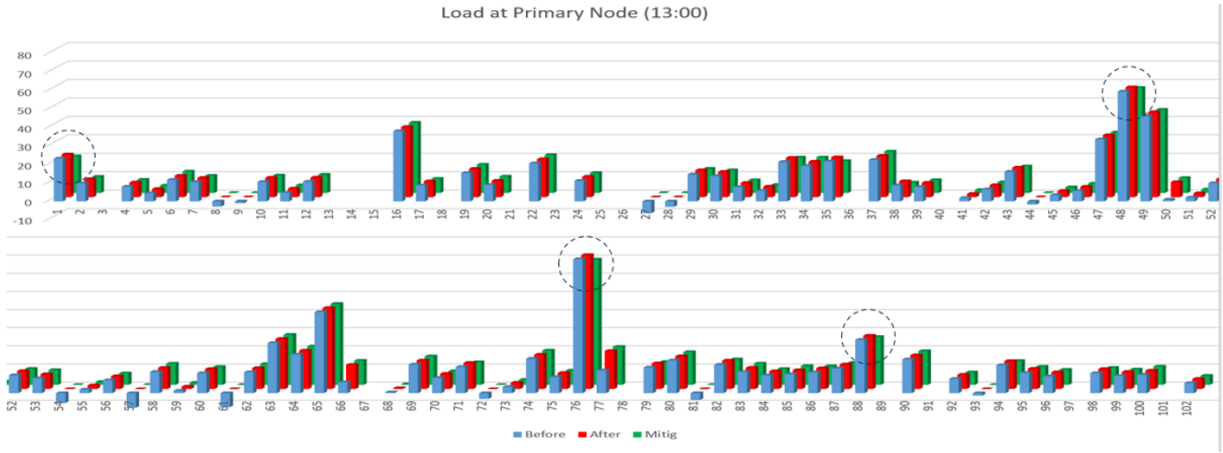

Figure 5.29: Load change at primary nodes during attack 1b. The values (i) without attack, (ii) with attack, and (iii) with attack mitigation are shown in the blue, red, and green bars, respectively. The SMO nodes providing the most flexibility are circled.

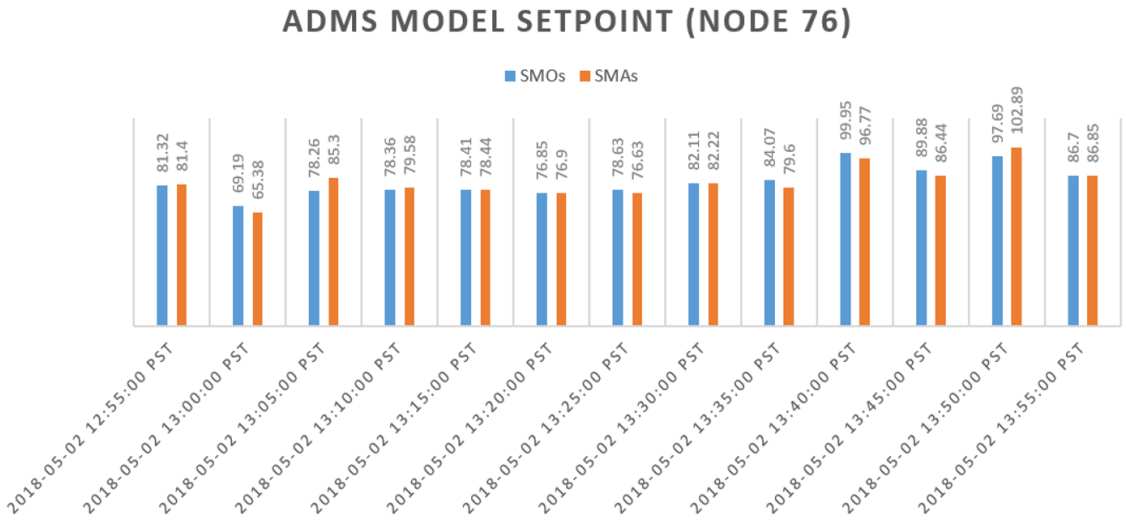

Figure 5.30: Forecasted values of SMO and SMA setpoints at primary node 76 during attack 1b mitigation.

Fig. 5.31 compares the contributions of the setpoints of the SMOs (5 minutes) and the SMAs (1 minute). The blue bar shows the 5-minute setpoint changes expected from the SMOs, while the orange bar shows the 1-minute setpoint changes at the SMA level. We see that the SM clearing every minute and the associated SMA setpoint changes contribute more toward the overall primary load adjustment when compared to the SMO-level changes alone.

5.13.4 Attack 1c mitigation

Attack 1c is a more distributed attack where individual SMAs representing secondary feeders are attacked directly. We considered a case where a large number of DGs, including solar PV and batteries, are attacked. A total of 53 SMA nodes with DGs were compromised and taken offline, resulting in a total loss in generation capacity of 157 kW. This leads to a decrease in the net injections across all the SMOs as seen in Fig. 5.32a - there are no longer any SMOs with net generation after the attack, and the loss of local generation also increases the net load at the SMOs. This leads to an increase in power import from the main transmission grid, as in Fig. 5.32b.

In the case of all other attacks, the mitigation strategy involves the PM redispatch occurring 1st, followed by the SM redispatch. There, only the PM is directly involved in attack mitigation, while the SM is only used to disaggregate the new SMO setpoints amongst their SMAs. However, in the case of attack 1c, the SM redispatch occurs first at the secondary feeder level and is then followed by the PM redispatch at the primary feeder level. Thus, both the SM and PM are actively involved in attack mitigation here. We see in Fig. 5.32b that we can partially mitigate the attack by leveraging the flexibility of SMAs in the SM. However, SM mitigation alone is not sufficient. We also need to utilize the inter-SMO flexibility in the PM to fully mitigate and restore the feeder import back down to the pre-attack level. A summary of the attack metrics is shown in Table 5.4.

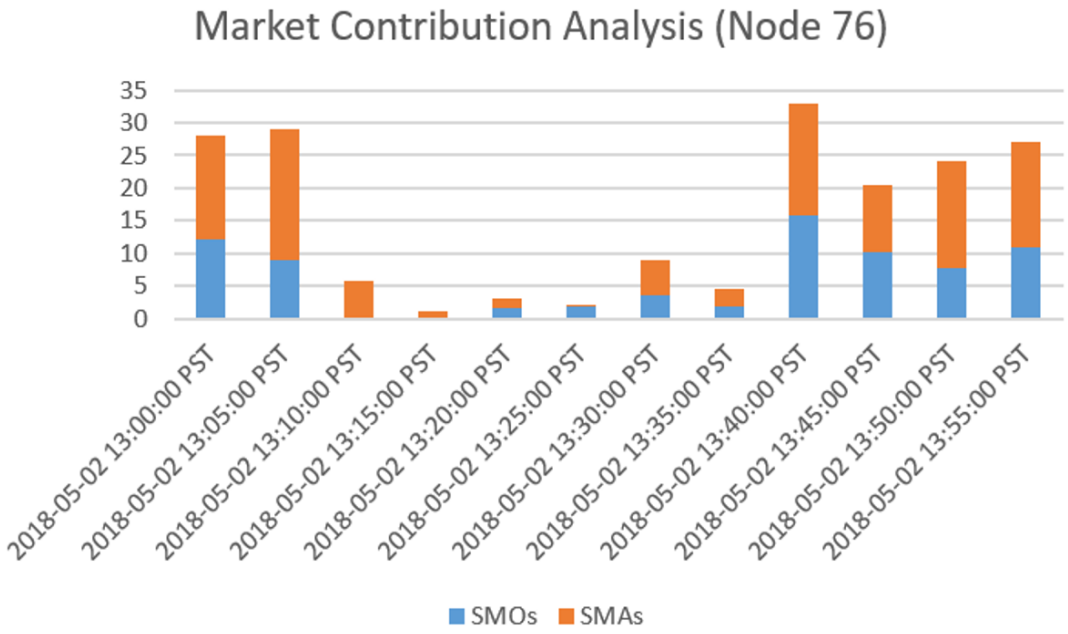

Figure 5.31: Comparison of forecasted changes in SMO and SMA setpoints due to attack 1b mitigation.

Table 5.4: Attack 1c summary.

	Power import [kW]	**Total net load [kW]**
Pre-attack	1412	1457
Post-attack	1722	1716
SM mitigation only	1553	1547
SM + PM mitigation	1422	1417

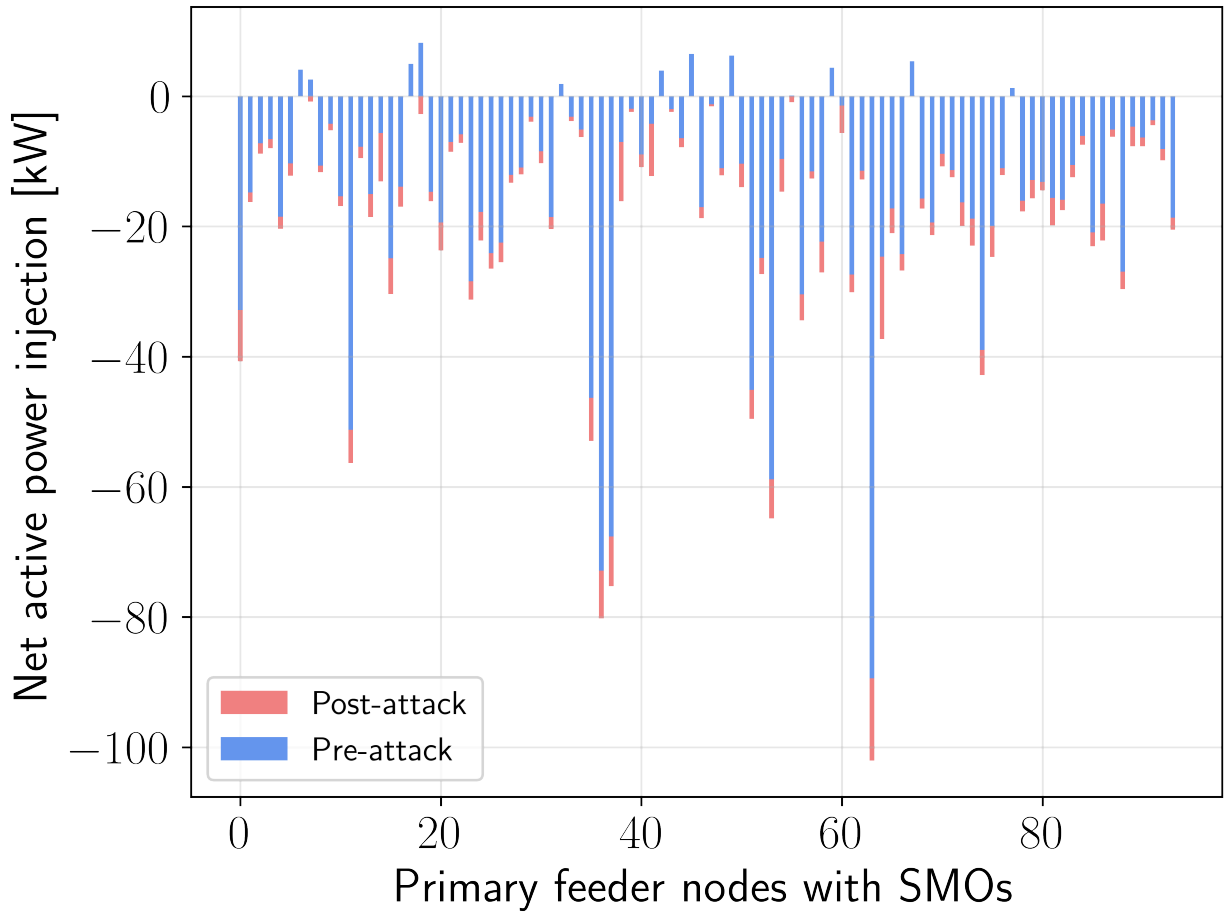

(a) SMO injections.

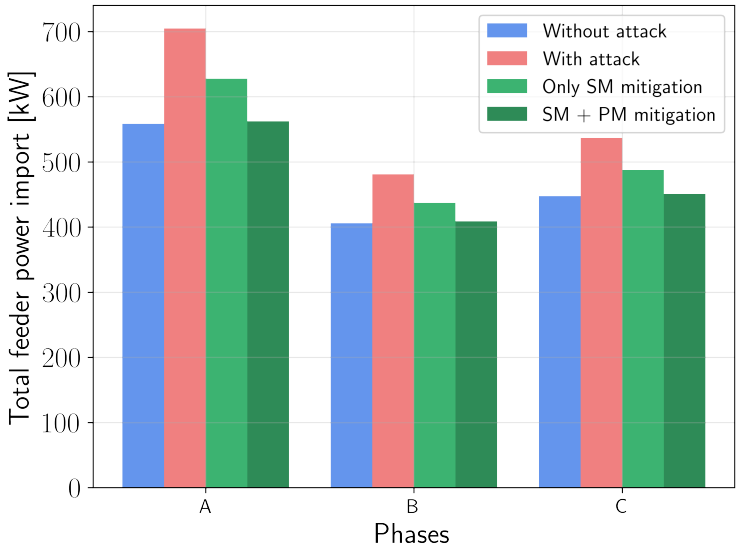

(b) 3-phase power imports from the main grid.

Figure 5.32: Effects of attack 1c on SMO injections and power import.

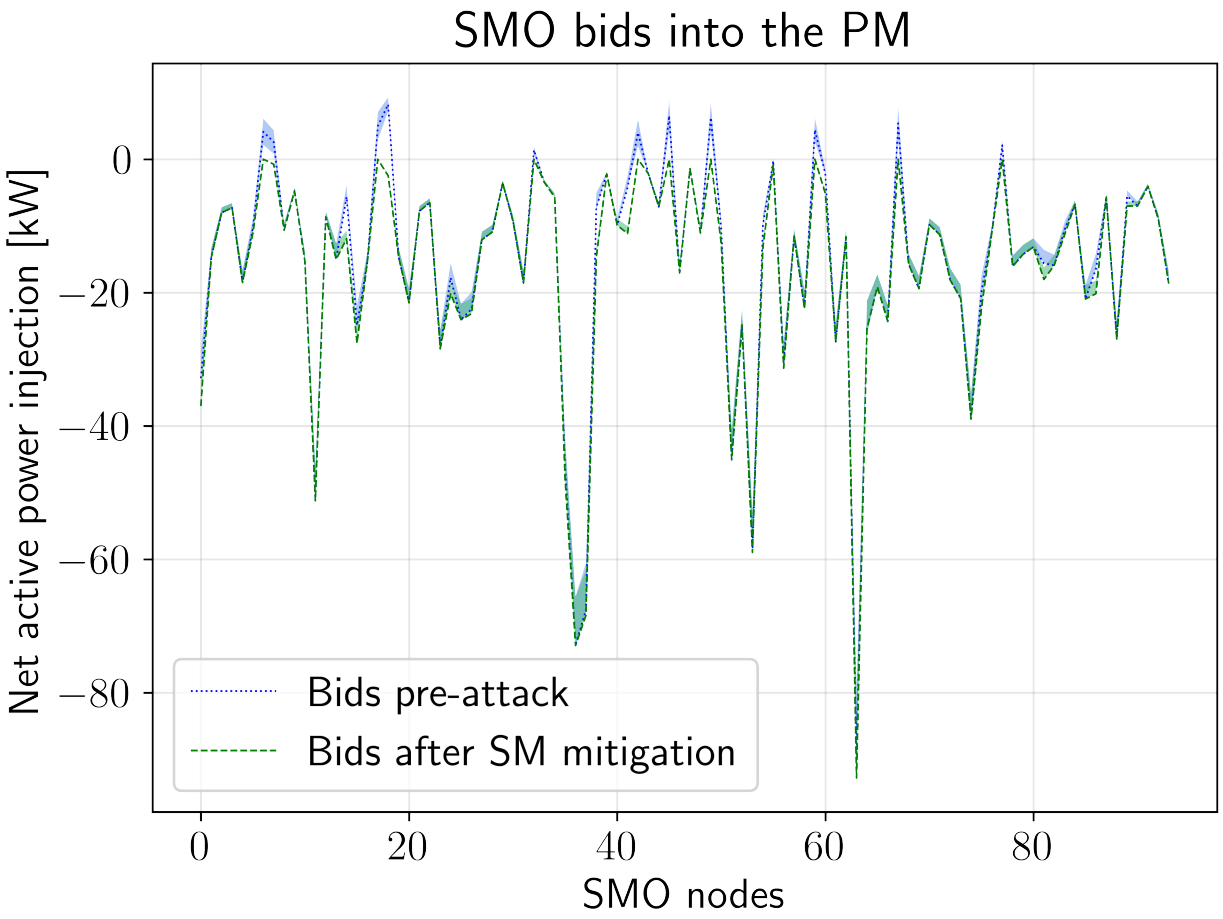

Figure 5.33: Comparison of SMO flexibility bids into the PM before and after the attack. The dashed and dotted lines indicate the baseline values, while the shaded regions are the flexibility bids around the baseline.

We also compare the flexibility bids of the SMOs before and after the attack in Fig. 5.33. As expected, the net load of the bids generally increases across all SMOs due to the loss of local DGs at their respective SMAs. However, by leveraging their SMA flexibilities, the SMOs are still able to offer some flexibility to the PM to help mitigate the attack.

Contributions of SM and PM to attack 1c mitigation

Attack 1c is a more distributed attack where individual SMAs are attacked directly. Here, we show how both the SM and PM flexibility are needed to fully mitigate the attack. Fig. 5.34 shows the contributions of the SM and PM toward attack mitigation. We see that for most of the SMO nodes, both the SM and PM flexibility play a significant role in reducing the net load compared to the post-attack case. At the SM level, we utilize the available upward flexibility of any SMAs with remaining online DGs and use the downward flexibility of all net load SMAs. At the PM level, we utilize the downward load flexibility of the SMOs (which are all net loads after the attack).

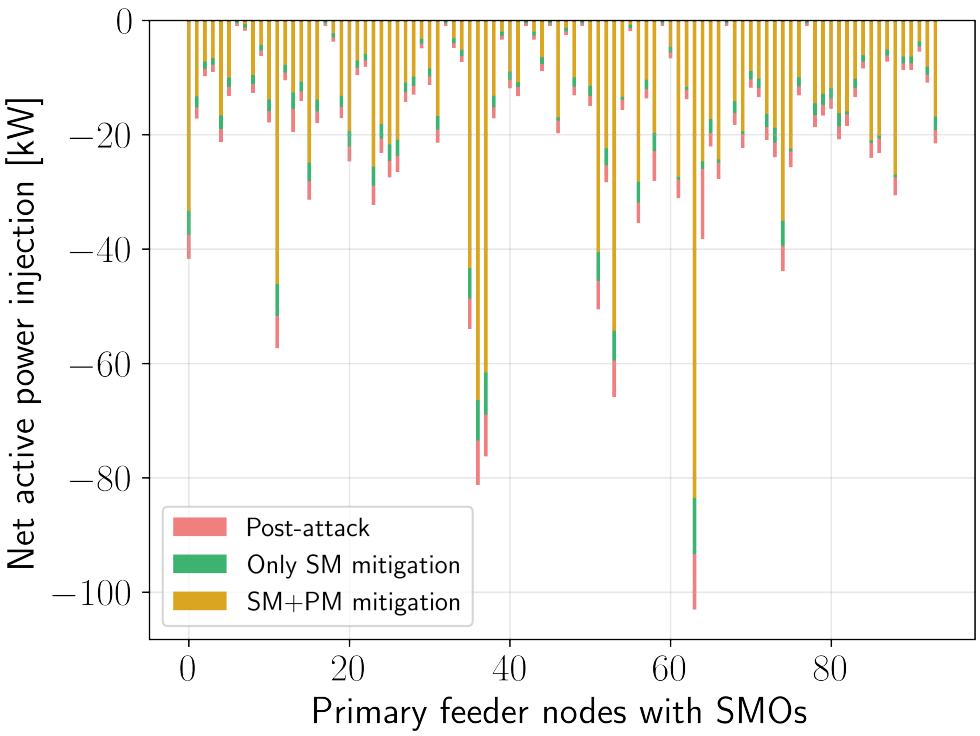

Figure 5.34: Contributions of SM and PM flexibility for attack 1c mitigation.

5.13.5 Mitigation of attack 2

Here, we describe the mitigation of two attacks at the primary feeder level that are relatively broader in scope, one is a medium-scale attack, and the second is a large-scale attack. Both are disruption attacks where the attacker shuts down one or more of the large DGs in the network. We only consider a single primary market time step to study the effects of an instantaneous attack. Mitigation can use P dispatch from batteries, P and Q curtailment from flexible loads, limited P dispatch from PV, Q support from smart inverters (connected to PV and batteries), as well as conventional dispatchable fossil fuel sources like diesel generators.

Mitigation of Attack 2a

This corresponds to a case where there are five large distributed generators in the modified IEEE-123 system, one of which (at SMO node 94) is taken offline. Here, we see that the remaining four SMO nodes (25, 40, 67, 81) have more than enough remaining generation capacity to meet the shortfall caused by the attack. Without mitigation, the attack would have resulted in an additional import of about 261 kW from the main grid. However, by utilizing the upward flexibility of remaining SMOs, we are able to fully resolve the attack and bring the total power imported back to pre-attack levels. The left figure in Fig. 5.35 shows the results of the PM dispatch before the attack and after attack mitigation for the five key SMO nodes of interest. The plot also shows the SMO's bids into the PM, with the dashed blue line being the baseline injection bid and the blue-shaded region representing the upward/downward flexibility around it. The right figure shows the results of the SM

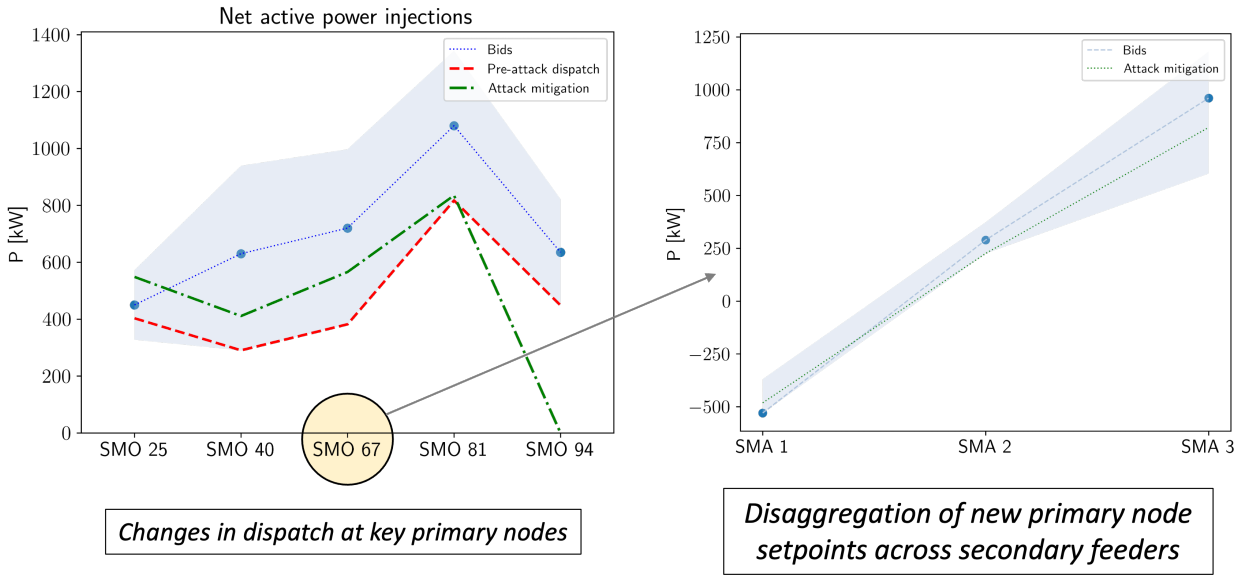

Figure 5.35: Mitigation of small-scale attack 2a.

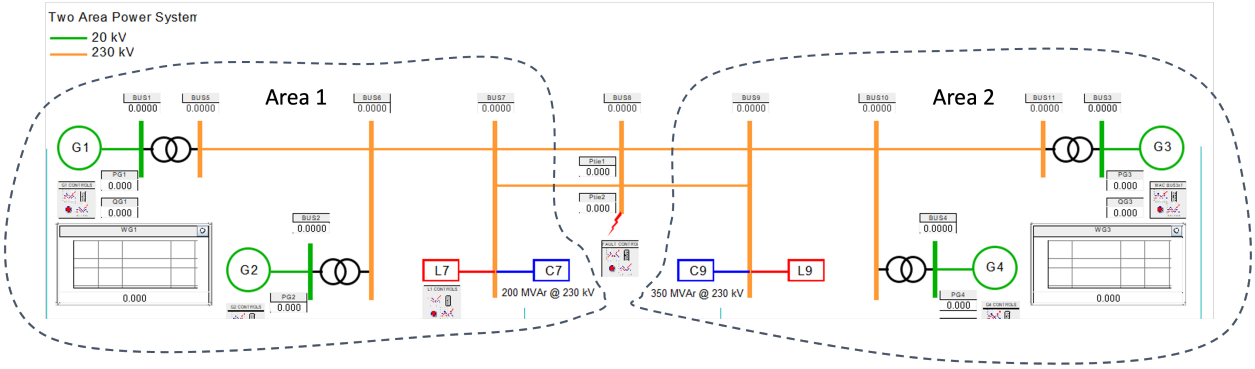

Figure 5.36: Schematic of Kundur 2-area power system

re-dispatch after the attack mitigation and PM re-dispatch for SMO 67 as an example. It disaggregates the new setpoints among its three SMAs, with SMA 1 being a net load while SMAs 2 and 3 are net generators.

Large-scale attack 2b

Here, we adopted a top-down approach in emulating an attack and started with a Kundur 2-area transmission model, with the attack occurring in Area 2. Fig. 5.36 is a diagram of the Kundur 2-area transmission system commonly used as a test case to study dynamic stability, power interchange, oscillation damping, etc. The system contains 11 buses, four generators, and two areas. The two areas are connected with weak tie lines [138].

Noting that Area 2 (which consists of a load of 1767 M) can be broken down into 552 IEEE-123 feeders, each with approximately 3.2 MW, we assume that an attack that compromises about 650 kW of generation occurs in each of these 552 feeders. This, in turn, corresponds to an overall shortfall of 359 MW at the transmission level. We introduced this 650 kW shortfall

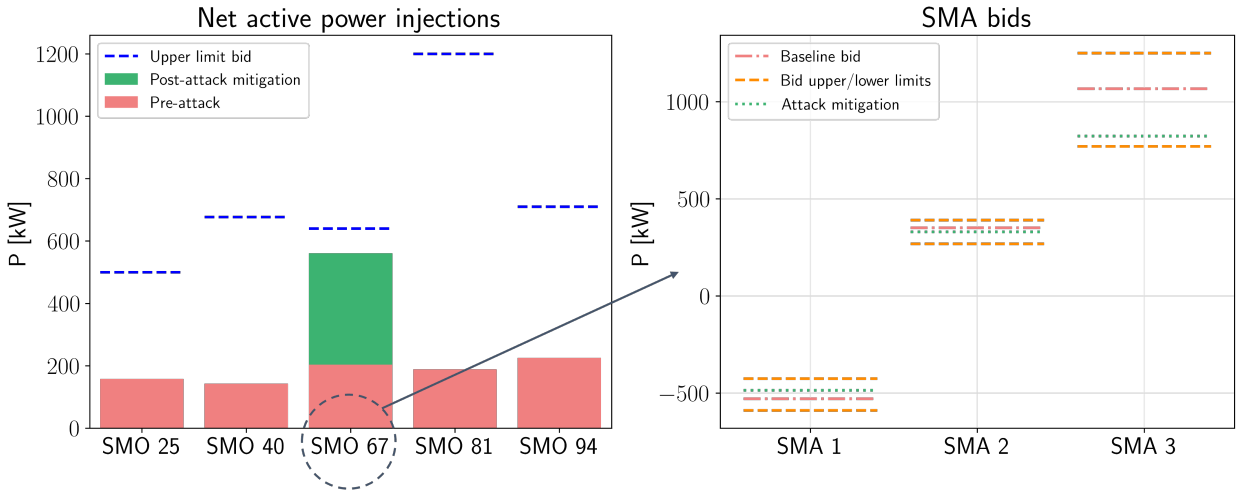

Figure 5.37: Mitigation of large-scale attack 2b.

in the form of a generation loss at four nodes, 25, 40, 81, and 94, in each of the 552 primary feeders. The only remaining SMO with significant generation capability is at node 67. With the same procedure as outlined in the previous scenarios, the use of our proposed EUREICA framework leads to the results in Fig. 5.37. In order to mitigate the attack, we need to leverage the upward generation flexibility of the remaining SMO 67 to increase its output injection after attack mitigation, while the net injections for all the other four attacked SMOs drop to zero, as seen in the left plot. The right plot shows the new SMA schedules resulting from the revised SM clearing.

However, due to the larger scale of the attack, re-dispatching the generator SMOs is no longer sufficient to fully meet the shortfall. Furthermore, as seen in Fig. 5.37, we are not able to utilize all the upward flexibility of the remaining online SMO 67 since its dispatch is limited by power flow constraints, on nodal voltages and line currents in particular. Thus, we also need to perform some shifting and curtailment of high-wattage flexible loads. These could include EVs and thermostatically controlled loads like HVAC and WHs. In addition, it could also involve some discharging of battery storage systems to reduce the net load. The distribution of net load reductions across the remaining SMOs is shown in Fig. 5.38, with a total decrease of around 14% as seen in Table 5.5. From Table 5.5 we also see that the attack would have potentially increased the power import from the transmission grid by over 37%, but the combination of increased local generation and load curtailment helps keep the imported amount almost the same as before.

Effects at the transmission level

The overall impact of the generation shortfall and mitigation using EUREICA is simulated in the RTDS using a proxy where the individual feeders are not modeled, but the aggregated effect is studied at the transmission scale. A combined shortfall of 359 MW, corresponding to a simultaneous compromise and outage of 650 kW in all 552 primary feeders in Area 2, triggers a frequency event (see Fig. 5.39). Left unchecked, this can potentially lead to drastic load shedding or parts of the system being blacked out. To mitigate this situation, the

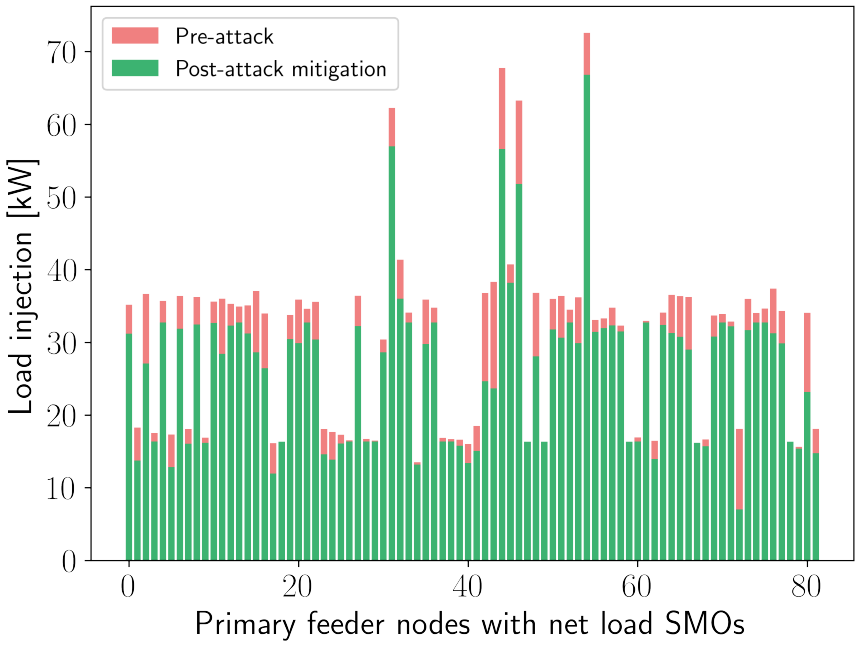

Figure 5.38: Flexible load curtailment for large attack mitigation.

power flow from Area 1 to Area 2 needs to be increased, which was observed in the RTDS, through the action of the governor system, which responds in the timescale of seconds, by increasing generation from the other generators present in the system proportionately based on a droop value. This increases the power flow from the generation-rich Area 1 to Area 2. However, changing the tie-line power flow creates a frequency imbalance, resulting in the system frequency oscillating, and settling at a lower/higher frequency, as shown in Fig. 5.39. With the EUREICA framework, the frequency mismatch is mitigated by suitably leveraging the flexibility of the remaining generation as well as demand response (DR) mechanisms from flexible loads at both the SMO and SMA levels (see Fig. 5.40). Once the governor response is completed and the system settles at a sub-optimal frequency, a combination of intelligent DR and generation redispatch in Area 2 facilitated by the EUREICA framework allows the system frequency to be restored to normal, ensuring grid resilience, avoiding system stress, and increased operational costs.

Key system metrics, economic, and distributional impacts

In our simulations, we find that attack mitigation comes at the expense of increased operational costs for the PMO since it needs to dispatch more expensive local resources to a greater extent, rather than importing cheaper power from the main grid (at the LMP rate). The PMOs and SMOs also need to adequately compensate agents for the critical flexibility they provide. As shown in Table 5.5 for attack 2b, the attack increases the system operating costs by around 7%, and the mitigation steps raise the cost by over 31%, both relative to the pre-attack case. However, the PMO could recoup this through other revenue streams and cost savings. For example, the transmission system operators may compensate PMOs for locally containing attacks. Being able to leverage local DER flexibility through markets could

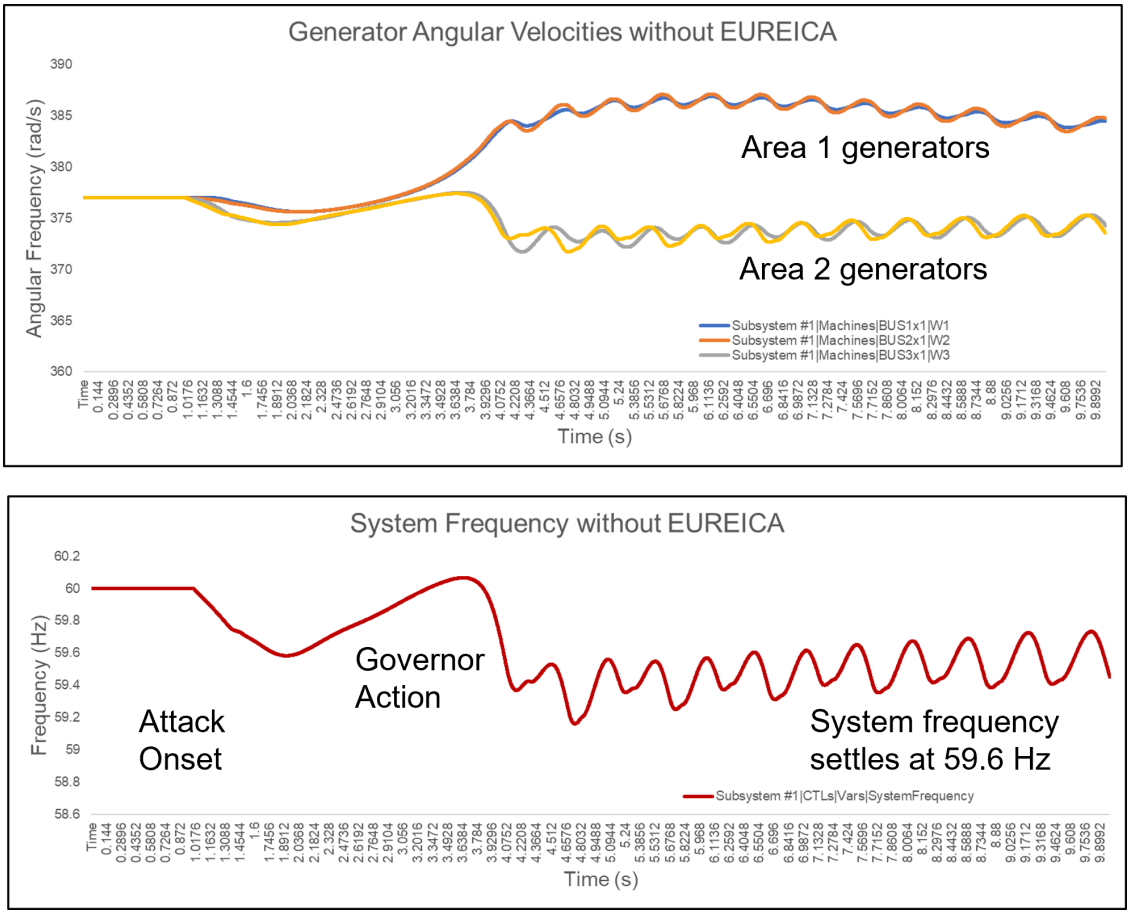

Figure 5.39: Response without EUREICA; system settles at sub-optimal frequency

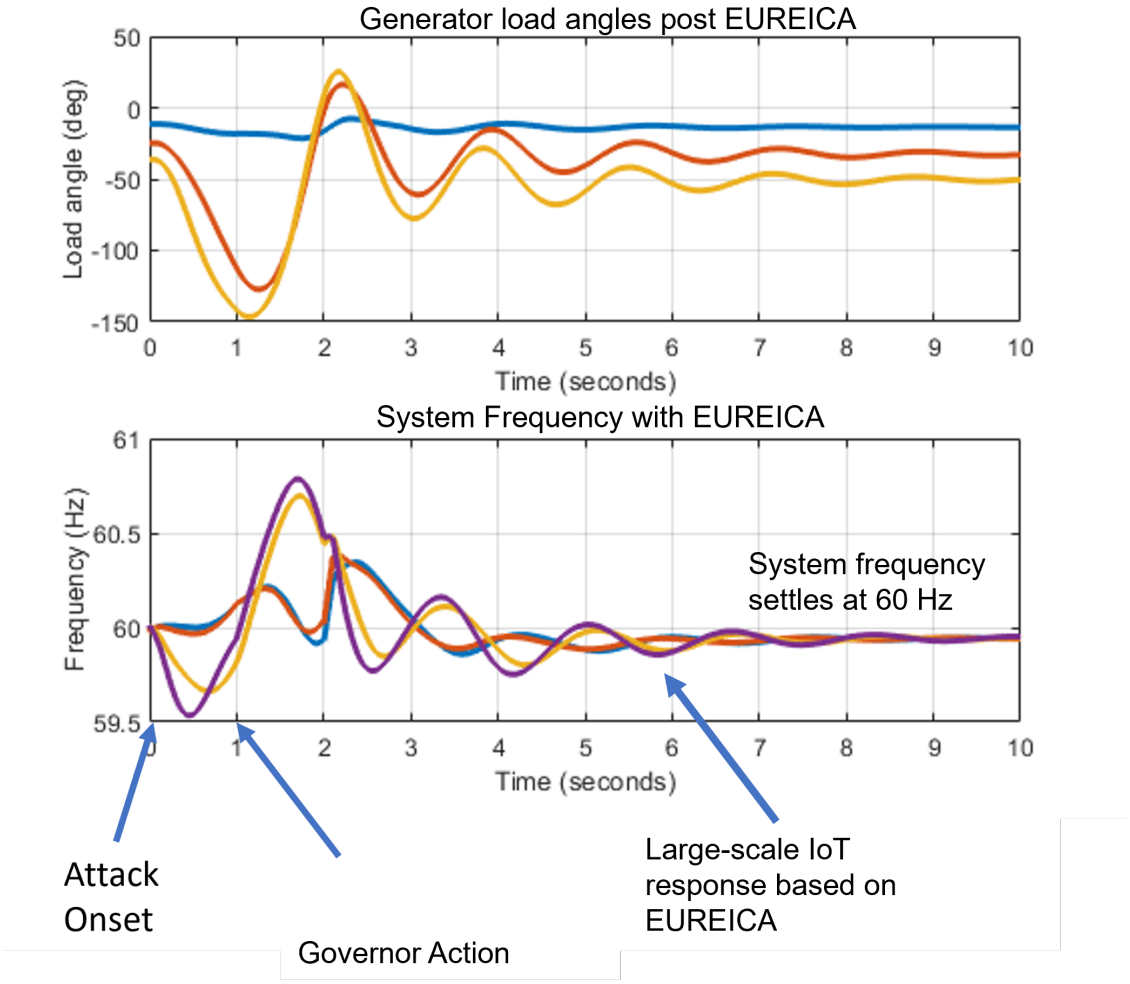

Figure 5.40: Frequency response with EUREICA; system settles at 60 Hz following demand response and load shedding enabled by the EUREICA framework

Table 5.5: Summary of metrics for large-scale attack 2b scenario.

	Pre-attack	**Post-attack**	**Attack mitigation**
Power import from main grid [kW]	1,325	1,821 (+37.4%)	1,328
Total cost [$]	10,752	11,500 (+7%)	14,156 (+31.7%)
Total load [kW]	2,064	2,023 (-0.02%)	1,775 (-14%)

also reduce the amount of auxiliary backup generation that the PMO needs to maintain, and lower the reserves it may have to otherwise procure from capacity or ancillary service markets. The PMO, in turn, could also redistribute some of these benefits among the SMOs and SMAs.

We can also obtain the electricity prices in the PM from the dual variables associated with the power balance constraints in Eq. (5.6). We refer to these as distribution-LMPs (d-LMPs) at each node (with an SMO) in the primary feeder. We compared the normalized d-LMPs for active power before and after the attack, as well as post-attack mitigation, shown in Fig. 5.41. As intuitively expected, we see that nodal prices increase throughout the grid after the attack and rise even further after implementing the attack mitigation steps, indicating that the loss of some local generation makes it more expensive to satisfy network constraints and results in sub-optimal solutions. The pre-attack and post-attack prices have nearly the same spatial profile across all the SMO nodes, with the post-attack values essentially being higher by an offset. This makes sense because the d-LMP variations between nodes are influenced by congestion on lines. In the attack case without mitigation, the shortfall caused by the attack would've been compensated for entirely by importing extra power from the grid, and thus the relative congestion variation over the rest of the network remains largely unchanged. The price trends after attack mitigation look different since the changes in power flow and congestion (resulting from the PM re-dispatch) are not uniform throughout the network. Notably, we see that the prices are significantly more volatile, especially around the nodes affected by the attack. The price also peaks at node 67 - this makes sense since it has the highest increase in injection after attack mitigation, which in turn worsens congestion in lines connected to it.

Another important consideration is the impact of our mitigation approach on the different market participants, i.e., the SMOs and SMAs themselves. The objective function update rules from Sections 5.10.1 and 5.10.2 generally imply that these local resources will be compensated less per unit (kW or kVAR) of grid support they provide, either in terms of load flexibility or generation dispatch. It may also lead to significant load shifting and curtailment in order to meet grid objectives, which can reduce the overall utility of end-users. We also need to more carefully study the distributional impacts of such methods since they may end up disproportionately negatively impacting certain groups of customers or prosumers, which could in turn have important implications for energy affordability, equity, and fairness.

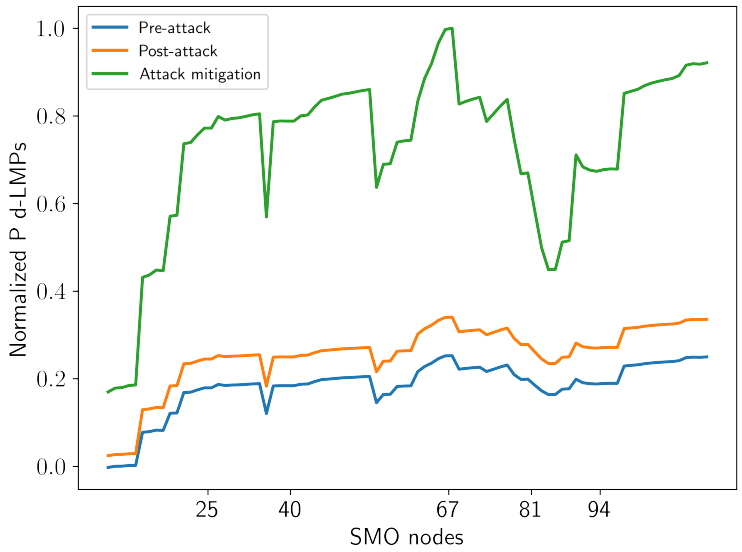

Figure 5.41: Effects of large-scale attack and mitigation on nodal d-LMPs at SMO nodes.

5.13.6 Mitigation of attack 3

We now consider the attack scenario where the distribution grid is isolated from the transmission system. In such a case, the distribution grid is fed through an alternate circuit, such as from node 350 (see Fig. 5.42). A typical response in such a case is that the distribution grid breaks into several "zones", creating smaller islands where only a portion of the load is fed through any DERs that may be present. We show below that with the increased awareness provided by the EUREICA framework, a much higher percentage of consumers remain unaffected, by suitably leveraging the DERs at node 48, the microgrid system connected at node 65 (marked by the red circles in Fig. 5.42, and DR methodologies. To ensure feasibility and supply-demand balance with islanding, we also introduce two large diesel generators located at nodes 48 and 65, which may only be called upon when the feeder is islanded. Three cases are presented.

Critical loads distributed across the feeder

In this case, through the proposed resilience-based IoT load restoration with DR optimization strategy (see Section 5.11, a feasible reconfiguration path is computed to open or close tie switches and completely or partially shed non-critical grid edge loads using reconfiguration to allow the available generation resources to cover approximately 30% of total load in the system. As seen in Fig. 5.43, with almost 70% of the load shed (second graph from the top) between 13:00 and 14:00, and batteries only allowed to discharge, if possible, to supply extra energy (third graph from the top), the burden on the diesel generators is significantly alleviated as they only need to ramp up to about 230 KW. These results were validated using the HELICS co-simulation platform at PNNL. Additional validation results using HELICS and LTDES are included in Section 5.13.6 and Section 5.13.6, respectively.

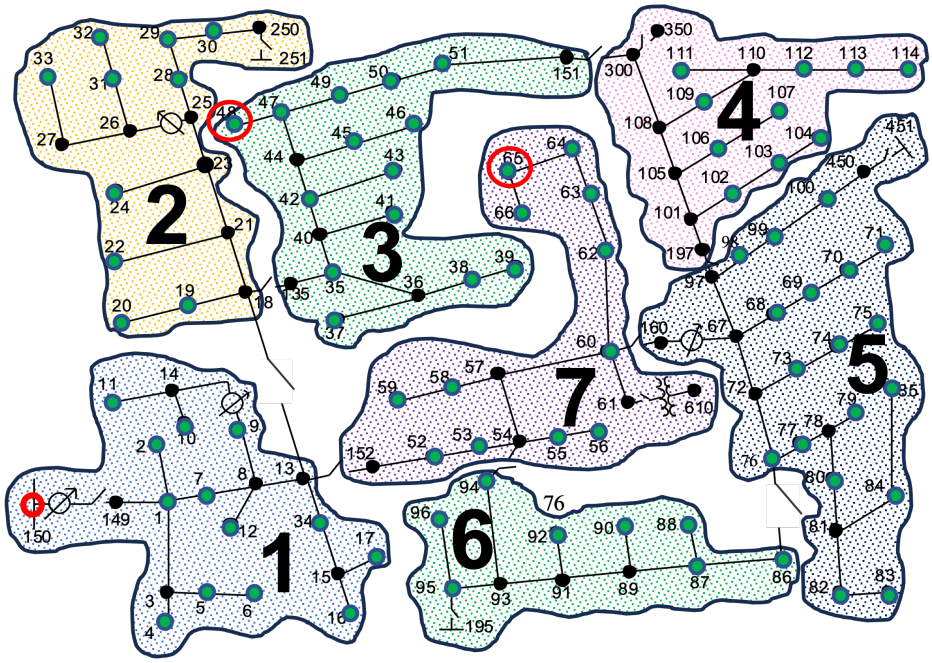

Figure 5.42: EUREICA IEEE 123-node feeder for reconfiguration module validation.

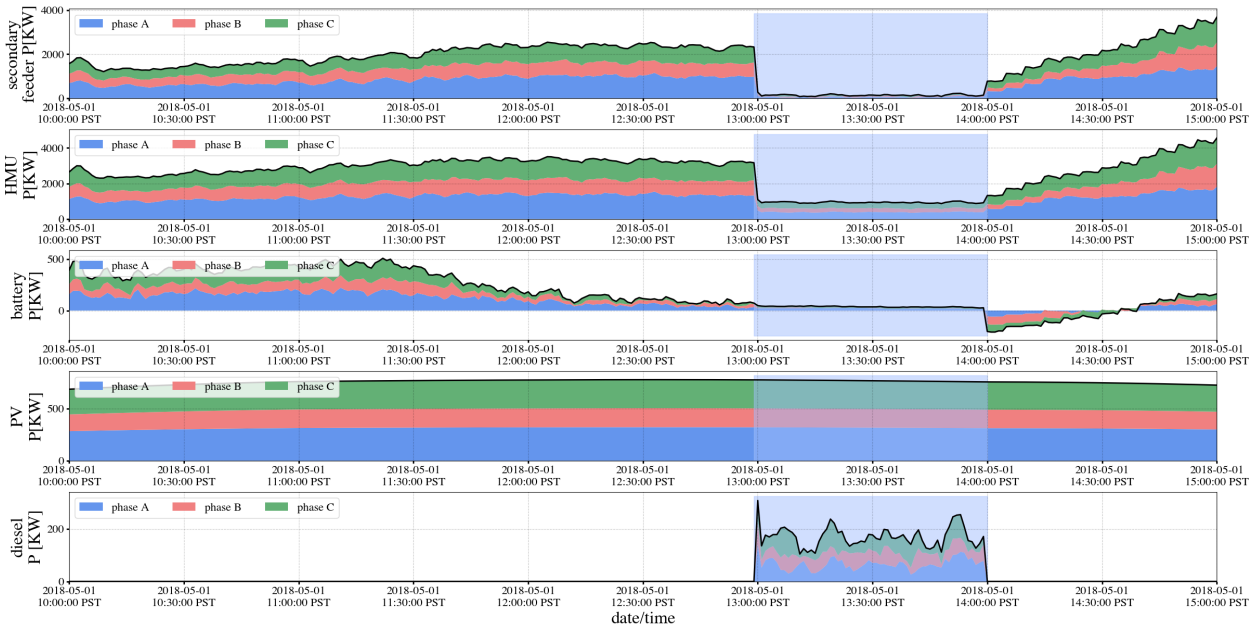

Figure 5.43: Demand and DG with resilience-based reconfiguration during attack 3.

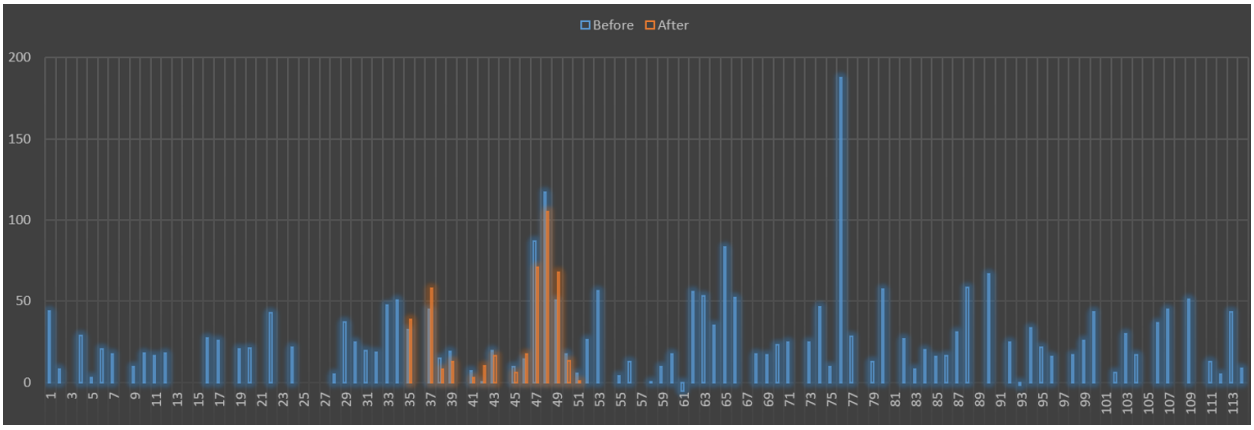

Figure 5.44: Primary node load change between 12:59 (before) and 13:00 (after attack).

Critical loads aggregated in a single zone

In this case, the SA from EUREICA helps the reconfiguration algorithm to disconnect or open the switches 18-135 and 151-300 to island zone 3 and pick up only the critical loads in this zone using the DG at node 48, which is a total of 430 kW. The results from this case are shown in Fig. 5.44. These results were validated using the DERIM and ADMS-DOTS software at LTDES. Additional results are included in Section 5.13.6.

Mitigation with a military microgrid

We assume that there is a military microgrid at node 66 in the primary circuit, which serves as a backup directly in the distribution system. Under current regulations, defense-critical systems have to be disconnected and isolated in the event of contingencies. Since EUREICA has the ability to identify trusted resources, our thesis is that there is confidence in the security of this resource as well as in meeting the power flow requirements, making it feasible to use this additional resource for attack 3 mitigation. First, the fault is isolated using reconfiguration based on the algorithm described in Fig. 5.9. The reconfiguration algorithm returns the most resilient path for implementation. In this case, only one feasible path is present, so that is chosen. This islands the feeder by opening the switch between nodes 150 and 149 and connecting the switches to the DG and microgrid at nodes 48 and 66, respectively. Then, a combination of ≈ 1.7 MW from the microgrid at node 66, 560 kW from the DG at node 48, and customer-side DR is utilized to pick up approximately 80% of the total load of the feeder. Some results from this case are shown in Fig. 5.45, validated using the ARIES platform at NREL. The complete set of results can be found in Section 5.13.6.

Validation of Attack 3 by PNNL using HELICS

The enhanced EUREICA IEEE 123-node feeder is covered partially by the local distributed generation resources, that is, the PVSs and BESSs, and mainly from the main transmission and generation grid through a connection at node 150, as shown in Fig. 5.42. Moreover, the system has available 2 large diesel generators at buses 48 (150 kVA rated capacity) and 66 (1 MVA rated capacity), respectively, that could be called upon to serve loads in case of

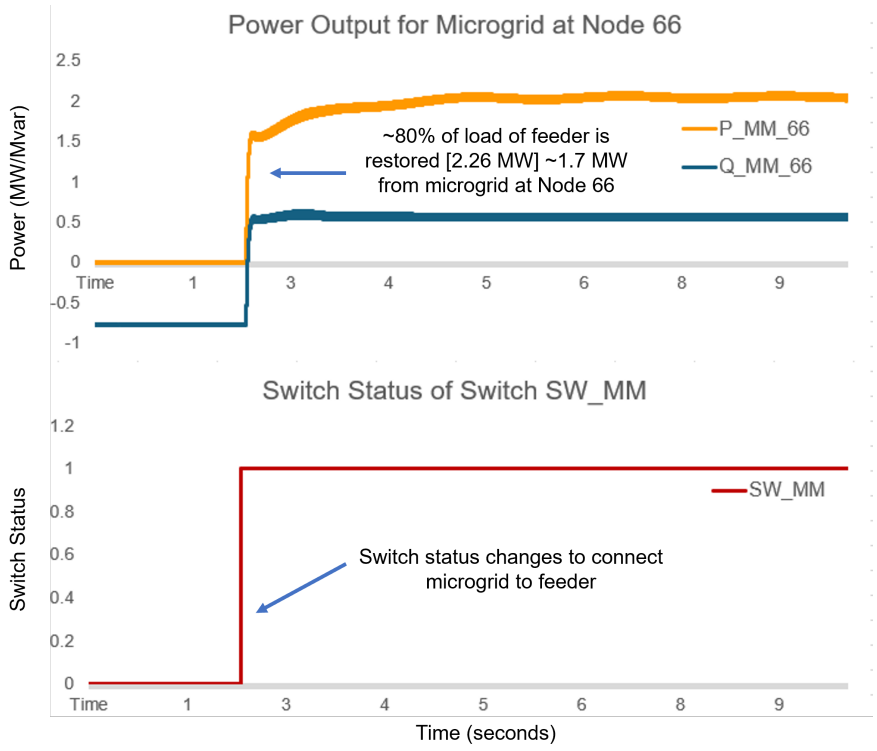

(a) Microgrid response after reconfiguration.

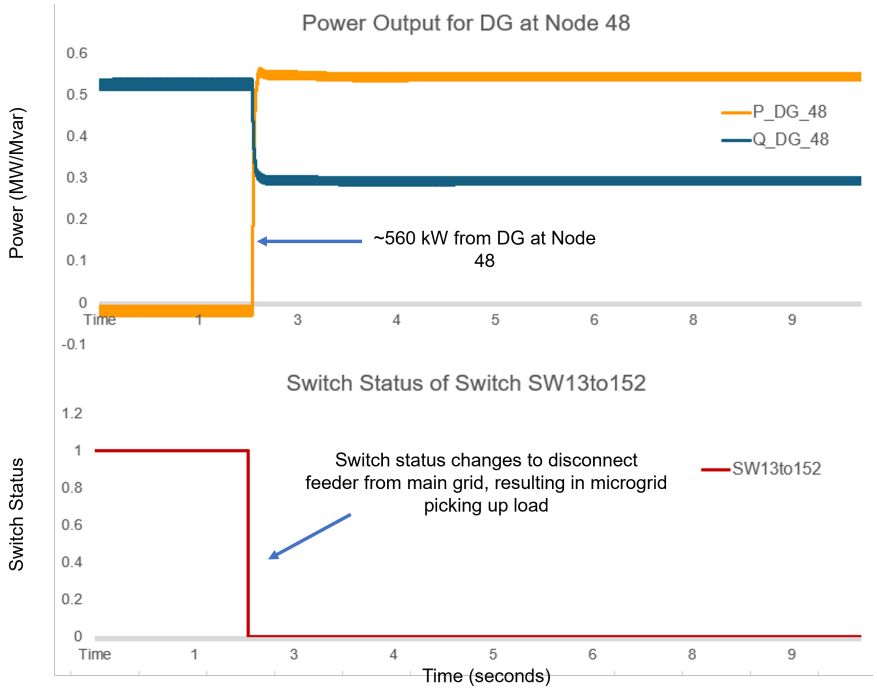

(b) DG response after reconfiguration.

Figure 5.45: System response after reconfiguration with microgrid.

Table 5.6: Original switch configuration in the EUREICA IEEE 123-node test feeder.

Node A	Node B	Switch status
13	152	CLOSED
18	135	CLOSED
60	160	CLOSED
61	610	CLOSED
97	197	CLOSED
150	149	CLOSED
250	251	OPEN
450	451	OPEN
300	350	OPEN
95	195	OPEN
54	94	OPEN
151	300	OPEN
13	18	CLOSED
86	76	CLOSED
48	48_dg	OPEN
65	65_dg	OPEN

adversarial events. Also, a set of switches between certain nodes of the system configures it into 7 areas that could be isolated in certain scenarios to be able to serve critical loads, as in Fig. 5.42. The initial configuration of the switches is given in Table 5.6.

The validation scenario assumes that due to an adversary event, either a cyber attack or a physical phenomenon, the distribution system gets islanded from the main grid, which is simulated by opening the switch between nodes 150 and 149 at 13:00. The system reconnection to the main grid is assumed to happen at 14:00. As expected, at 13:00 the system collapses, which is demonstrated by the sudden drop to 0 for all the spot-load bus voltages, as seen in Fig. 5.46.

The proposed reconfiguration and load shed approach addresses the situation created at 13:00 hours, creates situational awareness, and decides the switch statuses and loads that might need to be shed. If only the available diesel generators are brought online once the system is disconnected from the grid, by reconfiguring the corresponding switches' status, the system would not be in a blackout, as seen in Fig. 5.47.

However, as seen in Fig. 5.48, to supply the entire house population load (the total measurements from the house management units in the second graph from the top), even with the support of the PVs and batteries, the diesel generators would still need a total capacity of over 2 MW (as seen in bottom-most graph of Fig. 5.48), which is more than the maximum capacity of the model diesel generators.

Through the proposed resilience-based IoT load restoration with demand response optimization strategy, a feasible reconfiguration path is computed to open and/or close tie switches and shed either completely/partially grid edge loads to allow the available generation resources to cover the approximately 30% critical load in the system, as identified in Table 5.3. As seen in Fig. 5.43, with the almost 70% load shed (second graph from the top) between 13:00

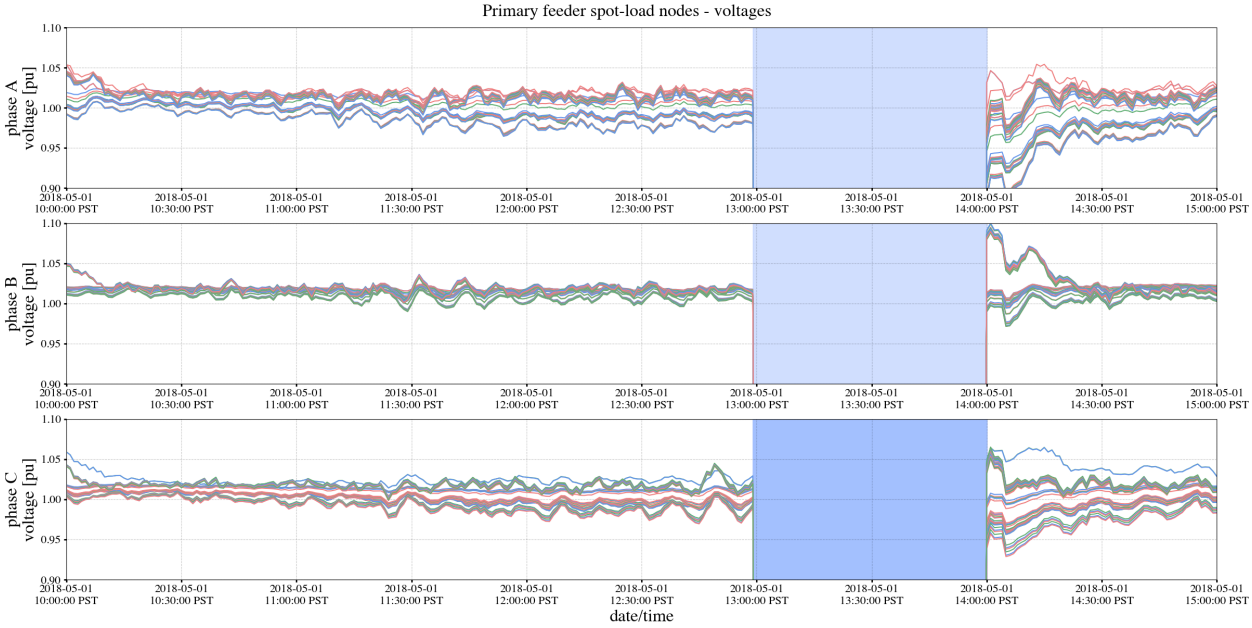

Figure 5.46: Black-out as a result of distribution system islanding in attack 3.

Figure 5.47: Voltage recovery after engaging diesel generators during attack 3.

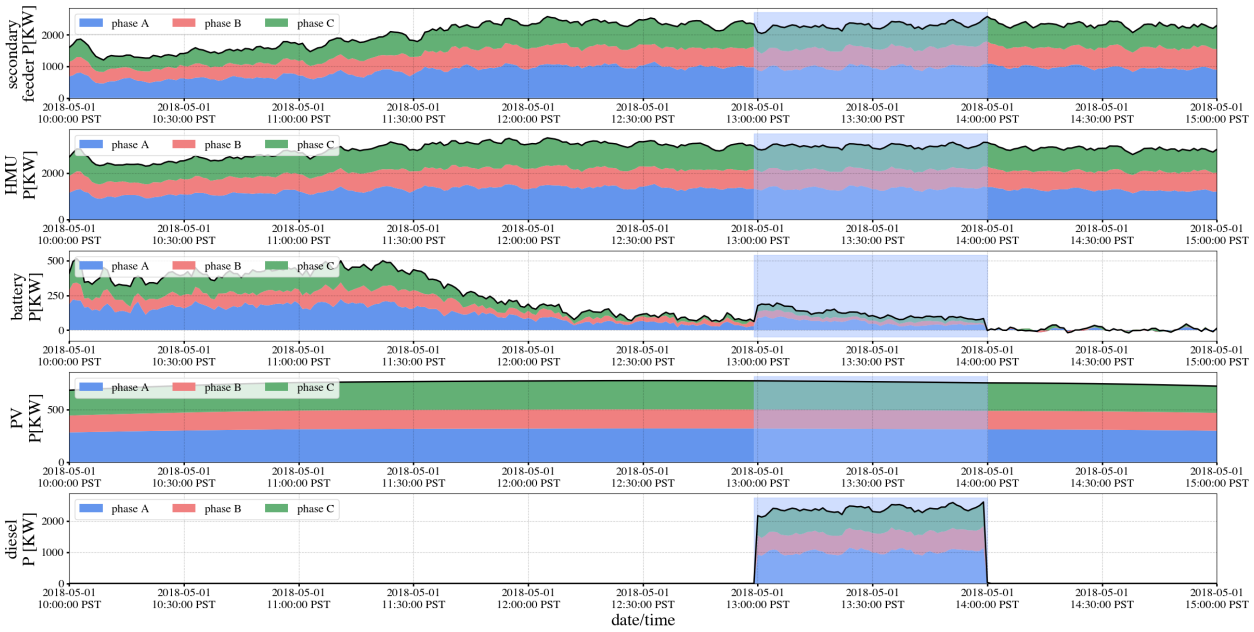

Figure 5.48: Demand and distributed generation without resilience-based reconfiguration during attack 3.

and 14:00 hours, and batteries only allowed to discharge, if possible, to supply extra energy (third graph from the top), the burden on the diesel generators is significantly alleviated as they only need to ramp up to about 230 KW.

For the islanded attack, the power flow redirection through switch reconfiguration as in Fig. 5.9 and load shed also helps with keeping the spot-load buses' voltages within the admissible limits during the attack (Fig. 5.49). Moreover, by bringing the loads back online sequentially after system recovery, under-voltage problems due to load rebound are also avoided.

Validation of Attack 3 by LTDES using DERIM and ADMS-DOTS

Case 1: Critical loads distributed across the feeder Fig. 5.50 shows the new switch settings and updated topology after applying the resilience-based reconfiguration during attack 3. In this case, we assume that there are critical loads distributed throughout the feeder. The system is islanded from the main grid at 13:00 PST and islanding ends at 14:00.

The DG at node 48 can output up to 270 kW while DG 65 outputs a constant 15 kW, and we use node 150 as the swing node (or slack bus) for the simulation. Fig. 5.51 shows the impact of the attack on the total feeder load, and Fig. 5.52 shows the changes in the net load at all primary nodes without the attack and with the attack (and associated reconfiguration). We see that the DGs at nodes 48 and 65 together pick up about 300 kW of the critical load, which represents about 20% of the total baseline load. The remaining 80% of the feeder load, which is non-critical, is shed (goes to zero after the attack in Fig. 5.52) to maintain feasibility.

Case 2: Critical loads aggregated in a single zone Fig. 5.53 shows the new switch settings and updated topology after applying the resilience-based reconfiguration during

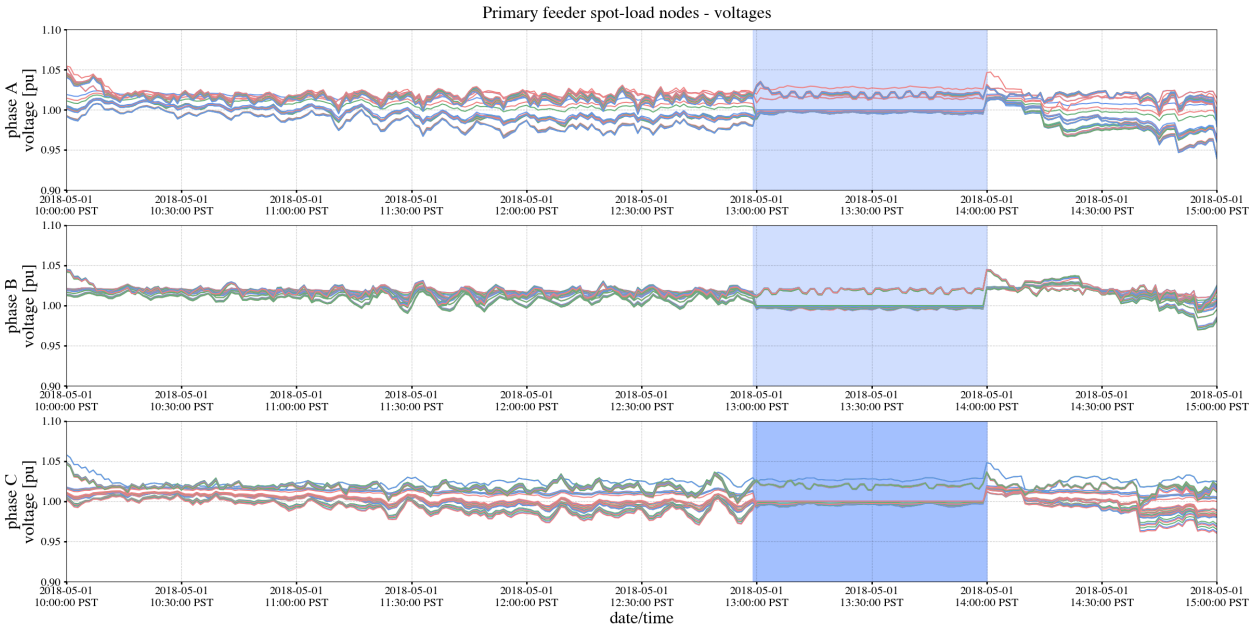

Figure 5.49: Voltage recovery after resilience-based reconfiguration during attack 3.

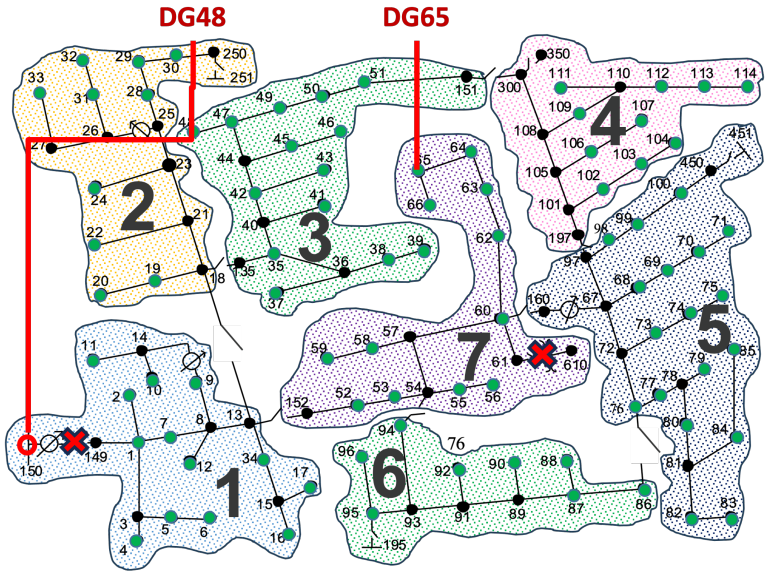

SWITCH	STATE
150-149	OPEN
61-610	OPEN

JUMP	STATE
150-48	CONNECTED

DG	STATE
DG48/65	CONNECTED

1. Islanding happens at 13:00 and ends at 14:00
2. DG 48 can output up to 270 kW
3. DG 65 can output constant 15 kW
4. Using Node 150 as swing node

Figure 5.50: Switch setting changes and network reconfiguration in the case when there are critical loads throughout the feeder.

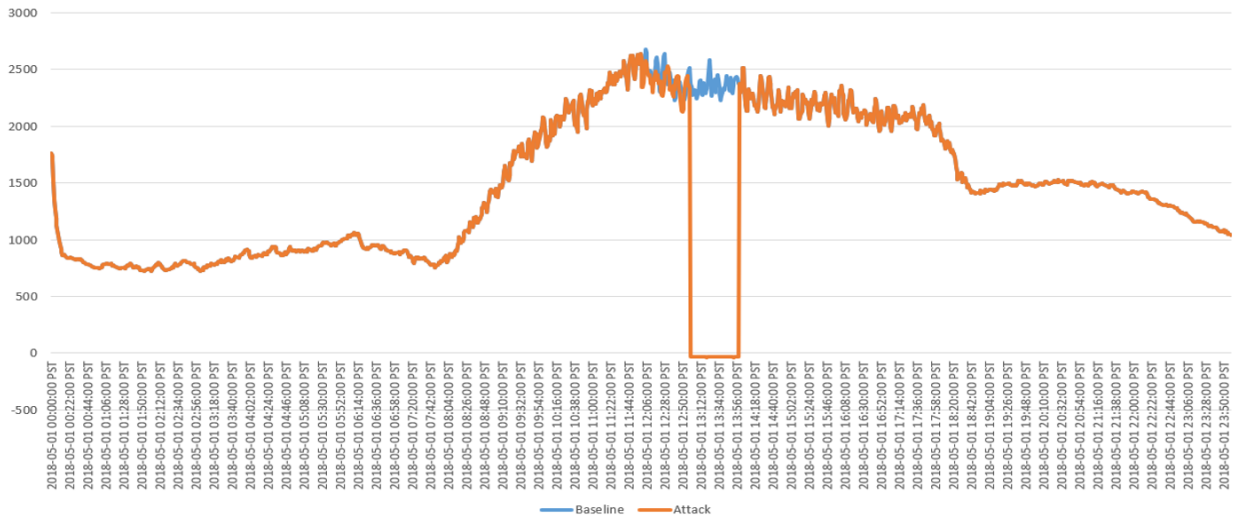

Figure 5.51: Total feeder head load over 24 h simulation, when there are critical loads throughout the feeder.

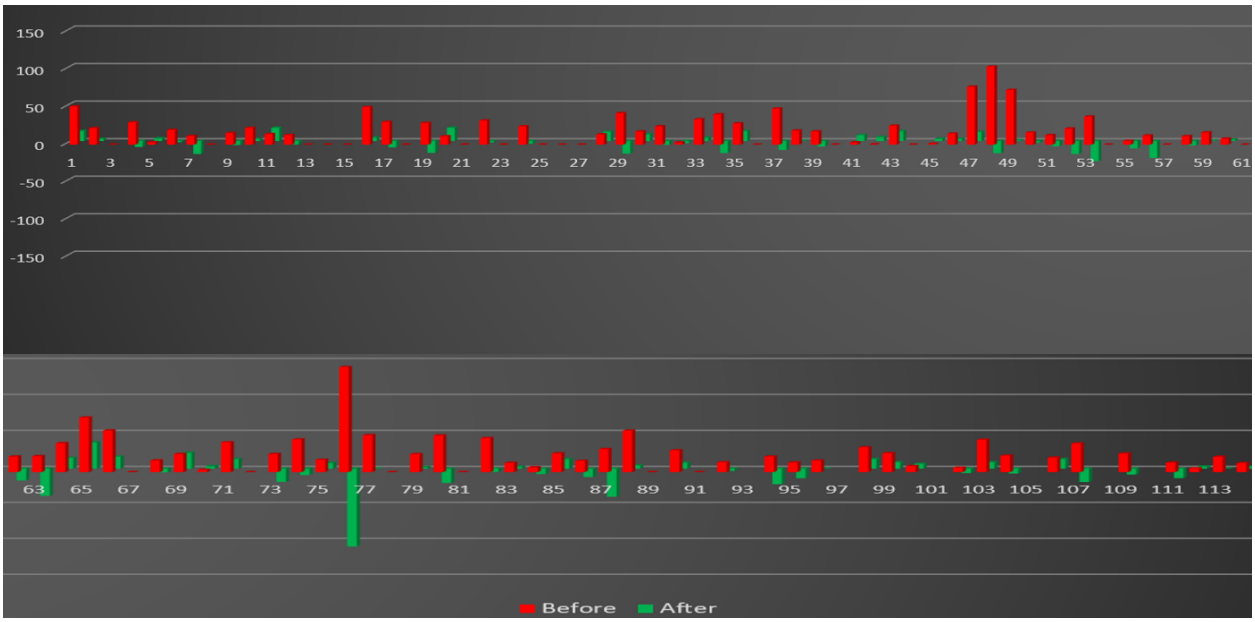

Figure 5.52: Primary node load change during attack 3 between 12:59 and 13:00 PST, when there are critical loads throughout the feeder.

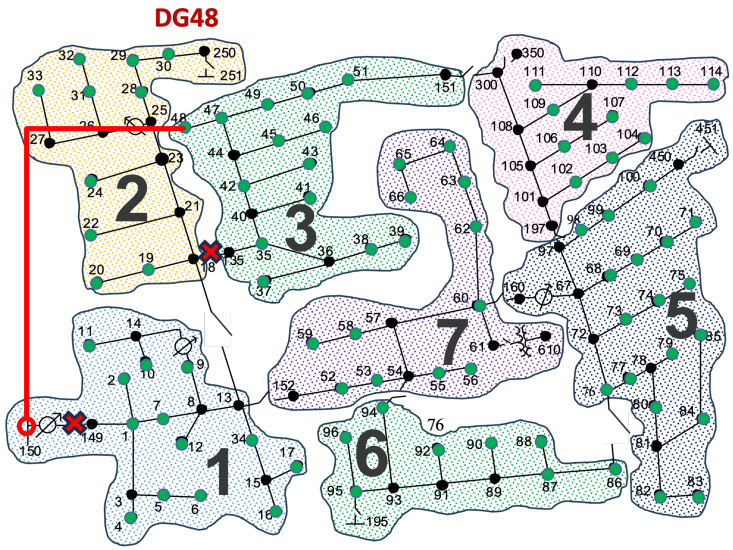

SWITCH	STATE
150-149	OPEN
18-135	OPEN

JUMP	STATE
150-48	CLOSED

DG	STATE
DG48	CONNECTED

1. Islanding happens at 13:00 and end at 14:00
2. Switch 18-135 open to create an microgrid
3. DG 48 has enough generation capacity to maintain region 3 load

Figure 5.53: Attack 3 case where critical loads are only located in zone 3 as a microgrid.

attack 3. In this case, we assume all the critical loads are concentrated only in zone 3. The system is islanded from the main grid at 13:00 PST and islanding ends at 14:00. During reconfiguration, switch 18-135 is opened so that cluster 3 becomes a microgrid. The DG at node 48 has sufficient capacity to meet all the load in zone 3 alone. Again, node 150 is used as the slack node for the simulation.

Fig. 5.54 shows the changes in the net load at all primary nodes without the attack and with the attack (and associated reconfiguration), as observed from the simulation results. We see that the DG at node 48 picks up all the expected load in zone 3 with 430 kW of generation output.

Thus, the main conclusions from the attack 3 validation using DERIM-ADMS-DOTS are as follows. Under case 1, the reconfiguration algorithm is able to restore all the critical loads throughout the islanded distribution circuit without relying on any power from the external transmission grid. In case 2, all the load in zone 3 (as a microgrid) was completely restored without any loss of load. In both cases, without the SA provided by the EUREICA framework, the control center operator at the substation would not have the necessary means to achieve restoration.

Validation of Attack 3 by NREL using ARIES

The mitigation of Attack 3 is validated using the RTDS at NREL-ARIES. The implementation of the reconfiguration algorithm in the RunTime environment of RTDS is shown in Fig. 5.55.

In the case where the EUREICA framework is not used, the frequency of the system

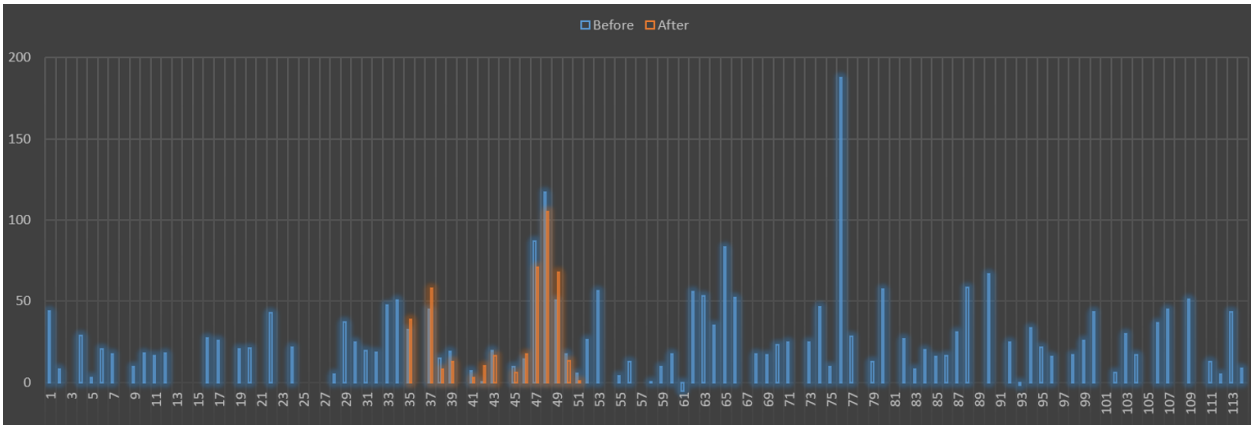

Figure 5.54: Primary node load change during attack 3 between 12:59 and 13:00 PST, when critical loads are only located in zone 3 as a microgrid.

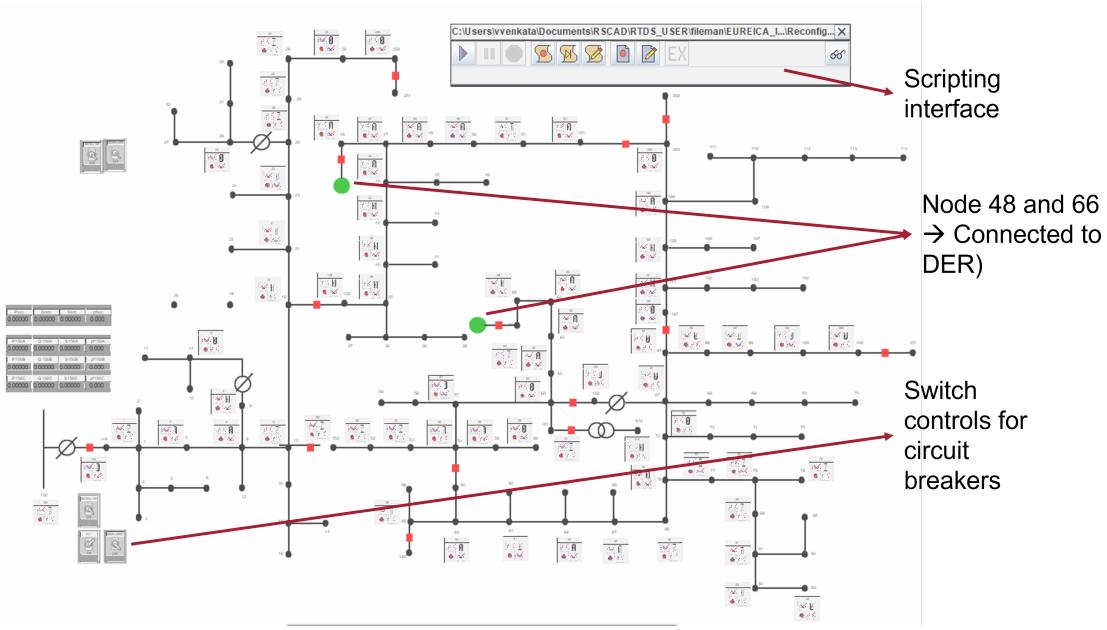

Figure 5.55: Implementation of reconfiguration algorithm in RTDS.

becomes unstable, and the distribution feeder is broken into islands, and only the loads in Zone 3 are picked by the DG in Node 48. This plot is shown in Fig. 5.56.

The case with the EUREICA framework, with the contributions from various DGs and the military microgrid connected at Node 66, has already been demonstrated in Section 5.13.6.

5.14 Conclusion

We have proposed a framework, EUREICA, for achieving grid resilience through the coordination of IoT-Coordinated Assets that are trustable. A local electricity market that has been previously shown to lead to grid reliability and provide services such as voltage support and overall power balance is leveraged in this framework to ensure grid resilience. The local market accomplishes this through SA to co-located operators. This SA consists of information about DERs and their power injections, as well as their levels of trustability, commitment, and resilience. With this SA, we have shown that a range of cyberattacks can be mitigated using local trustable resources without stressing the bulk grid. The demonstrations have been carried out using a variety of platforms with high fidelity, hardware-in-the-loop, and utility-friendly validation software.

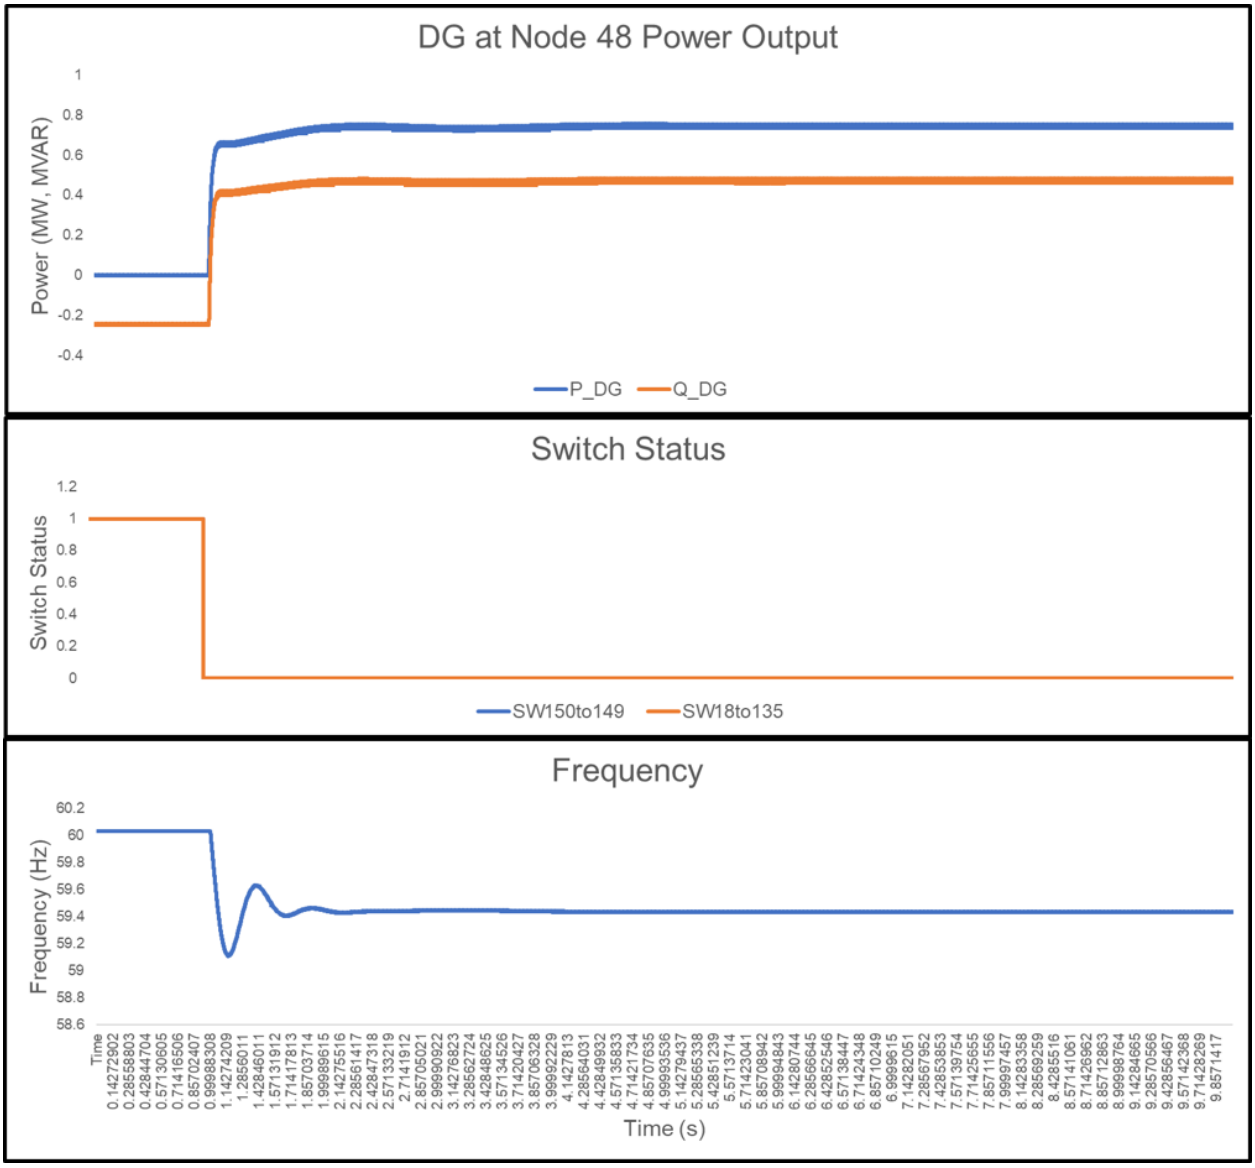

Figure 5.56: Distribution feeder broken into islands, with only Zone 3 load restored by DG at Node 48.

Chapter 6

DER coordination to boost grid hosting capacity

In order to decarbonize the grid, we need to rapidly integrate renewables like wind and solar, in addition to battery storage. Furthermore, the electrification of other sectors like heating and transportation is causing massive growth in electricity demand. Fig. 6.1 illustrates the fast growth of DERs, with the market value shown on the left and the total capacity growth on the right. We expect to see an additional 262 GW of DERs in the distribution grid between 2023 to 2027, which is comparable to the growth in utility-scale resources (+272 GW) at the transmission grid level over the same period. However, most of the conversation and industry focus today is centered largely on transmission grid capacity and utility-scale interconnection. However, it is crucial to ensure that the distribution grid also has enough capacity to accommodate the growth in installations of such small-scale resources. This chapter explores how our distributed coordination and optimization approach can help the distribution grid reliably and affordably accommodate more DERs. For the purpose of this work, we assumed that all the agents are independently owned and the coordination is enabled through our hierarchical local retail electricity market. However, note that our optimization approach and framework are more general and can be implemented even without the market in place. In this case, our decision-making tools would help entities like utilities, grid operators, or aggregators to optimally coordinate and manage their assets to enhance hosting capacity and grid reliability.

6.1 Introduction and background

Hosting capacity (HC) refers to the maximum amount of distributed energy resources (DERs) such as solar PV, battery storage, electric vehicles (EVs), and heat pumps (HPs) that can be integrated into the electric grid at a given location without causing reliability, power quality, or safety violations and without requiring major system retrofits upgrades. It is typically expressed in absolute terms in either absolute terms (kW) or in relative terms as a percentage of the peak load at that location. Hosting capacity is a critical metric for utilities and grid operators to assess the feasibility of integrating renewable energy sources and other distributed generation technologies, as well as new loads into the existing distribution grid.

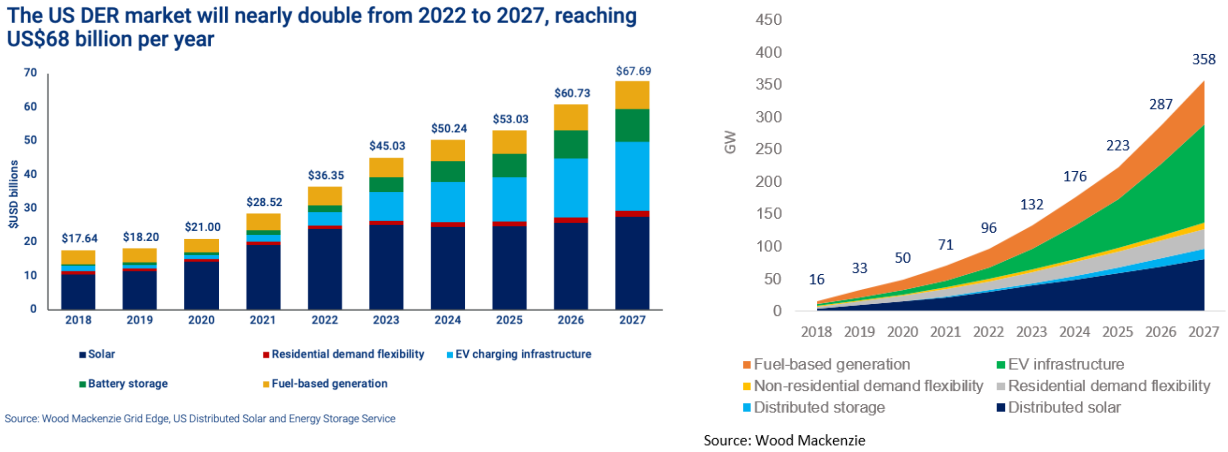

Figure 6.1: DER growth patterns.

Feeder hosting capacity (FHC) is the total capacity of a distribution feeder (i.e., primary distribution line from the substation) to host additional DERs without violating voltage, thermal, or protection limits [143]. It aggregates constraints over the whole feeder and is used for high-level planning. Locational hosting capacity (LHC) provides a more granular assessment of DER HC at specific locations (nodes or points of interconnection) and is more useful for utilities to assess individual interconnection requests, site-specific DER planning, and identifying optimal locations for DER deployment [144]. In this study, we focus mainly on FHC; extensions to LHC will be considered in future work. In fact, we can easily derive the FHC from our optimization formulation as well.

6.1.1 Factors limiting hosting capacity

HC may be limited by a variety of voltage, power quality, reliability, thermal, or operational constraints However, here we will mainly be focusing on voltage and line thermal constraints. Overvoltage and undervoltage limits are the most common voltage concerns for hosting capacity, since violating these upper or lower bounds can threaten grid stability. In particular, excess solar PV output during the middle of the day can lead to significant overvoltage issues. In addition to operational limits, utilities also prefer to maintain voltage profiles over the network as uniform as possible. Thus, undervoltage issues at the end of the feeder should also be avoided. This can make it more challenging to meet the massive demand growth that we're seeing today due to increasing electrification, as well as emerging issues like data centers. In addition to voltage bounds, voltage imbalances (across phases) can also be a concern for unbalanced three-phase systems. However, for our initial case studies, this does not apply since we assume that the systems are either already balanced or we convert them to this form. Apart from voltages, the other major limiting factors are the thermal limits on maximum line currents, or equivalently, their maximum allowable apparent power flows. Finally, the HC can also be limited by the substation transformer's rated capacity [in MVA]. Many utilities across the US are already severely HC-constrained as seen in Fig. 6.2, which delays or prevents the integration of new DERs in the distribution grid. The issue of hosting capacity is also closely related to the problem of renewables curtailment, wherein excess solar

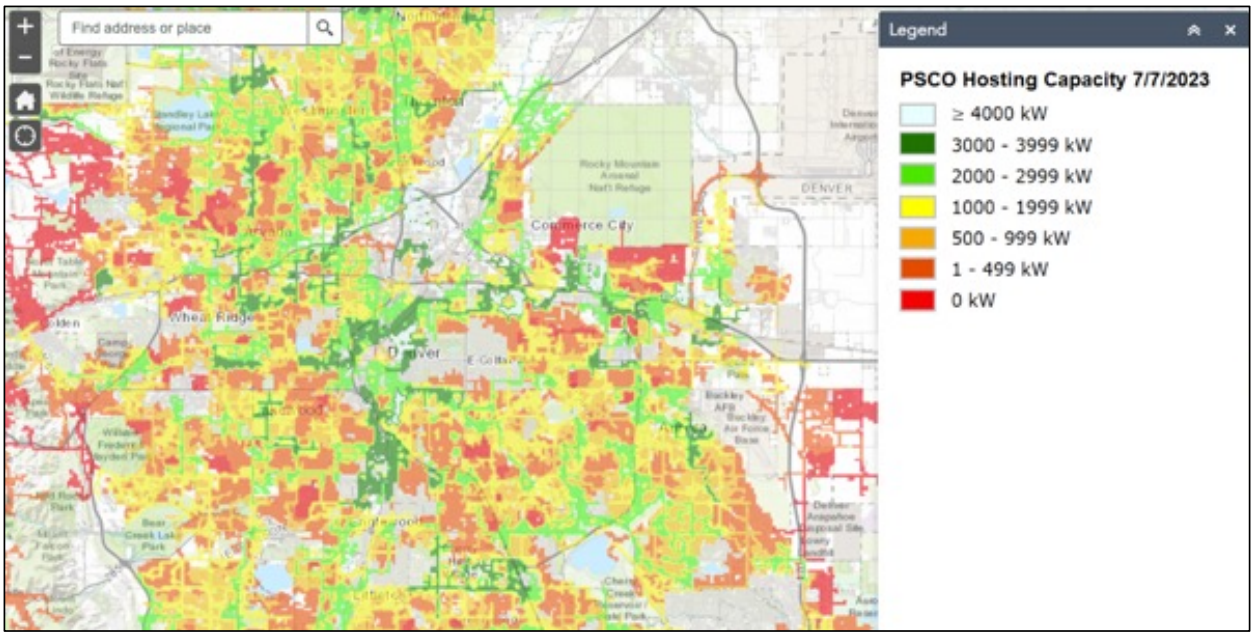

Figure 6.2: Example of HC map for Denver, CO.

PV output in the middle of the day is increasingly being diverted or wasted simply because the grid does not have sufficient capacity to transport that power. This leads to overbuilding of capacity, increased capital costs, and higher electricity rates. Ideally, we would like to avoid curtailment of clean power and instead utilize PV output productively (e.g., to charge batteries/EVs).

6.1.2 Different methods for HC analysis (HCA)

There are several different approaches to conducting HCA. These broadly fall under three types: Deterministic, non-deterministic, and time series methods [145]. The static hosting capacity analysis method evaluates the distribution system's ability to host DERs under specific, unchanging conditions, often considering the most extreme and worst-case scenarios. These result in generally lower, more conservative estimates for HC [146]. The time-series hosting capacity analysis assesses HC over longer time intervals, accounting for fluctuations in load and generation, which is particularly necessary for variable and intermittent renewables [147, 148]. Initially, we run the analysis for a shorter, 1h time horizon (12-1pm) since we're focused only on PV hosting capacity, and midday is when overvoltage issues are most likely to occur. However, we need to eventually consider longer time horizons (e.g., 1 day to 1 year) to fully capture different possible conditions, since other periods of the day could also have risks of limit violations. For example, periods of high load in the morning or evening/night (e.g., due to HP use or EV charging) could risk undervoltage or thermal overloading of lines. Note that each scenario's data consists of timeseries vectors of length (time dimension) = the time horizon T. To simplify the analysis, some studies take a worst-case time period approach and only consider specific hours of the day, e.g., those periods with peak PV output or peak load, etc. [149].

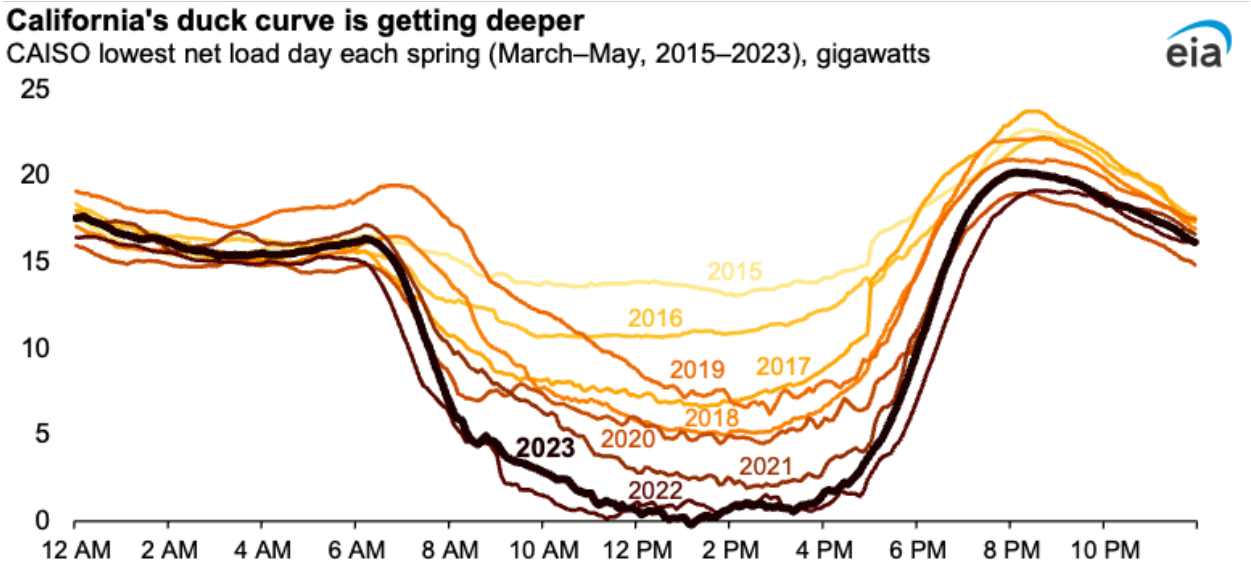

Figure 6.3: Net load "duck curve" showing solar curtailment in CA.

The final category consists of non-deterministic, stochastic, and probabilistic hosting capacity analyses. Here, we consider a non-deterministic approach for HCA, as opposed to a purely deterministic method [150–152]. This incorporates uncertainties in DER output (generation) and load variations (which are also influenced by weather uncertainty) to provide a probabilistic assessment. In addition to uncertainties in DER power profiles, we can also consider the stochasticity associated with (i) electricity prices (LMPs) and (ii) the distribution of DERs (both locations and sizes/capacities) over the feeder. Finally, there are more advanced methods that utilize real-time data and adaptive controls to assess and enhance the grid's ability to integrate DERs. These include using active learning techniques [153].

Literature gap and our contributions

The majority of prior HC works have focused on maximizing only solar PV HC [**Seuss2015ImprovingSuppo**
Some works have assessed the HC of either EV, HVAC, or HPs [154], but generally in isolation. However, no prior works have studied the HC of multiple DERs simultaneously. Our contributions, therefore, are the following: (i) We conduct HC analysis while considering all major types of DERs together, i.e. Solar PV, EVs, batteries, and heat pumps, (ii) We develop a flexible framework to co-optimize various DER types together, and (iii) We accurately model the device-level dynamics of DERs to capture their true flexibility and controllability. Furthermore, our method also allows seamless incorporation of other loads like data centers, which are a top priority for utilities currently.

6.2 Approach to populate DERs

In the iterative approach, we define a certain penetration level for each DER type (PV, HP, EV, BS) and then randomly allocate the total DER capacity among select nodes over the nodes in the network, while placing some restrictions on the minimum and maximum capacities at any one node. We also draw from a beta distribution while populating DERs and vary the α parameter to control how uniformly the DERs are spread out amongst the nodes.

6.2.1 Defining and computing DER penetration

DER penetration is defined slightly differently depending on the type of DER. PV and BS penetrations are defined relative to the peak feeder load, with no upper limit or cap. HP and EV penetrations are defined in terms of the proportion (%) of homes/buildings in the feeder that have been electrified, and the upper limit is 100%. All air conditioners for cooling are already electric. But heating (both space and water) is largely still gas-based and will transition to electric heat pumps. These heat pumps (either retrofitted or newly installed) can also act as cooling devices, and the average heat pump load for a home is about 5-6 kW. The average US house load is 1.2 kW with a peak load of 3-6 kW [155]. The base load (fixed) is 0.5-0.8 kW (from appliances like refrigerators, standby electronics, etc.), while the peak load is from thermostatically controlled loads (TCLs) like ACs, heat pumps, water heaters, electric dryers, etc.

IEEE-123 example system configuration [156]

Residential Load Distribution: Each bus in the system typically represents multiple households. On average, there are about 8 houses connected to each bus, with each house having a peak load of approximately 3.5 kW.
Total Peak Load Calculation: Given that the system comprises 123 buses, the total number of houses is approximately 984 (8 houses/bus * 123 buses). Multiplying this by the average peak load per house (3.5 kW) results in an estimated total peak load of around 3,444 kW, or approximately 3.4 MW.

6.3 AC power flow model

We use the branch flow model with second-order conic program convex relaxation [129, 157], where P_i^L refers to fixed loads. This is assuming that we have a balanced and radial grid, which is a reasonable assumption for distribution networks in the US. Extensions to meshed and unbalanced, multi-phase grids will be considered in future work. Note that $\{k_j\}$ = set of all nodes connected to i (but excluding i). Node 1 corresponds to the substation transformer at the point of common coupling (PCC). $\mathcal{N}$ and $\mathcal{E}$ are the sets of all nodes/buses and edges/branches in the network, respectively. The power flow model is specified by the following constraints:

$$v_j - v_i = (R_{ij}^2 + X_{ij}^2) l_{ij} - 2(R_{ij} P_{ij} + X_{ij} Q_{ij}) \quad \forall (i,j) \in \mathcal{E} \tag{6.1}$$

$$P_{ij} = R_{ij} l_{ij} - P_j + \sum_{k \in \{k_j\}} P_{jk} \quad \forall (i,j) \in \mathcal{E} \tag{6.2}$$

$$Q_{ij} = X_{ij} l_{ij} - Q_j + \sum_{k \in \{k_j\}} Q_{jk} \quad \forall (i,j) \in \mathcal{E} \tag{6.3}$$

$$P_{ij}^2 + Q_{ij}^2 \leq v_i l_{ij}, \quad P_{ij}^2 + Q_{ij}^2 \leq \overline{S}_{ij}^2, \quad 0 \leq l_{ij} \leq \overline{S}_{ij}^2 / v_i \quad \forall (i,j) \in \mathcal{E} \tag{6.4}$$

$$P_j \in [\underline{P}_j, \overline{P}_j], \quad Q_j \in [\underline{Q}_j, \overline{Q}_j], \quad v_j \in [\underline{v}_j, \overline{v}_j] \quad \forall i \in \mathcal{N} \tag{6.5}$$

$$\text{where } l_{ij} = |I_{ij}|^2 \text{ and } v_i = |V_i|^2 \tag{6.6}$$

$$P_i(t) = P_i^{PV}(t) + P_i^{BS}(t) + P_i^{HP}(t) + P_i^{EV}(t) - P_i^{L}(t), \; Q_i = Q_i^{PV} + Q_i^{BS} + Q_i^{EV}(t) - Q_i^{L}$$

We also add the following additional constraints:

Transformer capacity constraints

$$P_{1i}^2 + Q_{ij_1}^2 \leq S_{tran}^2, \quad 0 \leq l_{ij_1} \leq \frac{S_{tran}^2}{v_{max}^2} \tag{6.7}$$

PCC power balance

$$P_1 = \sum_{j \in j_1} P_{1j}, \quad Q_1 = \sum_{j \in j_1} Q_{1j} \tag{6.8}$$

Feeder large capacity (slack bus at node 1)

$$\underline{P}^{PCC} \leq P_1 \leq \overline{P}^{PCC}, \quad \underline{Q}^{PCC} \leq Q_1 \leq \overline{Q}^{PCC} \tag{6.9}$$

where j_1 denotes all the nodes connected directly to the PCC at node 1. These underlying constraints apply to both the static and dynamic cases. We can solve this feasibility problem (without any objective function) to check whether the DER power injections and dispatch results are feasible to satisfy the grid physics.

6.4 Increasing hosting capacity with coordination

Fig. 6.4 provides an overview of the approach we use to increase the hosting capacity. We leverage complementary relationships among different types of DERS (distributed generation, storage, flexible demand) and leverage the flexibility available at the grid edge to dynamically increase hosting capacity. This does require flexible interconnection agreements, which are not yet widely adopted by utilities. This is in contrast to the static HC analysis, which takes a conservative or worst-case approach based on fixed interconnection agreements. In what follows, we show that dynamic coordination and distributed optimization can also help reduce solar curtailment and costs. Rather than seeing the rapid demand growth as a problem, we can take advantage of this growing flexible resource to benefit both grid planning and operation.

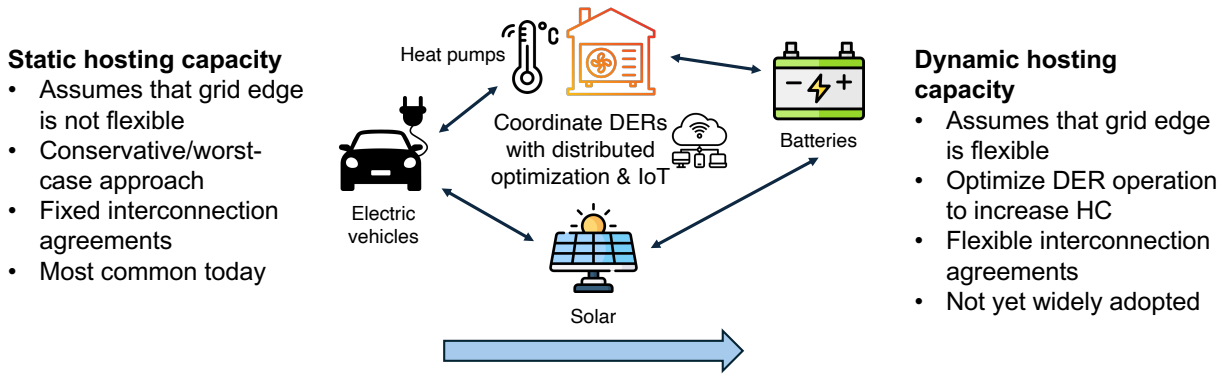

Figure 6.4: Hosting capacity increase overview.

6.4.1 Static case

For the static case, we just use baseline profiles for fixed loads, heat pumps, batteries, and solar radiation profiles for PV. We derived several profiles from historical data. Considering many scenarios also allows us to move towards non-deterministic or probabilistic results as well. We assume that these injections are not flexible and that there's no coordination or optimization among the DERs. We just run the power flow model (from Section 6.3) and check if the DER power injections are AC-feasible.

6.4.2 Local control case

Here, we assume that resources do have some flexibility and DERs can be optimized to increase their dynamic HC. However, only local optimization is done at each node. For e.g., each node tries to optimize charging its own batteries and EVs and/or curtailing HP load or PV output. This could be done to maximally utilize its PV self-generation to charge batteries and EVs, or to reduce their net load based on electricity prices. The primary options for control in the local case are as follows.

Smart inverter control for PV, batteries, and EV chargers

This involves local inverter control via either:

- Variable power factor control.
- $Q(P)$ droop curve-based Volt-VAR control: This will also entail some oversizing of the inverter capacity relative to its maximum rated active power.

6.4.3 Local DER active power optimization

This involves each node locally optimizing their BS and/or EV charging/discharging and flexible loads (HP) based on only *local* nodal information and the LMP data. No other information from other nodes is available. Thus, this local approach does not directly account for power flow constraints or grid physics during the control decision-making, which may

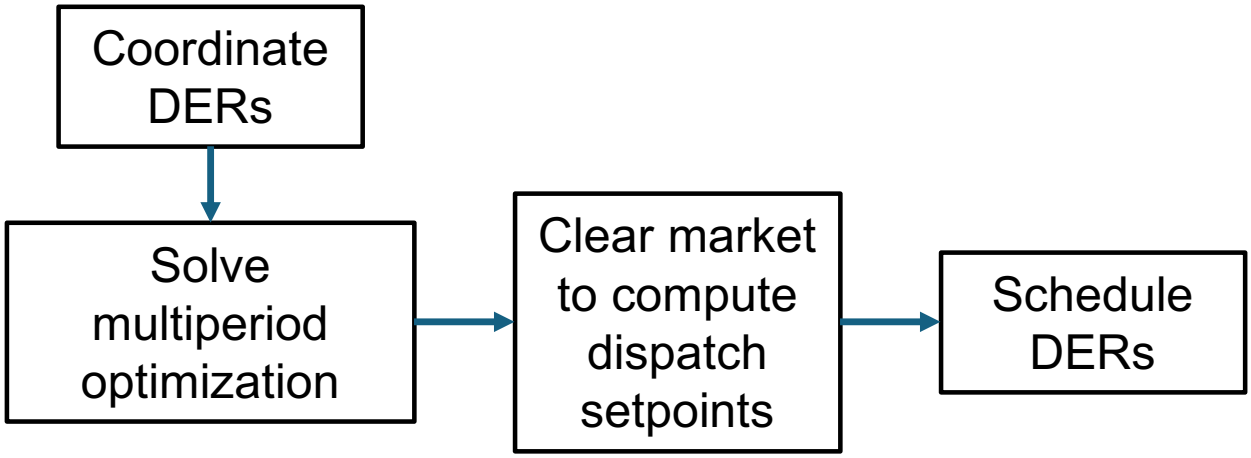

Figure 6.5: Dynamic optimization steps.

cause it to run into infeasibility more easily and also sacrifice some optimality. This motivates the need for coordination and a more global view.

AC Optimal Power Flow - dynamic coordinated case

In this case, we leverage network-wide global coordination over the whole feeder. This allows collaboration among different nodes rather than each agent just optimizing its own DERs locally. This is enabled by solving the global optimal power flow problem. The overall feeder is overseen by a distribution system operator (DSO), but the global optimization can also be solved in a fully distributed, privacy-preserving, and computationally efficient manner using a proximal atomic coordination algorithm [132]. A distribution-level retail market structure could be one approach to achieve this coordination, which would also allow us to compensate resources for their flexibility through pricing. However, we can also do this without markets, as long as there is infrastructure in place for communication and coordination among the different operators and agents. The ACOPF problem minimizes the following objective function in the dynamic case, subject to the power flow constraints from Section 6.3.

6.4.4 ACOPF objective (dynamic case)

$$\min f^{OPF} = \tag{6.10}$$

minimize PV curtailment (6.11)

$$\sum_{i \in bus} \sum_{t \in t_{H_p}} \beta_{pv} (P_i^{PV}(t) - \overline{P}_i^{PV} \alpha^{PV}(t))^2$$

minimize thermal line losses (6.12)

$$\sum_{(i,j) \in \mathcal{E}} \sum_{t \in t_{H_p}} l_{ij} R_{ij}$$

maximize PV usage for BS/EV charging (6.13)

$$+ \sum_{i \in bus} \beta_{coupl,bs} \alpha^{PV}(t) P_i^{BS}(t) + \beta_{coupl,ev} \alpha^{PV}(t) P_i^{EV}(t)$$

minimize absolute thermal discomfort + HP temp tracking (6.14)

$$+ \sum_{i \in bus} \sum_{t \in t_{H_p}} \beta_T (T_i^{in}(t) - T_i^*)^2$$

minimize HP and BS cycling (6.15)

$$+ \sum_{i \in bus} \sum_{t \in t_{H_s}} \beta_{hp} (P_i^{HP}(t+1) - P_i^{HP}(t))^2 + \beta_{bs} (P_i^{BS}(t+1) - P_i^{BS}(t))^2$$

minimize EV cycling and track desired SOC (6.16)

$$+ \sum_{i \in bus} \sum_{t \in t_{H_s}} \beta_{ev1} (P_i^{EV}(t+1) - P_i^{EV}(t))^2 + \beta_{ev1} (SOC_i^{EV}(t^*) - SOC_i^{EV*})^2$$

minimize cost of power import at PCC (at LMP rate) (6.17)

$$+ \sum_{t \in t_{H_p}} \lambda(t) P_{1,t}$$

encourage charging during low-price periods (when LMP is below average) (6.18)

or encourage discharging when LMP is high (use more local gen) (6.19)

$$- \sum_{i \in bus} \sum_{t \in t_{H_p}} (lmp(t) - mean(lmp)) P_i^{BS}(t) \tag{6.20}$$

Note that Eq. (6.14) encourages battery charging during daylight hours when solar PV is available. Thus, for the dynamic coordinated case, the overall optimization problem solved for each iteration (DER penetration level) is a mixed integer second-order cone program (MISOCP). In addition to ACOPF, the dynamic coordination problem also includes several more constraints (including intertemporal constraints) to represent each of the DER models. These will be described in the next section. Note that adding such detailed DER models requires the use of integer variables, for e.g., to prevent simultaneous BS charging and discharging, or to switch between HP heating and cooling modes.

6.5 DER modeling

We perform a multiperiod optimization (MPO) that accounts for all the device-level constraints of individual DERs, including time-coupled (or intertemporal) state constraints for the BS, EV, and HP. The simulation timestep is Δt, the total simulation time is $\mathcal{T} = [0, T]$ and the planning horizon for the MPO is given by $\mathcal{H} = [t_H, t_H + (H-1)\Delta t] \subset \mathcal{T}$. We assume that the market-clearing timestep is also equal to Δt. Here, H is the length (number of timesteps) of the planning horizon. Note that throughout this paper, P_i denotes net active power injections (generation minus load). Thus, $P > 0$ implies net generation while net loads correspond to $P_i < 0$. We do not consider reactive power in this work.

6.5.1 PV model

The maximum PV generation output is determined by the forecasted, time-varying solar irradiance profile $\alpha^{PV}(t)$ along with its maximum rated capacity $\overline{P}_i^{PV}$, which can be curtailed if needed.

$$0 \leq P_i^{PV}(t) \leq \alpha^{PV}(t)\overline{P}_i^{PV} \tag{6.21}$$

The objective here is to minimize the amount of clean power that is curtailed.

$$f_i^{PV} = \xi_{pv} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (\alpha^{PV}(t)\overline{P}_i^{PV} - P_i^{PV}(t))^2 \tag{6.22}$$

Reactive power of PV is determined by variable power factor control:

$$-P_i^{PV} \tan(\cos^{-1}(\underline{pf})) \leq Q_i^{PV} \leq P_i^{PV} \tan(\cos^{-1}(\underline{pf}))$$

6.5.2 BS model

The state of charge (SOC) dynamics of the battery are:

$$SOC_i^{BS}(t+1) = (1 - \delta_{BS}^i)SOC_i^{BS}(t) - \frac{P_i^{BS}(t)\Delta t \eta_i^{BS}}{\overline{E}_i^{BS}} \tag{6.23}$$

$$\underline{P}_i^{BS} \leq P_i^{BS}(t) \leq \overline{P}_i^{BS} \tag{6.24}$$

$$\underline{SOC}_i^{BS} \leq SOC_i^{BS}(t) \leq \overline{SOC}_i^{BS} \tag{6.25}$$

$$SOC_i^{BS}(0) = SOC_i^{BS}(T) \tag{6.26}$$

where $\delta_i^{BS}, \eta_i^{BS}$ and $\overline{E}_i^{BS}$ are the BS self-discharge rate, round-trip efficiency, and maximum capacity, respectively. We also have a terminal constraint to ensure that the state of charge at the start and end of the simulation must be equal. Note that generally, $\underline{P}_i^{BS} = -\overline{P}_i^{BS}$ since we assume the BS has the same charging and discharging power capabilities. During

BS operation, the CMA also aims to minimize the cycling cost to avoid excessive charge and discharge cycles, which can degrade the battery's lifetime.

$$f_i^{BS}(P_i^{BS}) = \alpha_{cyc} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (P_i^{BS}(t+1) - P_i^{BS}(t))^2 \tag{6.27}$$

where we sum over all the timesteps in the planning horizon starting at time t_H.

We also have variable power factor control for battery reactive power. We need to reformulate this as an MIQP since P_i^{BS} can be either > 0 (discharging) or < 0 (charging).

$$-|P_i^{BS}|\tan(\cos^{-1}(\underline{pf})) \leq Q_i^{BS} \leq |P_i^{BS}|\tan(\cos^{-1}(\underline{pf})) \tag{6.28}$$

$$|P_i^{BS}| = P_i^{BS,+} + P_i^{BS,-}, \quad P_i^{BS} = P_i^{BS,+} - P_i^{BS,-} \tag{6.29}$$

$$0 \leq P_i^{BS,+} \leq z_i \overline{P}_i^{BS}, \quad 0 \leq P_i^{BS,-} \leq (1-z_i)\overline{P}_i^{BS} \tag{6.30}$$

$$z_i \in \{0,1\} \tag{6.31}$$

The additional binary variables z_i and binary constraints ensure that the BS does not simultaneously charge and discharge (i.e., only at most one of either $P_i^{BS,+}$ or $P_i^{BS,+=-}$ can be > 0 at any given time). Note that we don't need to reformulate like this for PV power factor since $P_{PV}^i > 0$ always.

6.5.3 EV model

The EV also has SOC constraints on its battery similar to the BS. In addition, we place restrictions on EV availability. We assume the EV is not present at the building or home during the period $[t_1, t_2]$, e.g. between 9am and 5pm when the owner is at work.

$$SOC_i^{EV}(t+1) = (1-\delta_{EV}^i)SOC_i^{EV}(t) - \frac{P_i^{EV}(t)\Delta t \eta_i^{EV}}{\overline{E}_i^{EV}} \tag{6.32}$$

$$\underline{P}_i^{EV} \leq P_i^{EV}(t) \leq \overline{P}_i^{EV} \tag{6.33}$$

$$\underline{SOC}_i^{EV} \leq SOC_i^{EV}(t) \leq \overline{SOC}_i^{EV} \tag{6.34}$$

$$SOC_i^{EV}(0) = SOC_i^{EV}(T) \tag{6.35}$$

$$P_i^{EV}(t) = 0 \; \forall \; t \in [t_1, t_2] \tag{6.36}$$

We can impose a similar cycling cost on the EV to extend its lifetime. We also add a tracking objective that the EV owner would like to achieve a certain desired SOC (SOC_i^*) by a specific time (t^*), say the owner needs the EV to be 90% charged by 9am before work. Thus, the EV objective function is

$$f_i^{EV} = \alpha_{cyc} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (P_i^{EV}(t+1) - P_i^{EV}(t))^2 + \; \xi_{ev}(SOC_i^{EV}(t^*) - SOC_i^{EV^*})^2 \tag{6.37}$$

We can also place additional constraints based on whether we enable only V1G (managed charging) or vehicle-to-grid V2G (where the EV can also inject power back into the grid).

Under V1G mode, the EV can only discharge while being driven/operated (disconnected from the charger) - it cannot discharge to the grid while connected to the charger.

$$P_i^{EV}(t) = P_i^{EV,+} - P_i^{EV,-}, |P_i^{EV}(t)| = P_i^{EV,+} + P_i^{EV,-} \tag{6.38}$$

$$\text{If only V1G allowed: } P_i^{EV}(t) \leq 0 \implies P_i^{EV,+} = 0 \tag{6.39}$$

$$\text{If bidirectional power flow allowed (can be +ve): } 0 \leq P_i^{EV,+} \leq z_i \overline{P}_i^{EV} \tag{6.40}$$

We can also have other bidirectional charging modes in addition to V2G, such as V2L and V2H. In addition to the tracking term in the objective, we can enforce the following stricter constraint to ensure the EV reaches the desired charge by the deadline time t_{dl}.

$$SOC_i^{EV}(t_{dl}) \geq SOC_i^{EV^*}$$

For simulations, we assumed that the homes could be installed with a 9.6 kW EV charger, which is the most common level 2 residential charging setup in the US. We also assume 70 kWh energy storage capacity, allowing for full charge in about 6-8 hours.

6.5.4 HVAC or heat pump model

The HVAC dynamics describe how the power drawn affects the indoor air temperature in the home or building. In this work, we consider the HVAC unit to be a heat pump (HP), which can serve as either a heating or cooling device depending on the ambient temperature. These thermostatically controlled loads (including heat pumps, HVAC units, and water heaters) can provide load flexibility by modifying their temperature setpoints [158]. If $T_i^{out}(t) > T_i^{in}(t)$, the temperature dynamics of the HP in cooling mode (i.e., when it acts as an air conditioner) are ([99]):

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)\left(T_i^{out}(t) + \rho_i P_i^{HP}(t)\right)$$

where $\theta_i = e^{\frac{-\Delta t}{R_i^{th} C_i^{th}}} \approx 1 - \frac{\Delta t}{R_i^{th}, C_i^{th}}, \rho_i = R_i^{th}\eta_i$ and $R_i^{th}, C_i^{th}, \eta_i$ are the equivalent thermal resistance, thermal capacitance, and coefficient of performance of the system, respectively. The temperature dynamics in heating mode, when $T_i^{out}(t) < T_i^{in}(t)$, are:

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)\left(T_i^{out}(t) - \rho_i P_i^{HP}(t)\right)$$

The HP operation is also subject to operational limits on power draw and indoor temperature:

$$\underline{P}_i^{HP} \leq P_i^{HP}(t) \leq 0 \tag{6.41}$$

$$\underline{T}_i^{in} \leq T_i^{in}(t) \leq \overline{T}_i^{in} \tag{6.42}$$

Note that here the lower limit is determined by the maximum power consumption rating of the HVAC unit, i.e. $\underline{P}_i^{HP} = -P_{rated,i}^{HP}$ since the HP always acts as a load.

Parameter	Description	Value	Unit
C	thermal capacitance	2	kWh/°C
R	thermal resistance	2	°C/kW
P_m	rated electrical power	5.6	kW
η	coefficient of performance	2.5	
θ_r	temperature set-point	22.5	°C
Δ	temperature deadband	0.3125	°C
θ_a	ambient temperature	32	°C

Table 6.1: Typical Parameter Values for a Residential HP unit

Note: We can also simplify the HP model using a single equation, by allowing HP to automatically switch between heating and cooling modes without explicitly using conditional statements:

$$\begin{aligned} T_i^{in}(t+1) &= \theta_i T_i^{in}(t) + (1-\theta_i)\left(T_i^{out}(t) + \rho_i P_i^{HP}(t)\right) \\ \text{where } & -P_i^{HP,r} \leq P_i^{HP}(t) \leq P_i^{HP,r} \end{aligned} \tag{6.43}$$

where $P_i^{HP,r}$ is the rated power of the HP. P_i^{HP} is negative when the HP is in heating mode and positive when it is in cooling mode. This HP model can also be used for water heating, but we will not consider that in this work.

For the objective function, the CMA would also like to track a desired temperature setpoint to maximize the thermal comfort of occupants (first term), and avoid excessive cycling (or hysteresis) between heating and cooling modes (second term).

$$f_i^{HP} = \xi_{ac} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (T_i^{in}(t) - T_i^{in^*})^2 + \alpha_{cyc} \sum_{t=t_H}^{t_H+(H-1)\Delta t} (P_i^{HP}(t+1) - P_i^{HP}(t))^2 \tag{6.44}$$

In addition to HVAC units, using a similar approach, we may also consider other types of thermostatically controlled loads (TCLs) such as water heaters (WH). We assume HPs don't inject or absorb any reactive power, i.e., $Q_i^{HP} = 0$.

The thermal parameters for a residential (single-family) HP unit are given by Table 6.1 [159]. Note that since we're aggregating all individual DER loads up to the primary feeder node level, we need to appropriately scale these parameters based on the number of homes/buildings connected to each bus, i.e., the total rated HP load capacity at each node.

6.5.5 Data centers

A possible extension is to also include small, modular data centers in distribution grids. Most modern data centers range between 1-50 MW, for those used for cryptocurrency mining or AI model training. Capacities < 1 MW are quite rare now - these would correspond to smaller HPC clusters rather than LLMs. The average size is around 20 MW. Thus, these could also be present at the medium-voltage level in distribution, depending on the size and peak load

rating of the feeder. For example, a 1 MW center for the IEEE-123 (peak load 3.6 MVA) or larger units for feeders like the 8500-node.

The research hypothesis here is: *Can we accommodate more data center load growth and with greater reliability and lower costs, using dynamic coordination?*. We can derive data center power consumption profiles using real-world data. Some potential sources include:

1. MIT Supercloud data set for Lincoln Labs High-Performance Computing cluster [160] accessible at https://dcc.mit.edu/data/

2. Google data center power traces from 2019 [161], accessible at https://github.com/google/cluster-data/blob/master/PowerData2019.md

Several prior works have studied the flexibilities of conventional data centers, but there is very little work looking at the load flexibility potential of newer data centers used for cryptocurrency mining and AI/ML model training. In order to simplify, we can start by treating data centers as inflexible loads with custom profiles and later extend to the flexible case (this would require good quality data for data center flexibility profiles).

6.6 Deterministic HCA

We first consider a deterministic case where we don't account for uncertainty and only consider a single scenario of inputs. One simplification is to only restrict ourselves to a specific time of day, which corresponds to the worst possible inputs. For example, if we're mainly focused on PV hosting capacity, we can consider just mid-day hours (e.g., 12-1pm) during peak PV output since the main concern is overvoltage due to excess generation [162]. However, while estimating HC for other DERs like EV, BS, or HP, we may also need to include other time periods since there are also risks of undervoltage (due to excess load) or line overloading.

6.6.1 Iterative approach

Initially, we implement an iterative approach where we gradually increase the DER penetration in small increments and run both the static and dynamic models until we hit power flow infeasibility. This allows us to keep track of maximum bus voltages and line loadings at each intermediate penetration level. Note that most papers based on iterative approaches focus on estimating the HC of only one type of DER at a time, e.g., maximizing PV only or HP only [154]. However, it is still relatively more computationally expensive and may not give as granular HC estimates due to the fixed step size increments at each iteration.

6.6.2 Optimization-based approach

Here, we formulate the HC estimation directly as a convex mixed integer optimization problem. We maximize the penetrations (i.e., the maximum upper capacity limits) of different types of DERs as decision variables, subject to power flow and other constraints (risk, uncertainty, etc.)

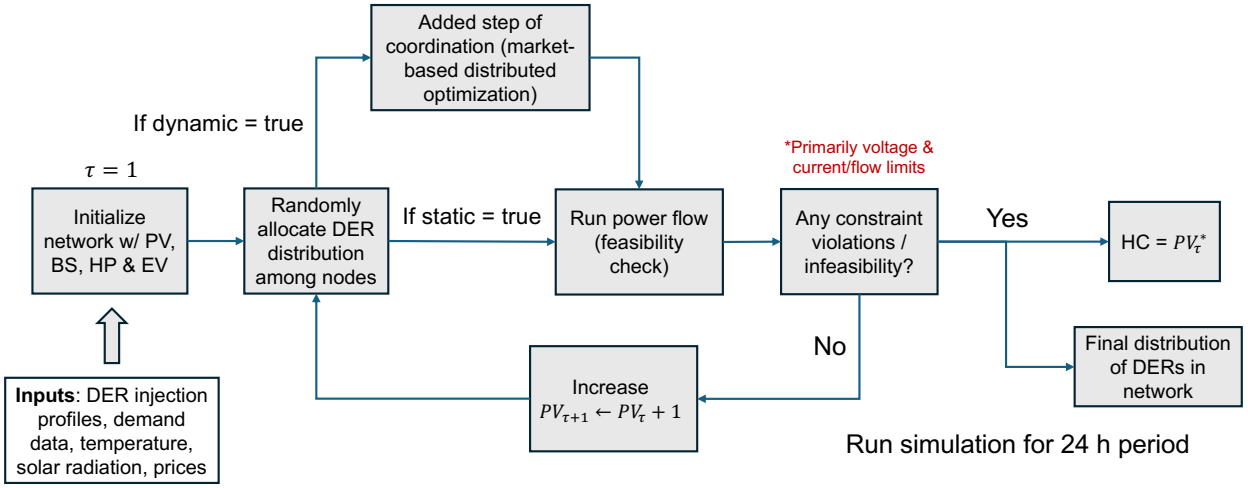

Figure 6.6: Deterministic iteration results.

$$\max_{\overline{P}_i^d, \overline{Q}_i^d, \mathbf{x}} \sum_{d \in \mathcal{D}, i \in \mathcal{N}} \overline{P}_i^d, \quad d \in \mathcal{D} = \{PV, BS, EV, HP\} \tag{6.45}$$

$$\text{s.t. } \underline{P}^d \leq \overline{P}_i^d \leq \overline{P}^d, \underline{Q}^d \leq \overline{Q}_i^d \leq \overline{Q}^d \tag{6.46}$$

$$\mathbf{x} \in \mathcal{X}, \ f(\mathbf{x}) = 0 \tag{6.47}$$

Here $\mathbf{x}$ denotes all the states and decision variables involved (e.g. nodal voltage magnitudes, line currents, SOC of BS and EV, temperature for HP, etc.) $\mathcal{X}$ denotes the feasible set determined by all inequality constraints (including power flow), along with the equality constraints f (including intertemporal or time-coupled constraints). Some potential advantages of this approach include:

1. This can be more computationally efficient and faster, especially when considering non-deterministic problems over longer time horizons.

2. It offers more flexibility, and it is easier to control which DER penetrations I allow to be varied in the simulation.

Eq. (6.45) is also in the form of a multiobjective optimization problem, so the optimal combination of different DER types likely won't be unique. Instead, we obtain a Pareto front of different DER capacities. This would also show correlations among different DER penetrations. Solving the optimization problem in Eq. (6.45) also solves for the optimal siting and sizing of DERs. Instead of randomly allocating DERs, given any desired (feasible) DER penetration level, this will optimally determine the nodal locations of different types of DERs, along with their associated sizes. The siting of new DER units depends on various factors such as network topology and parameters, existing loads and generators, etc. We can also modify the optimization problem to meet different objectives like maximizing HC, minimizing solar curtailment, or ensuring more uniform distributions of DERs over the feeder. This optimization-based approach will be studied further in Section 6.9.1.

6.7 Non-deterministic version

6.7.1 Sources of uncertainty or non-determinism

In the static case, we have stochasticity in the baseline power profiles of DERs (BS/EV/HP) that we assume for the simulation, along with inherent uncertainty in solar output, fixed load consumption, and prices. In the dynamic case, we have uncertainty in temperature (affecting HP load and flexibility), fixed load profiles, solar radiation, and prices. In addition, there is stochasticity in the time availability windows of EVs, i.e., EVs will differ in when they're grid-connected and available for V1G and/or V2G [163]. It is reasonable to assume that all these variables can be treated as roughly independent random variables. Thus, we can just randomly sample and generate each variable separately and then combine them to create each scenario. This input uncertainty captures both forecast uncertainty (i.e., epistemic uncertainty) as well as the inherent variability (i.e., aleatoric uncertainty) arising from factors like seasonal weather patterns [145]. Grid planners need to accurately account for uncertainty for more realistic, conservative HC estimates, which are generally lower than deterministic results

6.7.2 Simplest rudimentary non-deterministic approach

This is done by repeatedly running the iterative method from Section 6.6.1 for many scenarios of timeseries profiles, in order to obtain distributions of results across scenarios. Thus, we convert the deterministic iterative method to a stochastic iterative method using Monte Carlo sampling. The approach is summarized in Fig. 6.7 and consists of the following steps.

1. Generate many scenarios or realizations of all input timeseries profiles. In our case, we generated these by adding random Gaussian noise to certain baseline profiles.

2. We then run both the static and dynamic iterative HCA by considering each scenario at a time. Since the scenarios can all be run independently, we can parallelize these to significantly speed up the simulation..

3. Pick the lowest worst-case HC over all the scenarios (in order to be conservative) and also record other metrics like the mean and median. We can then plot the estimated HCs as distributions and obtain their kernel density estimates as well. Quantifying uncertainty in this manner also provides valuable information for grid planners, operators, and load-serving entities.

One approach to improve the computational efficiency of this manual sweep approach is to reduce the number of scenarios that need to be run. We can perform unsupervised clustering methods (e.g., k-means) to identify the most representative/characteristic scenarios and only run the simulation for these. It is crucial that this reduced scenario set must still include the most severe or extreme cases that will place more stress on the grid and make it challenging to satisfy constraints.

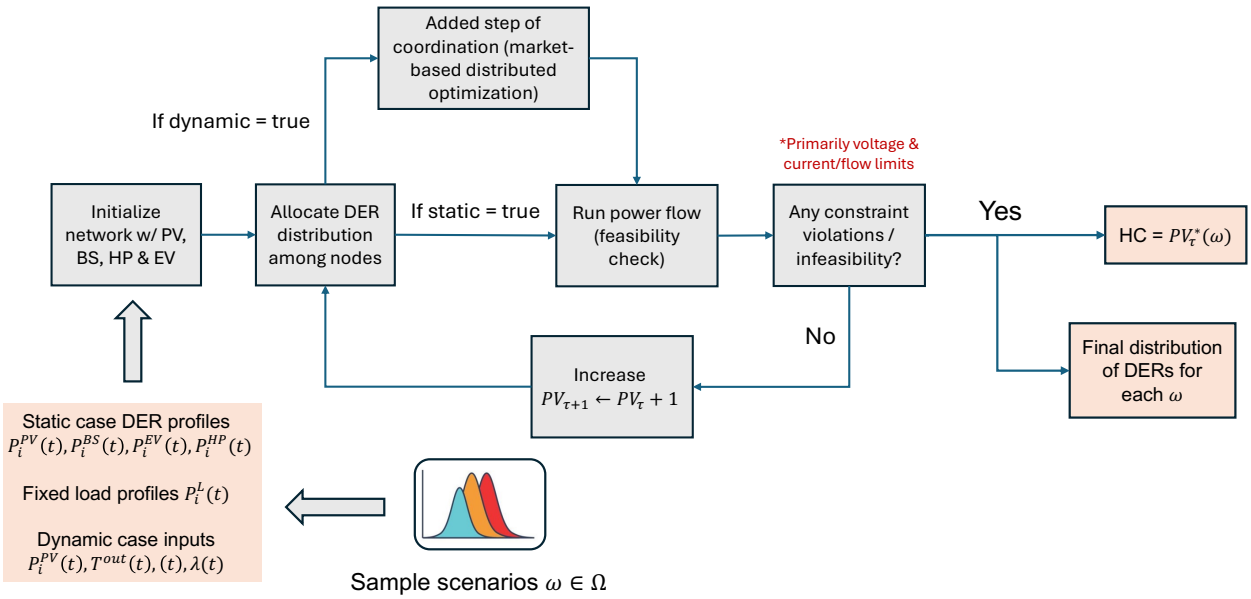

Figure 6.7: Stochastic iterative approach.

6.7.3 Stochastic programming (SP)

Here, we use a stochastic optimization approach [164] and generate the scenarios needed using Monte Carlo simulations. The Monte Carlo sampling method is to construct a finite set of possible scenarios (or random realizations), usually by assuming that the samples are independent and identically distributed.

6.7.4 Robust optimization (RO)

Here, we assume that the objective functions and constraints are not known exactly, but belong to a certain uncertainty set. We aim to find a solution that is feasible for all possible realizations of the uncertain parameters. The goal is to make decisions that are feasible no matter what the constraints turn out to be, and are optimal for the worst-case objective function. This is a more conservative approach than stochastic programming, which aims to optimize the expected value of the objective function. We can always assume, without loss of generality, that there is no uncertainty in the objective function. This is because we can always reformulate uncertainty in the objective function to constrain uncertainty by introducing additional constraints for the worst-case objective function value. The main challenge with RO is efficiently computing and representing the uncertainty set. We can use ellipsoidal uncertainty sets, polyhedral uncertainty sets, or more general conic uncertainty sets. The choice of the uncertainty set can have a significant impact on the computational complexity and tractability of the optimization problem. Furthermore, RO tends to be too conservative for many applications and may trade off more optimality than desired.

6.7.5 Information gap decision theory (IGDT)

This is a non-probabilistic approach that accounts for the worst possible scenarios [165]. It is based on the concept of information gap, which measures the lack of knowledge about the true values of uncertain parameters. The goal is to make decisions that are robust to the worst possible realization of the uncertain parameters. This approach is particularly useful when the probability distributions of the uncertain parameters are unknown or difficult to estimate.

6.7.6 Chance constraints

Chance constraints enforce constraints in a probabilistic manner. This approach is safer than stochastic programming since it allows the decision-maker to specify the probability of violating a constraint. It can also improve feasibility for problems containing constraints that are hard to strictly satisfy. The goal is to find a solution that satisfies the constraints with a high probability. Chance constraints can be formulated as follows:

$$\mathbb{P}(g(x,\omega) \leq 0) \geq 1 - \epsilon \tag{6.48}$$

This is a safer approach than SP and less conservative than RO, but chance constraints can be hard to compute efficiently, and we also need accurate probability distributions (which may not always be available). For these reasons, we decided that stochastic programming is the best approach for our HC application.

6.8 Two-stage stochastic program (2-SSP)

A two-stage stochastic program is a bilevel problem with the master problem being solved at the 1st decision stage and several sub-problems solved at the 2nd operational stage (one for each scenario or realization). The key challenge here is to ensure that our solutions for the 1st stage (decisions made before information on realized scenarios) are feasible even for the 2nd stage. This requires coupling between the two stages.

6.8.1 General formulation (infinite)

A standard two-stage stochastic programming problem can be formulated as follows. This formulation considers an infinite, continuous scenario set. The following is the master problem:

$$\min_{x \in X} \quad f^I(x) + \mathbb{E}_\omega[Q(x,\omega)] \tag{6.49}$$

$$\text{subject to} \quad Ax = b, \quad h(x) \leq 0 \tag{6.50}$$

$$x \geq 0 \tag{6.51}$$

where $Q(x,\omega)$ is the optimal value of the second-stage (or sub) problem:

$$Q(x,\omega) = \min_{y \in Y} \quad f^{II}(y,\omega) \tag{6.52}$$

$$\text{subject to} \quad T(\omega)x + W(\omega)y = d(\omega) \tag{6.53}$$

$$g(\omega, y) \leq 0, \quad y \geq 0 \tag{6.54}$$

Here:

- $x \in \mathbb{R}^{n_1}$ represents the first-stage decision variables
- $y \in \mathbb{R}^{n_2}$ represents the second-stage decision variables
- ω represents the random event from probability space $(\Omega, \mathcal{F}, P)$. Here, the sample space Ω is the set of all possible scenarios or states, i.e., all possible outcomes of the random variable ω. $\mathcal{F}$ is the sigma-algebra of events, i.e., the collection of all events for which we can assign probabilities, and P is the probability measure. For any event $A \in \mathcal{F}$, $0 \leq P(A) \leq 1$ is the probability that the event A occurs.

Note that this formulation with continuous probability spaces is an infinite optimization problem. Discretizing the distribution and the scenario set results in a tractable, finite optimization problem. Sampling a large number of diverse scenarios helps ensure that the finite case is still a good approximation for the infinite case. We can use Monte Carlo sampling to generate such scenarios.

Risk measures

A risk measure is a function that maps a random variable to a real number. In the above formulation, we used the mean (or expectation) of the 2nd stage cost as our risk measure. In order to be more risk-averse, we could instead use other risk measures like the conditional value at risk (CVaR) [166]. The value at risk (VaR) at confidence level α is defined as the α-quantile of the distribution of the random variable X. Suppose X is a loss we're aiming to minimize, i.e., larger values are worse outcomes. VaR gives the worst expected loss at a given confidence level. It is a threshold, with a probability (α = confidence level), the loss will not exceed VaR_α. Thus, the VaR is the smallest value z such that the CDF hits or exceeds α, i.e., $P(X \leq z) \geq \alpha$. The formula for VaR is given by:

$$VaR_\alpha(X) = \min\{z | F_X(z) \geq \alpha\}$$

where $F_X(z) = P(X \leq z)$ is the cumulative distribution function (CDF) of the random variable X. Note that using the VaR measure does not control loss scenarios exceeding the VaR cutoff. On the other hand, CVaR or Expected Shortfall is a risk assessment measure that quantifies the expected loss in the worst-case scenarios beyond a specified confidence level. Unlike Value at Risk (VaR), which estimates the maximum potential loss at a given confidence level, CVaR provides the average loss occurring beyond the VaR threshold, offering a more comprehensive view of tail risk. It also has better mathematical properties than VaR. The CVaR risk measure is given by:

$$CVaR_\alpha(X) = \mathbb{E}[X | X \geq VaR_\alpha(X)] = \frac{1}{\alpha} \int_0^\alpha VaR_\epsilon(X) d\epsilon$$

Equivalently, CVaR can also be expressed as the solution to the following optimization problem:

$$CVaR_\alpha(X) = \min_{t\in\mathbb{R}} \left\{ t + \frac{1}{1-\alpha}\mathbb{E}[(X-t)^+] \right\} \tag{6.55}$$

$$= \min_{t\in\mathbb{R}} \left\{ t + \frac{1}{1-\alpha}\mathbb{E}[\max\{X-t, 0\}] \right\} \tag{6.56}$$

This can, in turn, be represented by the following linear program, with additional constraints to linearize the max function:

$$CVaR_\alpha(X) = \min_{t\in\mathbb{R}, u\in\mathbb{R}^n} \left\{ t + \frac{1}{1-\alpha}\mathbb{E}(u(\omega)) \right\} \tag{6.57}$$

$$\text{s.t. } u(\omega) \geq X - t,\ u(\omega) \geq 0,\ \forall \omega \in \Omega \tag{6.58}$$

where $X(\omega), u(\omega)$ are the random variables and auxiliary variables, respectively. In the case of a discrete distribution, we can replace the expectation with a summation over all scenarios weighted by their probabilities.

$$CVaR_\alpha(X) = \min_{t\in\mathbb{R}, u_i\in\mathbb{R}^n} \left\{ t + \frac{1}{1-\alpha}\sum_i p_i u_i \right\} \tag{6.59}$$

$$\text{s.t. } u_i \geq X_i - t,\ u_i \geq 0,\ \forall i \in \Omega \tag{6.60}$$

6.8.2 Discrete Distribution Case (finite)

If the random variable ω has a discrete distribution with finite support $\{\omega_1, \omega_2, \ldots, \omega_N\}$ with corresponding probabilities $p_1, p_2, \ldots, p_N$, then the problem can be reformulated as a large-scale deterministic nonlinear program. This is also known as the discretized 2-SSP or sample average approximation approach. Note that this optimization problem solves for separate sets of 2nd-stage decision variables for each scenario, and the goal is to obtain 1st-stage decision variables (or design decisions) that are feasible for all scenarios.

$$\min_{x, y_1, \ldots, y_K} \quad f^I(x) + \sum_{k=1}^{K} p_n f^{II}(y_n, \omega_n) \tag{6.61}$$

$$\text{subject to} \quad Ax = b \tag{6.62}$$

$$h(x) \leq 0 \tag{6.63}$$

$$T_k x + W_k y_k = d_k, \quad k = 1, 2, \ldots, K \tag{6.64}$$

$$g(\omega_k, y_k) \leq 0, \quad k = 1, 2, \ldots, K \tag{6.65}$$

$$x \geq 0, \quad y_k \geq 0, \quad k = 1, 2, \ldots, K \tag{6.66}$$

$$\tag{6.67}$$

where for scenario n:

$$d_k = d(\omega_k) \tag{6.68}$$
$$T_k = T(\omega_k) \tag{6.69}$$
$$W_k = W(\omega_k) \tag{6.70}$$

We will use a discretized 2-stage SP for our HCA problem. Here, each event ω_k corresponds to a scenario or realization. Generally, we assume that all scenarios are equally likely, i.e., $p_k = \frac{1}{K}$. However, it is also straightforward to use scenario-specific probabilities if we have access to the accurate probability distribution of the inputs over possible scenarios. We can either use Julia packages like `StochasticPrograms.jl`, `SDDP.jl` [167] or `InfiniteOpt.jl` [168] to solve this problem. However, StochasticPrograms.jl is not actively maintained anymore (last updated in 2022), and SDDP.jl is primarily designed for linear programs - it is hard to extend to mixed integer nonlinear problems. As a result, we decided to implement the 2-SSP ourselves from scratch natively in JuMP[1]. While this approach has worked reasonably well for our moderately-sized problems, it probably won't be practical for larger-scale problems and networks. For problems with many more decision variables and constraints, this standard implementation would be very computationally expensive and not practical since it does not use specialized solution algorithms (like Benders decomposition or SDDP described in Section 6.10) that can exploit the special multi-stage and bilevel temporal structure of the problem.

6.9 2-stage SP for HCA

The HCA can be written as a 2-stage stochastic optimization problem with fixed recourse. In the 1st stage (upper level master problem), we choose the overall DER penetrations before scenario realization, and in the 2nd stage (lower level sub-problems, we run power flow simulations for all realized scenarios. We can then solve this problem efficiently using tools like Benders decomposition [164]. If we also include optimal siting in the optimization, the 1st stage (planning) decisions would be the DER penetrations along with their locations and capacities at nodes. The lower level 2nd stage problem is an AC OPF that checks feasibility for each scenario $\omega \in \Omega$. Thus, the upper problem is concerned with design, planning, and investment - the grid planner has to make these decisions before seeing the input uncertainty. The lower problem is concerned with actual grid operation to set the power flow solutions and DER dispatch. These can only be decided after the uncertain input scenarios ω are realized.

6.9.1 First stage planning problem (design stage)

Here we aim to maximize the capacities of PV, EV, and HP, for a given BS level ($\overline{P}^{BS} = \sum_i \overline{P}_i^{BS}$). We consider the *total* maximum BS capacity (both rated power and energy storage capacity) as exogenous variables or inputs (specified by the planner, utility, or grid operator) since one can essentially arbitrarily increase BS levels to as high as desired. This is because

[1] https://jump.dev/JuMP.jl/stable/tutorials/applications/two_stage_stochastic/

we can always control whether to charge, discharge, or just turn off (disconnect) batteries as needed, depending on grid conditions. Thus, we don't define the notion of HC for batteries. Higher BS penetrations will allow us to increase dynamic HC further, but come at an added cost. One of the expected results is to show that dynamic coordination will allow us to achieve hosting capacities similar to (or more than) static values, but with lower storage requirements. While the total BS penetration across the entire feeder is given, its distribution amongst the nodes is determined optimally using our program.

The decision variables represent the penetration levels of each DER type at each node. In the case of PV and BS, x_i^{PV} and x_i^{BS} refer to the penetration of PV capacity and BS relative to the total baseline/nominal load at that node i (as specified by the IEEE datasheet). However, in the case of EVs and HPs, the decision variables x_i^{EV}, x_i^{HP} represent the proportion (%) of homes/buildings (connected to primary feeder node i) that have been electrified with EVs and HPs, respectively. Formulating all the decision variables as proportions (between 0 and 1) makes the first 4 terms in the objective function Eq. (6.71) to be in the same units (dimensionless) and have similar magnitudes. The last term f_{cost} is in units of \$, so we would need to transform this to be comparable in magnitude. The objective function is to maximize the total DER capacity, while minimizing fixed costs[2]. These primarily include capital costs such as for DER installations or for any grid infrastructure upgrades or retrofits needed. We also introduce binary variables indicating whether a DER is present at a particular node i or not, $z_i^{PV}, z_i^{BS}, z_i^{EV}, z_i^{HP} \in \{0, 1\}$. The first stage problem is formulated as:

$$\max_{x_i, z_i} f^I = \sum_i z_i^{PV} x_i^{PV} + z_i^{BS} x_i^{BS} + z_i^{EV} x_i^{EV} + z_i^{HP} x_i^{HP} - f_{cost}(x_i^{PV}, x_i^{EV}, x_i^{HP}) \tag{6.71}$$

$$\text{s.t. } \underline{x}^{PV} \leq x_i^{PV} \leq \overline{x}^{PV}, \quad \underline{x}^{BS} \leq x_i^{BS} \leq \overline{x}^{BS}, \quad \underline{x}^{EV} \leq x_i^{EV} \leq \overline{x}^{EV}, \quad \underline{x}^{HP} \leq x_i^{HP} \leq \overline{x}^{HP} \tag{6.72}$$

$$\sum_i z_i^{PV} \leq N^{PV}, \quad \sum_i z_i^{BS} \leq N^{BS}, \quad \sum_i z_i^{EV} \leq N^{EV}, \quad \sum_i z_i^{HP} \leq N^{HP}, \quad z_i \in \{0, 1\} \tag{6.73}$$

$$\sum_i x_i^{BS} z_i^{BS} \overline{P^L}_{nom,i} \leq \overline{x}^{BS} \overline{P^L}_{nom} \tag{6.74}$$

$$\tag{6.75}$$

where the constraints are upper and lower limits on the penetrations of each DER type at each node, and the total number of DERs of each type (across the entire network) is limited to N^{PV}, N^{EV}, N^{HP}. The last constraint Eq. (6.74) restricts the maximum BS capacity to the level pre-determined by the planner. The right-hand side of this constraint sets the maximum battery power capacity [kW] based on the overall BS penetration level (%) and the total baseline feeder load $\overline{P^L}_{nom} = \sum_i \overline{P^L}_{nom,i}$. We can then run a sensitivity analysis to assess the impacts of storage capacity/availability on HC. Note that this problem may have multiple globally optimal solutions due to the presence of integer variables and constraints. The 1st stage decision variables are $\mathbf{x} = \{\mathbf{x}_i\} = \{x_i^{PV}, x_i^{BS}, x_i^{EV}, x_i^{HP}, z_i^{PV}, z_i^{BS}, z_i^{EV}, z_i^{HP}\}$. x represents the continuous variables where y represents the binary variables.

[2] nrel.gov/solar/market-research-analysis/distribution-grid-integration-unit-cost-database.html

Note that bilinear products of decision variables ($x_i z_i$) are nonconvex, but state-of-the-art solvers like Gurobi can still solve such problems to global optimality, although it's computationally quite expensive. However, if we don't have access to such commercial solvers or if computational requirements are a concern, we can simplify by removing the z_i binary variables. But note that the problem is still mixed integer due to the binary variables needed for the BS and HP modeling. The revised, simplified first-stage problem is given by:

$$\max_{x_i} f^I = \sum_i x_i^{PV} + x_i^{BS} + x_i^{EV} + x_i^{HP} - f_{cost}(x_i^{PV}, x_i^{EV}, x_i^{HP}) \tag{6.76}$$

$$\text{s.t. } \underline{x}^{PV} \le x_i^{PV} \le \overline{x}^{PV}, \quad \underline{x}^{BS} \le x_i^{BS} \le \overline{x}^{BS}, \quad \underline{x}^{EV} \le x_i^{EV} \le \overline{x}^{EV}, \quad \underline{x}^{HP} \le x_i^{HP} \le \overline{x}^{HP}$$

$$\sum_i x_i^{BS} \overline{PL}_{nom,i} \le \overline{x}^{BS} \overline{PL}_{nom}$$

The upper and lower bounds on the DER penetrations are given by:

$$0 \le x_i^{PV}, x_i^{BS} \text{ (no upper limit since we can have PV penetration more than 100\%)} \tag{6.77}$$

$$0 \le x_i^{HP}, x_i^{EV} \le 1 \text{ (max when 100\% of homes are electrified)} \tag{6.78}$$

We can also consider a case where we explicitly fix the penetrations of certain DERs for sensitivity studies, by using equality constraints instead of inequalities:

$$\sum_i x_i^{BS} \overline{PL}_{nom,i} = \overline{x}^{BS} \overline{PL}_{nom}, \quad \sum_i x_i^{HP} n_i^h P_r^{HP} = \overline{x}^{HP} \sum_i n_i^h P_r^{HP}, \quad \sum_i x_i^{EV} n_i^h P_r^{EV} = \overline{x}^{EV} \sum_i n_i^h P_r^{EV}$$

Comment: Framing 1st stage as multiobjective optimizaton

The 1st stage problem could also be viewed from the lens of a multiobjective optimization problem suitable for Pareto analysis. The objective function shows a tradeoff between maximizing benefits (from PV, HP, and EV installations) versus minimizing the total system costs. A Pareto analysis would identify the set of non-dominated solutions (i.e., on the Pareto frontier) where you cannot improve one objective without worsening another. Each point would represent a different optimal allocation of DERs that gives a distinct balance between benefits and costs. In addition, if we include the option of installing some conventional fossil fuel sites (e.g., diesel generators or gas microturbines), we would also have a tradeoff with a 3rd term minimizing carbon (or GHG) emissions. This would result in the following multiobjective cost f^I for stage 1:

$$\max_{x_i, z_i} f^I = \sum_i z_i^{PV} x_i^{PV} + z_i^{BS} x_i^{BS} + z_i^{EV} x_i^{EV} + z_i^{HP} x_i^{HP} - f_{cost}(x_i^{PV}, x_i^{BS}, x_i^{EV}, x_i^{HP}) - f_{ghg}(x_i^{DG})$$

6.9.2 Feeder configuration

For now, we assume that the network topology is known and fixed, and the distribution of houses or buildings (and their baseline maximum loads) over the feeder is also fixed. Later

on, we can extend this to run the HCA with multiple possible feeder configurations. Thus, the number of houses at each node i is given by n_i^h. For simplicity, we assume the same standardized sizes of EV and HP units for each home. P_r^{EV} and E_r^{EV} denote the common rated power and energy storage capacities of EVs, and P_r^{HP} is the standard rated HP load. Thus, the total actual capacities of EVs and HP at each node are given by:

$$\overline{P}_i^{HP} = x_i^{HP} n_i^h P_r^{HP}, \quad \overline{P}_i^{EV} = x_i^{EV} n_i^h P_r^{EV}, \quad \overline{E}_i^{EV} = x_i^{EV} n_i^h E_r^{EV} \tag{6.79}$$

$$\overline{P}_i^{BS} = x_i^{BS} \overline{P^L}_{nom,i} \quad \overline{E}_i^{BS} = \overline{P}_i^{BS} E_r^{BS} \tag{6.80}$$

Similarly, given the total maximum baseline nodal load $\overline{P^L}_{nom,i}$ specified by the IEEE data sheet, the nodal PV capacity at node i in kW is:

$$\overline{P}_i^{PV} = x_i^{PV} \overline{P^L}_{nom,i}$$

Thus, the 1st stage decision variables $\mathbf{x}$ couple the two stages since they set the maximum DER capacities at each node, which in turn determine constraints in the 2nd stage problem.

6.9.3 Second stage problem (operation stage)

The 2nd-stage formulation is different for the static and dynamic cases.

Static case

In the static case, we assume that all the DER injections are inflexible. Thus, we only solve a feasibility problem to check if the given DER injections satisfy all the power flow constraints from Section 6.3. In this case, the scenarios consist of different profiles for inflexible DER timeseries profiles, including HP load consumption, EV/BS charge-discharge cycles, fixed loads, and PV radiation. Thus, there are a total of 5 random variables, and we can assume these to be independent, but not necessarily identically distributed. There are two possible approaches to generating these scenarios:

- We can estimate probability distributions (PDFs) for each of the variables from real datasets and then sample from these in a Monte Carlo approach.
- Or we can directly generate multiple scenarios or realizations by randomly perturbing or adding noise to some baseline or nominal profiles derived from actual historical data. This is the approach we adopted for this study.

Thus, the 2-stage static SP is a recourse problem with only a feasibility check in the 2nd stage. The feasibility of first-stage decisions under uncertainty is assessed in the second stage. The 2nd stage doesn't have an objective function $f^{II} = 0$ and the decision variables are $\mathbf{y} = \{P_i(t), Q_i(t), P_{ij}(t), Q_{ij}(t), l_{ij}(t), v_i(t)\}\ \forall t \in \mathcal{T}$, with a separate set of variables for each scenario $\omega_k \in \Omega$. Note that the 2nd stage is a multiperiod optimization problem. The overall problem is then given by:

$$\min_{\mathbf{x},\mathbf{y}_1,\ldots\mathbf{y}_K} \quad f^I(\mathbf{x}) \tag{6.81}$$

$$\text{s.t.: Constraints on } \mathbf{x} \text{ in Eq. (6.76)} \tag{6.82}$$

$$h_{pf}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) = 0, \quad g_{pf}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) \leq 0, \quad k = 1, 2, \ldots K \tag{6.83}$$

$$P_i(\omega_k, t) = P_i^{PV}(\omega_k, t) + P_i^{EV}(\omega_k, t) + P_i^{BS}(\omega_k, t) - P_i^L(\omega_k, t) \tag{6.84}$$

$$Q_i(\omega_k, t) = -Q_i^L(\omega_k, t) \; \forall t \in \mathcal{T} \tag{6.85}$$

where h_{pf}, g_{pf} are the AC power flow equality and equality constraints, respectively, from Section 6.3, and K is the total number of sampled scenarios. ω_k denotes the data for each scenario. In case we allow solar PV curtailment, we would have an additional decision variable for the PV power injection P_i^{PV}. The coupling between the stages is through the maximum capacities for each of the DERs, which are determined in the 1st stage (see Section 6.9.2) and utilized in the 2nd. As an example for PV and EV:

$$P_i^{PV}(\omega_k) = \alpha^{PV}\overline{P}_i^{PV} = \alpha^{PV} x_i^{PV} \overline{P^L}_{nom,i} \tag{6.86}$$

$$P_i^{EV}(\omega_k) = \alpha_i^{EV}\overline{P}_i^{EV} = \alpha^{EV} x_i^{EV} n_i^h P_r^{EV} \tag{6.87}$$

where α^{PV}, α^{EV} are the stochastic baseline profiles for the PV and EV power injections, respectively. Again, these are derived from actual data and assumed to be inflexible.

Dynamic case

During the second stage, we solve the multiperiod optimal power flow problem that determines the dispatch after the realization of the uncertain scenarios. Solve the ACOPF problem (from Section 6.4.3 with objective from Section 6.4.4 and with all DER models/constraints from Section 6.5) to find feasible solutions for each scenario $\omega \in \Omega$. Each scenario introduces uncertainty in the following:

$$P_i^{PV}(\omega) = \overline{P}_i^{PV}\alpha(\omega) \quad \text{(without curtailment)} \tag{6.88}$$

$$P_i^{PV}(\omega) \leq \overline{P}_i^{PV}\alpha(\omega) \quad \text{(without curtailment)} \tag{6.89}$$

$$T_i^{out} = T_i^{out}(\omega) \quad \text{(assume common outdoor temp for all nodes)} \tag{6.90}$$

$$\lambda = \lambda(\omega), \quad P_i^L = P_i^L(\omega) \tag{6.91}$$

Here $\alpha(\omega)$ is the stochastic PV output based on solar irradiation in scenario ω. $T(\omega)$ and $\lambda(\omega)$ are the external temperature profile and LMP, respectively. These three variables depend on the scenario but are the same for all nodes. The stochastic fixed load profiles $P_i^L(\omega)$ differ for each node. The 2nd stage decision variables are The 2nd stage decision variables are $\mathbf{y} = \{P_i(t), Q_i(t), P_{ij}(t), Q_{ij}(t), l_{ij}(t), v_i(t)\}$ along with the additional DER-specific injection variables $\{P_i^{PV}, P_i^{BS}(t), P_i^{EV}(t), P_i^{HP}(t), Q_i^{PV}(t), Q_i^{BS}(t), Q^{EV}(t)_i\}$ and the controllable state variables $\{T_i^{in}(t), SOC_i^{BS}(t), SOC_i^{EV}(t)\}$. The 2nd stage objective function is the ACOPF cost function $f^{II} = f^{opf}$. The coupling between the two stages is through the maximum

capacities for each of the DERs (see Section 6.9.2), which are determined in the 1st stage and set the constraints for stage 2.

We can first write this 2-stage SP in a standard bilevel form to illustrate. The separable objective function is split into (i) a deterministic term representing decisions at the design stage and (ii) the expectation of a stochastic term which depends on the realization of uncertain parameters at the operation stage [169]. The 1st stage problem is the master problem, and the 2nd stage problem is the subproblem. In the continuous case, the master problem is given by:

$$\min_{\mathbf{x}} \quad f^I(\mathbf{x}) + \mathbb{E}_\omega[V(\mathbf{x}, \omega)] \tag{6.92}$$

$$\text{s.t.: Constraints on } \mathbf{x} \text{ in Eq. (6.76)} \tag{6.93}$$

where the subproblem is given by:

$$V(\mathbf{x}, \omega) = \min_{\mathbf{y}(\omega)} \quad f^{II}(\mathbf{x}, \mathbf{y}(\omega), \lambda) \tag{6.94}$$

$$\text{s.t.: } h_{pf}(\omega, \mathbf{x}, \mathbf{y}(\omega), t) = 0, \quad g_{pf}(\omega, \mathbf{x}, \mathbf{y}(\omega), t) \leq 0 \tag{6.95}$$

$$P_i(\omega, t) = P_i^{PV}(\omega, t) + P_i^{EV}(\omega, t) + P_i^{BS}(\omega, t) - P_i^{L}(\omega, t) \tag{6.96}$$

$$Q_i(\omega, t) = Q_i^{PV}(\omega, t) + Q_i^{EV}(\omega, t) + Q_i^{BS}(\omega, t) - Q_i^{L}(\omega, t) \tag{6.97}$$

$$h_{der}(\omega, \mathbf{x}, \mathbf{y}(\omega), t) = 0, \quad g_{der}(\omega, \mathbf{x}, \mathbf{y}(\omega), t) \leq 0 \tag{6.98}$$

In the discrete case, if we use the sample average approximation (SAA) approach, we can replace the expectation with a sample average over the scenarios. The master problem is then given by:

$$\min_{\mathbf{x}} \quad f^I(\mathbf{x}) + \sum_k p_k V(\mathbf{x}, \omega_k) \tag{6.99}$$

$$\text{s.t.: Constraints on } \mathbf{x} \text{ in Eq. (6.76)} \tag{6.100}$$

where p_k is the probability of scenario ω_k. The subproblem is given by:

$$V(\mathbf{x}, \omega_k) = \min_{\mathbf{y}_k} \quad f^{II}(\mathbf{x}, \mathbf{y}, \lambda_k) \tag{6.101}$$

$$\text{s.t.: } h_{pf}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) = 0, \quad g_{pf}(\omega, \mathbf{x}, \mathbf{y}_k, t) \leq 0, \quad k = 1, 2, \dots K \tag{6.102}$$

$$P_i(\omega_k, t) = P_i^{PV}(\omega_k, t) + P_i^{EV}(\omega_k, t) + P_i^{BS}(\omega_k, t) - P_i^{L}(\omega_k, t) \tag{6.103}$$

$$Q_i(\omega_k, t) = Q_i^{PV}(\omega_k, t) + Q_i^{EV}(\omega_k, t) + Q_i^{BS}(\omega_k, t) - Q_i^{L}(\omega_k, t) \tag{6.104}$$

$$h_{der}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) = 0, \quad g_{der}(\omega, \mathbf{x}, \mathbf{y}_k, t) \leq 0, \quad k = 1, 2, \dots K \tag{6.105}$$

Note that the subproblem is multiperiod and at each timestep, it is repeatedly solved for a finite time horizon into the future $\mathcal{H}$, in order to cover the whole simulation period. By combining the two stages into a single large-scale optimization problem, we obtain a single-level deterministic equivalent for the bilevel stochastic program. The overall optimization problem is given by:

$$\min_{\mathbf{x},\mathbf{y}_1,\ldots\mathbf{y}_K} \quad f^I(\mathbf{x}) + \sum_{k=1}^{K} p_k f^{II}(\mathbf{x}, \mathbf{y}_k, \lambda) \tag{6.106}$$

$$\text{s.t.: Constraints on } \mathbf{x} \text{ in Eq. (6.76)} \tag{6.107}$$

$$h_{pf}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) = 0, \quad g_{pf}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) \leq 0, \quad k = 1, 2, \ldots K \tag{6.108}$$

$$P_i(\omega_k, t) = P_i^{PV}(\omega_k, t) + P_i^{EV}(\omega_k, t) + P_i^{BS}(\omega_k, t) - P_i^{L}(\omega_k, t) \tag{6.109}$$

$$Q_i(\omega_k, t) = Q_i^{PV}(\omega_k, t) + Q_i^{EV}(\omega_k, t) + Q_i^{BS}(\omega_k, t) - Q_(^{L}\omega_k, t) \tag{6.110}$$

$$h_{der}(\omega_k, \mathbf{x}, \mathbf{y}_k, t) = 0, \quad g_{der}(\omega, \mathbf{x}_k, \mathbf{y}, t) \leq 0, \quad k = 1, 2, \ldots K \tag{6.111}$$

where h_{der}, g_{der} are the equality and equality constraints for the DER models and state dynamics, respectively, from Section 6.5. Note that the subproblem decision variables $\mathbf{y}_k$ are solved separately for each scenario ω_k and the solution is used to update the master problem. The master problem is then solved iteratively until convergence.

Enforcing DER model constraints

An important consideration for the stochastic dynamic case is that we enforce DER model state constraints (equalities and inequalities) only at buses or nodes where they're actually present. Such conditional constraints are non-convex but can be implemented as a mixed-integer program using big-M constraints, and then solved using solvers like Gurobi. As an example, suppose we want to enforce the equality constraint for the BS SOC update only at nodes i where BS is present. Thus, we enforce the constraint only when $x_i^{BS} \neq 0$ or, equivalently,y when $x_i^{BS} > 0$. Then the big-M constraints to enforce this conditional constraint would be:

$$SOC_i^{BS}(t+1) = h_{bs}(SOC_i^{BS}(t)) \text{ if } x_i^{BS} > 0 \tag{6.112}$$

$$SOC_i^{BS}(t+1) \geq h_{bs}(SOC_i^{BS}(t)) - M\left(1 - \frac{x_i^{BS}}{\epsilon}\right) \tag{6.113}$$

$$SOC_i^{BS}(t+1) \leq h_{bs}(SOC_i^{BS}(t)) + M\left(1 - \frac{x_i^{BS}}{\epsilon}\right) \tag{6.114}$$

Similarly, for the bound inequalities:

$$\underline{SOC}_i^{BS} \leq SOC_i^{BS}(t+1) \leq \overline{SOC}_i^{BS} \text{ if } x_i^{BS} > 0 \tag{6.115}$$

$$SOC_i^{BS}(t+1) \geq \underline{SOC}_i^{BS} - M\left(1 - \frac{x_i^{BS}}{\epsilon}\right) \tag{6.116}$$

$$SOC_i^{BS}(t+1) \leq \overline{SOC}_i^{BS} + M\left(1 - \frac{x_i^{BS}}{\epsilon}\right) \tag{6.117}$$

Furthermore, we can force the SOC variables to be zero at nodes where BS is not present:

$$SOC_i^{BS}(t+1) = 0 \quad \text{if } x_i^{BS} = 0 \tag{6.118}$$

$$\text{In Big-M form: } SOC_i^{BS}(t) \leq M\frac{x_i^{BS}}{\epsilon} \tag{6.119}$$

Where M is a large constant and ϵ is a small positive number. The value of M should be chosen large enough to ensure that the constraints are not violated, but not too large to avoid numerical instability, e.g., $M = 1000$. The variable ϵ is used to avoid division by zero and can be set to a small value like 10^{-6}. Importantly, ϵ should be much smaller than the usual/expected values of x_i^{BS} to ensure that the constraints are effectively enforced and always feasible. The same approach can be used for other DERs like EVs and HPs, where we can enforce the state constraints only at nodes where they are present.

Another option is to use binary variables to indicate whether a DER is present at a node or not. For example, we can introduce binary variables $z_i^{BS,ind}$ to indicate whether BS is present at node i or not, i.e.:

$$z_i^{BS,ind} = \begin{cases} 1 & \text{if } x_i^{BS} > 0 \\ 0 & \text{if } x_i^{BS} = 0 \end{cases} \tag{6.120}$$

These constraints can be implemented using:

$$x_i^{BS} \leq z_i^{BS,ind}, \quad x_i^{BS} \geq \epsilon z_i^{BS,ind} \tag{6.121}$$

Then we can enforce the constraints only when $z_i^{BS} = 1$. This approach is more computationally expensive but allows for more flexibility in modeling the problem. Using such binary variables also allows us to automatically set states to zero for nodes with no DERs, without needing to explicitly set them to zero. This is useful for the BS SOC constraints, where we can set the SOC to zero at nodes with no BS, and for the HP, where the temperature would always be zero. The constraints would then be:

$$z_i^{BS}\underline{SOC}_i^{BS} \leq SOC_i^{BS}(t) \leq z_i^{BS}\overline{SOC}_i^{BS}, \quad z_i^{HP}\underline{T}_i^{in} \leq T_i^{in}(t) \leq z_i^{HP}\overline{T}_i^{in} \tag{6.122}$$

$$T_i^{in}(t+1) = \theta_i T_i^{in}(t) + (1-\theta_i)\left(z_i^{HP} T_i^{out}(t) + \rho_i P_i^{HP}(t)\right) \tag{6.123}$$

$$\overline{E}_i^{BS} SOC_i^{BS}(t+1) = (1-\delta_{BS}^i) SOC_i^{BS}(t)\overline{E}_i^{BS} - P_i^{BS}(t)\Delta t \eta_i^{BS} \tag{6.124}$$

And we know that $P_i^{HP}, P_i^{BS}, \overline{E}_i^{BS}$ are going to be zero at nodes with no DERs.

Stochastic programs also differ depending on whether we discretize the scenario set by sampling only a few from a known distribution (finite) or if we consider the continuous distribution directly while solving the optimization (infinite). In the finite case, we can use Monte Carlo sampling to generate scenarios. In the infinite case, we can use scenario tree generation methods like L-shaped decomposition or progressive hedging. For our study, we only consider the finite case.

6.10 Solving 2-stage stochastic problems

So far, we've considered solving the 2-SSP by converting it to a single large-scale deterministic equivalent problem. However, this approach can be prohibitively expensive and slow for larger problems and often limits the number of scenarios that can be considered. Thus, there are some specialized decomposition techniques to solve large-scale stochastic optimization

problems that can reduce the computational burden and accelerate the solution process. These often involve an iterative approach due to the coupling between the 1st and 2nd stages. We can first solve the upper-level problem and then use these decision variables to solve the lower-level problem. In case the 2nd stage problem is infeasible, we add extra constraints (feasibility cuts) to the 1st stage problem and repeat the process until both stages are feasible. Some works also add these feasibility cuts as chance constraints rather than strict constraints, which helps the algorithm converge faster.

6.10.1 Benders-decomposition based solution algorithms

Benders decomposition is one of the most commonly used approaches to solve two-stage stochastic programming problems [170]. It is a general decomposition approach for large-scale mixed-integer and continuous optimization problems. It decomposes the problem into a master problem and subproblems, iteratively solving them to find the optimal solution. The master problem involves the first-stage decision variables (which are the variables we're mainly interested in solving for), while the subproblems involve the second-stage decision variables for each scenario. The algorithm iterates between solving the master problem and the subproblems, adding Benders cuts to the master problem (based on the dual solutions of the subproblems) to refine the solutions and ensure both feasibility and optimality. The general steps of Benders decomposition are as follows:

1. **Initialization**: Start with an initial guess for the first-stage decision variables.
2. **Subproblem Solution**: For each scenario, solve the second-stage subproblem to obtain the optimal second-stage decision variables and the corresponding dual variables.
3. **Benders Cuts**: Use the dual variables to generate Benders cuts, which are added to the master problem.
4. **Master Problem Solution**: Solve the master problem with the added Benders cuts to update the first-stage decision variables.
5. **Convergence Check**: Check for convergence. If the solution has converged, stop. Otherwise, return to step 2.

The process continues until the solution converges, ensuring that the first-stage decision variables are optimal and feasible for all scenarios. In the context of the HCA problem, Benders decomposition can be used to iteratively solve the first-stage planning problem and the second-stage feasibility problem, ensuring that the optimal DER penetrations are feasible for all scenarios. The L-shaped method is a specific application of Benders for 2-stage stochastic **linear** (generally continuous) programs [171]. The name "L-shaped" comes from the geometric structure of the feasible region in the Benders decomposition iterations. In the first stage, you decide "here-and-now" decisions. In the second stage, the subproblem optimizes "recourse" actions based on different stochastic realizations.

The subproblems in the L-shaped method are normally linear programs, which allow the derivation of Benders cuts. If the second-stage problem contains integer variables, the

dual function is no longer convex, which breaks the standard Benders decomposition and L-shaped approach. However, extensions have been proposed, such as the integer L-shaped method to handle integer programming subproblems by generating valid inequalities (also called integer optimality cuts) instead of standard optimality cuts [172]. Generalized Benders decomposition also extends this to nonlinear and integer-constrained problems by solving a relaxed master problem and generating outer approximations instead of simple linear cuts [173].

There are also hybrid decomposition methods like combining branch-and-bound with Benders decomposition (Branch-and-Benders cut) [174] and scenario decomposition, where scenarios are solved separately using heuristics. The latter decomposes large-scale stochastic optimization problems into scenario-based subproblems, facilitating more efficient solutions [175]. The choice of method used depends on whether integer variables appear in the first stage, the second stage, or both.

6.10.2 Stochastic dual dynamic programming (SDDP)

SDDP is an extension of Benders Decomposition tailored for multistage convex stochastic optimization problems. It is particularly effective for high-dimensional problems with a large number of stages and scenarios. It operates by approximating the cost-to-go (or value) functions using cutting planes, enabling the handling of uncertainties over multiple stages. It involves a (i) forward pass that simulates or samples paths (forward in time) of the stochastic process variables to generate some trial (or candidate) solutions and (ii) a backward pass that uses the trial solutions to solve each of the subproblems backwards in time by generating Benders cuts that approximate the value functions [176]. The following is an overview of the SDDP algorithm:

1. **Initialization:**
   - Define the planning horizon T and the set of scenarios Ω.
   - Initialize the value function approximations $\hat{Q}_t(x)$ for all stages $t = 1, \ldots, T$.
2. **Forward Pass:**
   - Simulate sample paths of the stochastic process ω_t.
   - Solve the deterministic subproblem for each stage t along the sample path to generate trial solutions.
3. **Backward Pass:**
   - For each stage $t = T, T-1, \ldots, 1$, solve the subproblem (in reverse time) to compute the cost-to-go function.
   - Update the value function approximation $\hat{Q}_t(x)$ using cutting planes (Benders) derived from the dual solutions.
4. **Convergence Check:**

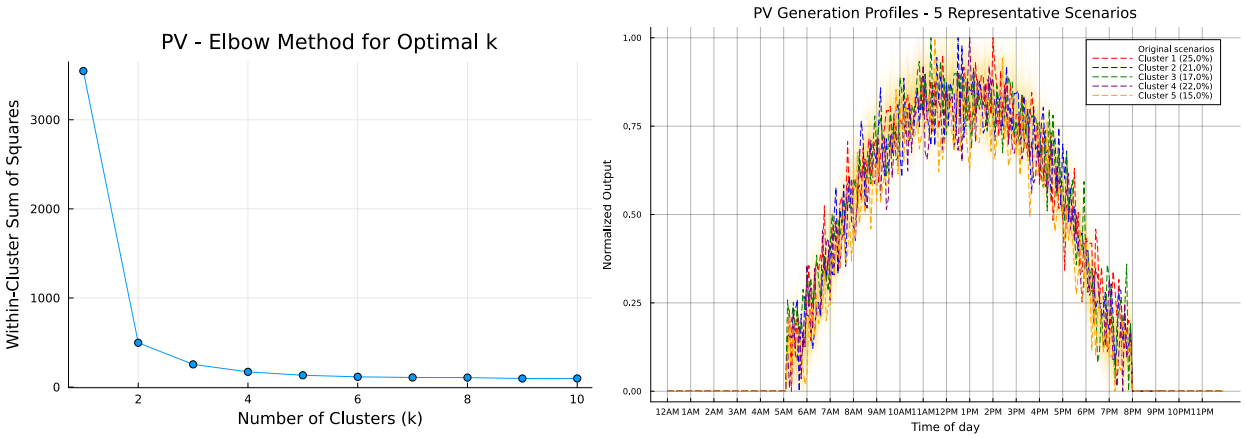

Figure 6.8: Reduced scenario analysis.

- Check if the value function approximations have converged.
- If not, return to the forward pass.

Note that our HCA problem only involves 2 stages, i.e., $T = 2$. SDDP is widely used for various applications in energy systems (e.g., long-term hydropower scheduling and management), financial portfolio management, and supply chain or inventory optimization under stochastic demand. In terms of computational complexity, SDDP's sampling-based approach can be more efficient for problems with many stages, as it avoids explicit enumeration of all scenarios, which is often required in traditional Benders Decomposition. However, even so, the worst-case complexity of SDDP scales exponentially in the number of decision variables [177]. Extensions of SDDP to nonlinear [178], mixed integer [176], and mixed integer nonlinear programs (MINLP) [179] have also been proposed.

However, it is quite challenging to adapt either Benders decomposition or SDDP for our specific problem since it is an MINLP with a complex structure that's not readily amenable to either algorithm. The presence of integer variables in the 2nd stage (instead of the 1st) makes it difficult to apply such cutting plane methods since they rely on the convexity of the 2nd stage value function (which is nonconvex in our case).

6.10.3 Accelerating 2-SSP for HCA

As an alternative to using decomposition algorithms, we continued to use the deterministic equivalent approach, but instead used scenario reduction to accelerate the simulation. We applied k-means clustering, an unsupervised machine learning method, to first cluster all the profiles in the full scenario set. This clustering allows us to identify the most representative scenarios that capture most of the trends in the original timeseries data. The left plot in Fig. 6.8 shows how we applied the elbow method to identify the optimal number of clusters for k-means. We see that the within-cluster variation stops decreasing after around 5 clusters. Thus, we chose to sample 5 scenarios out of the 100 total scenarios for our analysis. The right plot shows an example of what the representative scenarios look like for the PV radiation profiles, overlaid on the raw data.

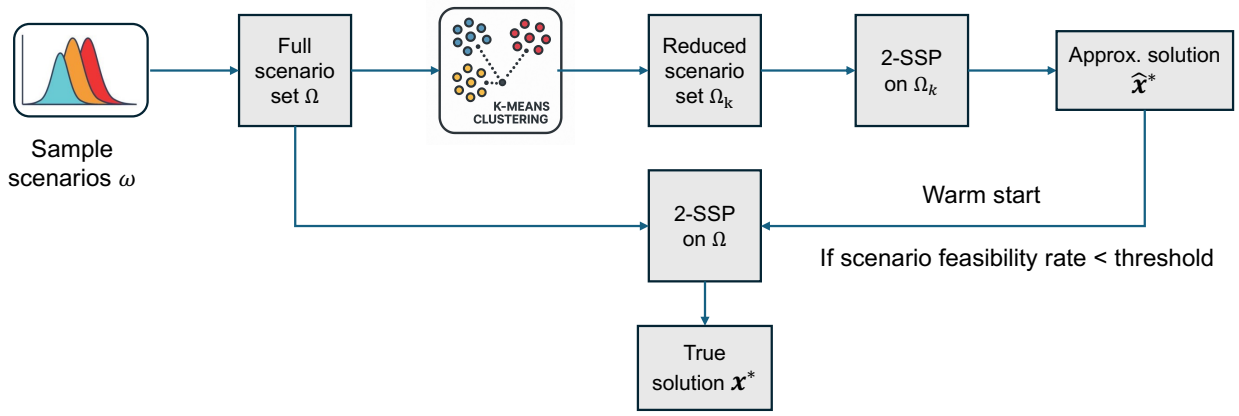

Figure 6.9: Accelerated 2SSP workflow.

We combined this scenario reduction with another algorithm to further accelerate the 2-SSP. This method is described in Fig. 6.9. The problem is first solved while considering the reduced scenario set. This can be solved much more quickly due to the smaller number of scenarios. This provides an approximate solution, but it may not be feasible for all the original scenarios. In practice, though, these crude solutions can perform reasonably well in terms of feasibility. However, if the scenario feasibility rate (i.e., the percentage of original scenarios for which the approximate solution is feasible) is below the threshold set by the planner, we can resolve the 2-SSP on the full set of scenarios. However, by using the crude solution to warm start the optimization, we can potentially accelerate convergence.

6.11 Simulation results and discussion

6.11.1 Input data

The key inputs into our simulation are data related to DER power injections for both generation and load. For the demand side, this entails load profiles and temperature data, which influence AC and HP use. On the supply side, these include weather data (solar irradiation) that affects the PV output. In addition to the baseline injection profiles, there are also distributions of parameters that influence the amount of flexibility (along with its temporal signature) that customers (or prosumers) are willing to provide. These include both downward and upward flexibilities in generation and load. For the static cases without flexibility, we need to come up with several representative scenarios and profiles for DER power injections, derived from real historical data. These specify how prosumers operate their DERs in the nominal case. It is especially important to have accurate estimates (net profiles) for BS, EV, and HP operation, while PV output will generally just follow radiation data. Finally, the other main input is locational marginal price (LMP) data from the ISOs, which influences how the DSO and agents balance importing power from the main transmission grid versus relying on the dispatch of local assets like batteries and EVs. The remaining parameters are generated randomly for each feeder and assumed to be the same for all scenarios of inputs. For example, distributions of the key EV parameters can be obtained from actual

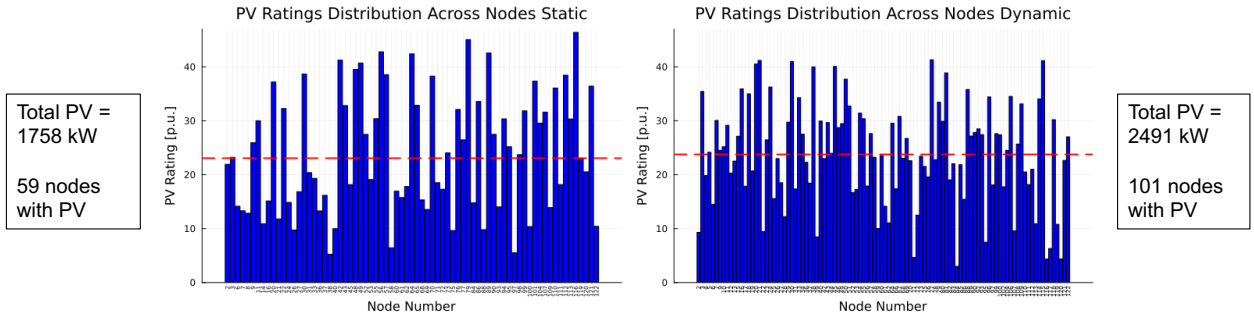

Figure 6.10: PV distribution histograms.

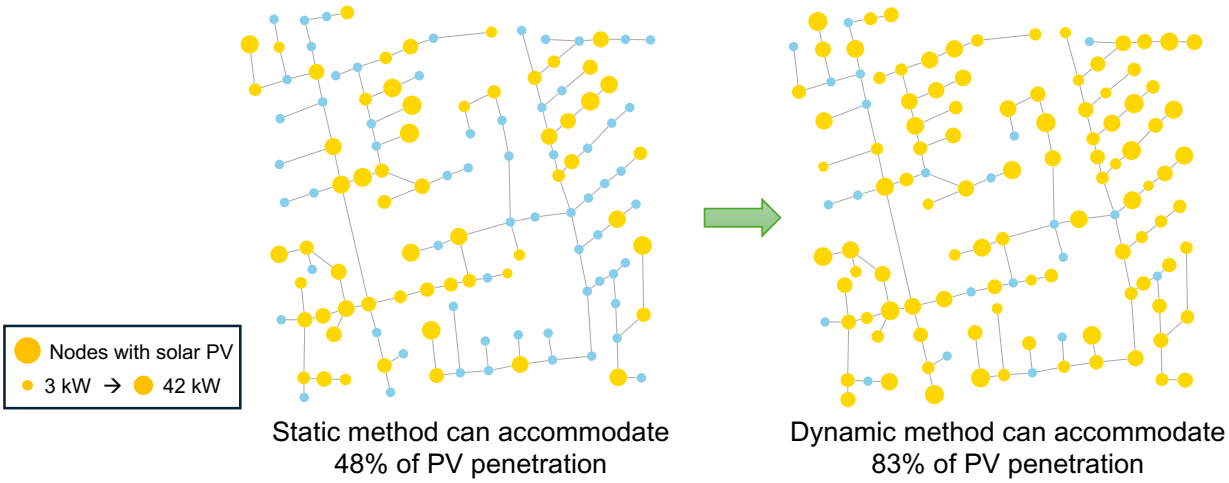

Figure 6.11: PV network plot - deterministic iteration.

load profiles of EV charging data. These include the EV unavailable time windows (i.e., release and deadline times [133]) and desired charge tracking time and SOC target value (e.g., achieve 85% SOC by 9am).

6.11.2 Deterministic iterative HCA results

We first consider the problem where we maximize PV penetration for fixed penetrations of BS, EVs, and HPs. BS penetration is set to 5% of total feeder peak load, while 5% of homes are assumed to have EVs and HPs. Fig. 6.10 shows the distributions of solar capacities over the network. We see that the dynamic approach significantly increases the total PV capacity to 83% compared to the 48% with the static case. Thus, coordination allows for a roughly 70% relative increase of solar penetration without curtailment. Although the mean capacity (shown by the red dashed line) is similar in both cases, the dynamic case has PV present at many more nodes. Fig. 6.11 shows the same results but overlaid on the network plot itself to show where the PV nodes (marked by the gold circles) are located, with the diameter representing the nodal PV capacity.

To further interpret these results, we inspect the key power flow metrics. From Fig. 6.12, we see that the dynamic approach in Fig. 6.12 is able to maintain relatively low voltages that

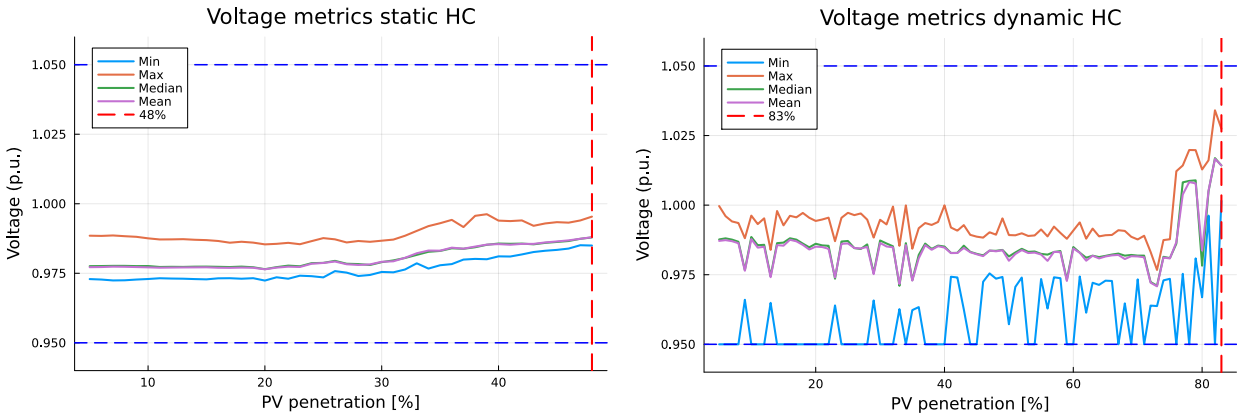

Figure 6.12: Voltage metrics - deterministic iteration.

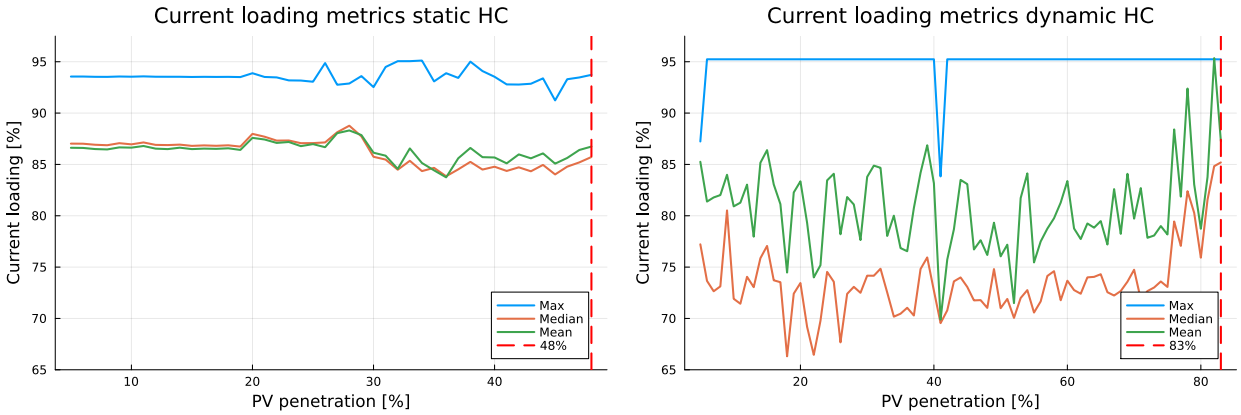

Figure 6.13: Current metrics for deterministic iteration.

are comparable to or even lower than the static case up until around 75% PV. The static case only remains feasible until 48% PV. As we increase PV beyond 75% in the dynamic case, the voltages start to rise rapidly until we hit infeasibility at 83% PV penetration. We also notice that there's more volatility in the voltage metrics in the dynamic case than in the static case. Comparing the two cases in Fig. 6.13, we find that although the maximum currents are similar, the mean and median network currents are significantly lower in the dynamic case. This indicates that leveraging flexibility with the dynamic approach has significant distributional benefits in lowering current loading throughout the network, even if the worst-case values are similar. Examining these metrics also helps us identify that nodal overvoltage and highly loaded lines are likely the main limiting factors for hosting capacity.

Fig. 6.14 shows the changes in the key metrics as we increase the PV penetration from 48% in the static case to 75% in the dynamic. We see that even though we push the network to much higher PV levels, the dynamic method is able to reduce overvoltage issues at all nodes (middle plot) and quite significantly bring down current loading in most lines (left plot), although there are increases in some lines, particularly near the substation where the distribution grid is connected to transmission. This is a good sign as it indicates that the network is able to handle higher levels of PV without overloading or overvoltage issues,

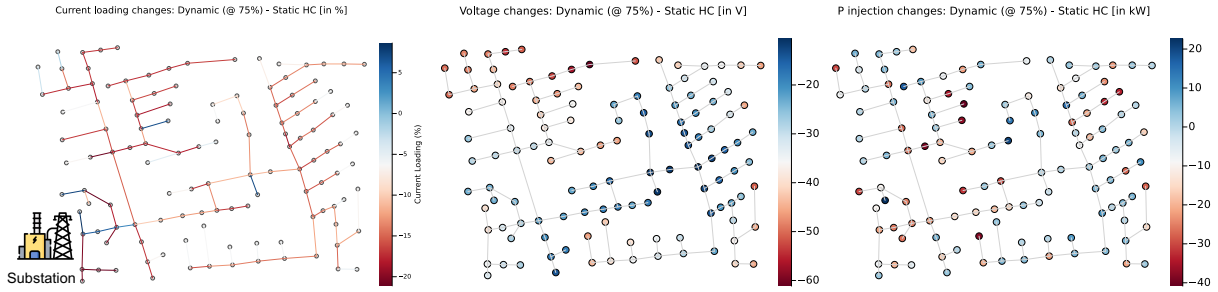

Figure 6.14: Metric changes from 48% PV (static) to 75% (dynamic).

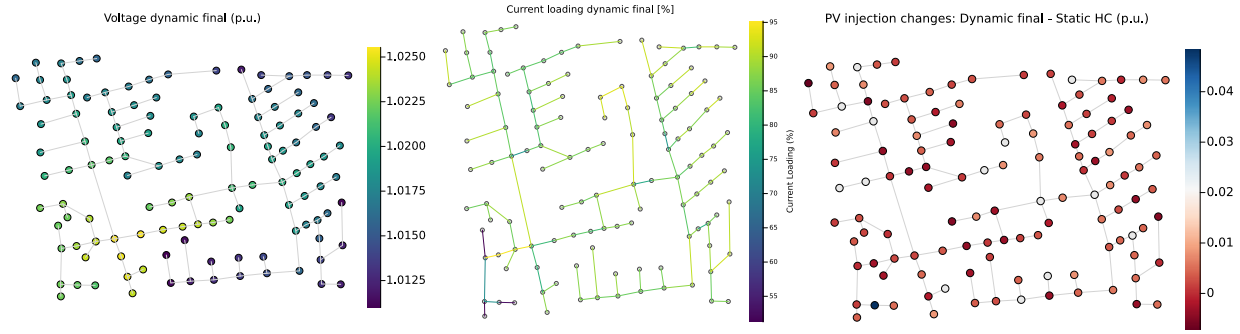

Figure 6.15: Final network dynamic results.

allowing higher active power injections at many nodes in the network (as seen in the right plot). Note that Figs. 6.14, 6.15 and 6.17 are all network plots for a single snapshot of time during midday (12:30pm ET). This time instant is of special interest to us since it corresponds to the peak PV output, which is most likely to cause overvoltage and current load issues.

In Fig. 6.15, we show the final state of the network at the maximum dynamic PV HC of 83%, and we see that remarkably, even as we push PV injections to the maximum feasible level, our method guarantees that voltages and currents are within limits.

Now, we provide some insights into exactly how the dynamic market-based coordination is able to support increased PV hosting capacity. Fig. 6.16 shows the changes in the PV capacities at nodes between the dynamic case (84%) and static case (48%). Comparing this against the BS and HP dispatch plots in Fig. 6.17, we see that the flexible batteries and heat pumps play a key role in supporting PV. Essentially, the heat pumps consume more power at several key nodes, and the batteries charge more in order to help absorb the excess PV output midday. The total BS power is 147 kW (5% of the total maximum load of 2930 kVA) with an energy storage capacity of 441 kWh, while the total HP capacity is 168 kW (representing 5% of electrified homes). The network also has significant EV capacity (396 kW, 2700 kWh), but these are only available for V1G (managed periods) for limited periods of the day. We assumed that most EVs were not available during the middle of the day when they're likely away from home, and thus they didn't play as important a role here. EV charging will likely be more helpful to reduce demand overnight when charging usually occurs. Comparing the clusters (dashed circles), we see that co-located batteries help boost PV at the same node, but batteries also support PV at nearby nodes due to the network effect.

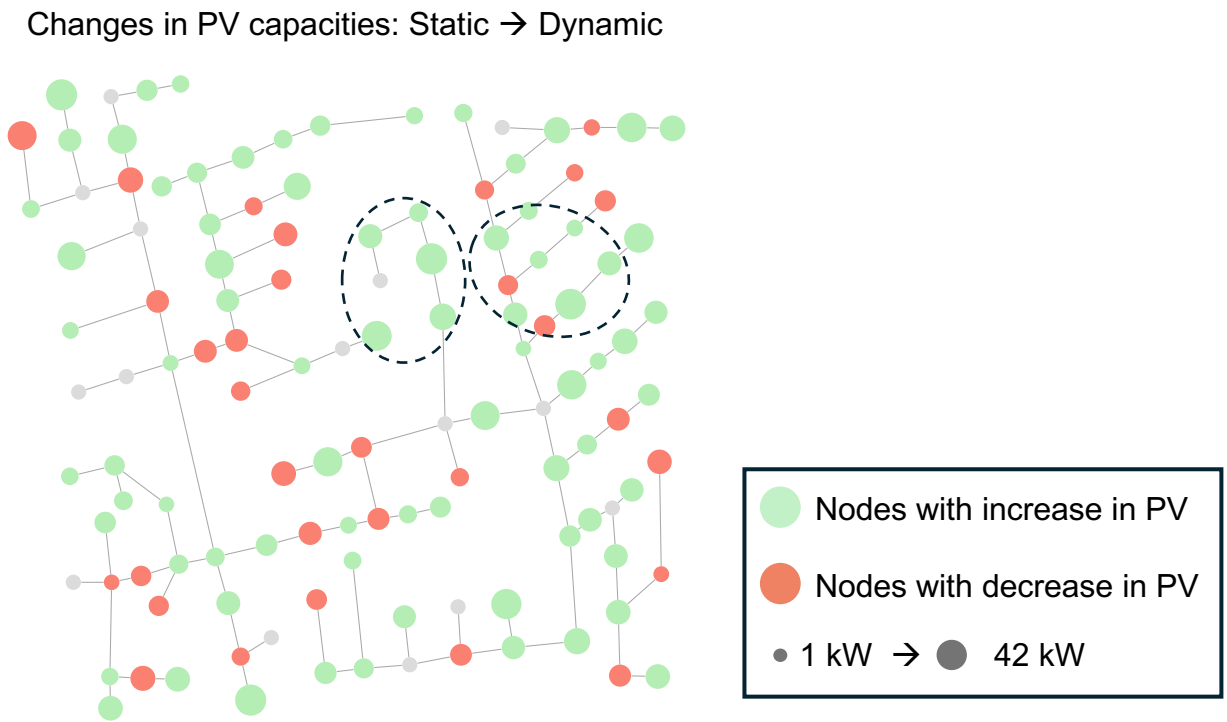

Figure 6.16: PV size changes in network - deterministic iteration.

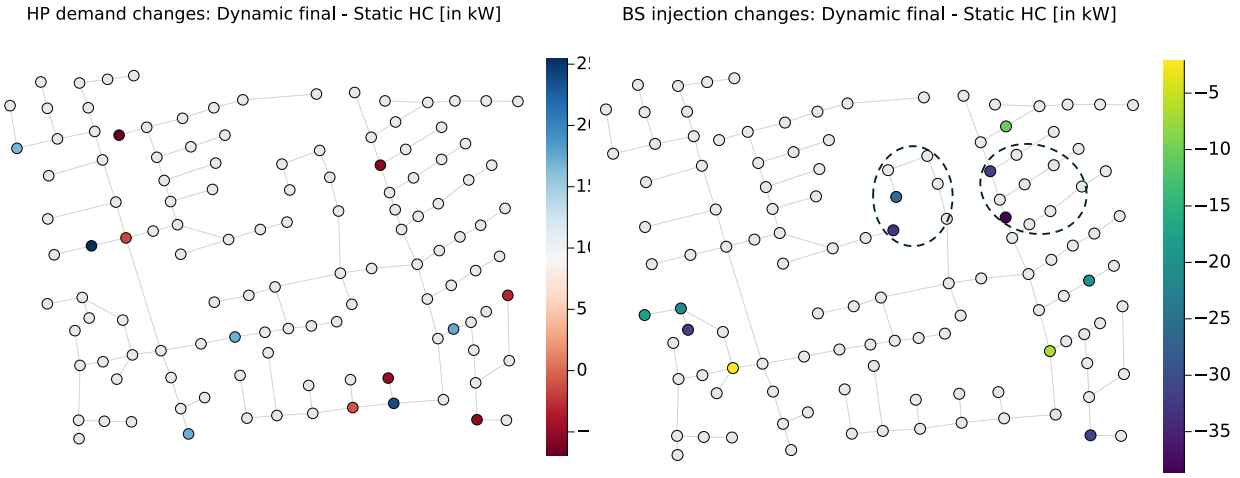

Figure 6.17: HP BS boost hosting capacity - deterministic iteration.

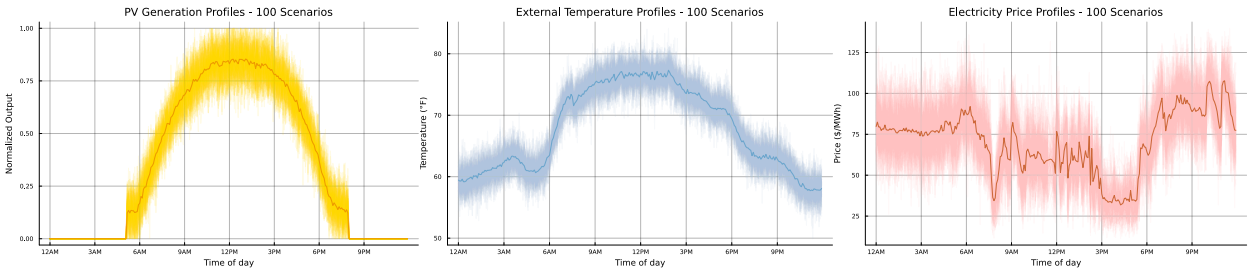

Figure 6.18: DER complementarity analysis.

	Static	Dynamic
Mean HC [%]	51.08	83.63
Standard deviation [%]	2.36	3.44

Table 6.2: Summary metrics of HC distributions with stochastic iterative approach.

Fig. 6.18 clearly shows the complementary relationships among different DERs based on the key inputs. We see that high PV output in the middle of the day also coincides with higher external temperatures and greater cooling demand, thereby allowing HPs to contribute their flexibility. During winter months, we expect EVs to play more of a role since their peak demand overnight also complements lower temperatures and higher heating demand, as well as peak wind power production. Similarly, electricity prices are also lower mid-day (likely due to excess PV charging), which further incentivizes batteries to charge more when cheap power is available.

6.11.3 Stochastic iterative HCA results

We now discuss the results of the stochastic iterative approach from Section 6.7.2. Fig. 6.19 shows the distributions of the HC estimates obtained by running the method for 100 randomly sampled scenarios. Similar to the deterministic case, we see that the dynamic approach pushes the entire distribution further towards the right with higher HC values. However, as seen from Table 6.2, the dynamic distribution is slightly wider and has a higher standard deviation. This indicates that the dynamic approach introduces more uncertainty in hosting capacities across scenarios. This is likely because the input uncertainty affects the DER flexibilities and thus the constraints of the underlying ACOPF problem solved in the dynamic case. Such uncertainty propagation through optimization problem constraints can often magnify or amplify the original uncertainties [180].

Similar to the HC distributions, the stochastic iterative approach also allows us to estimate distributions of the key power flow metrics. The left and right plots in Fig. 6.20 show the distributions of the maximum and mean voltage over the network, when averaged over all intermediate penetration levels (at each iteration), respectively. We see that both the mean and maximum voltages are higher in the dynamic case. This is expected since dynamism allows us to attain much higher PV penetrations. The dynamic maximum voltage has much higher uncertainty than the static maximum, while the dynamic and static means display

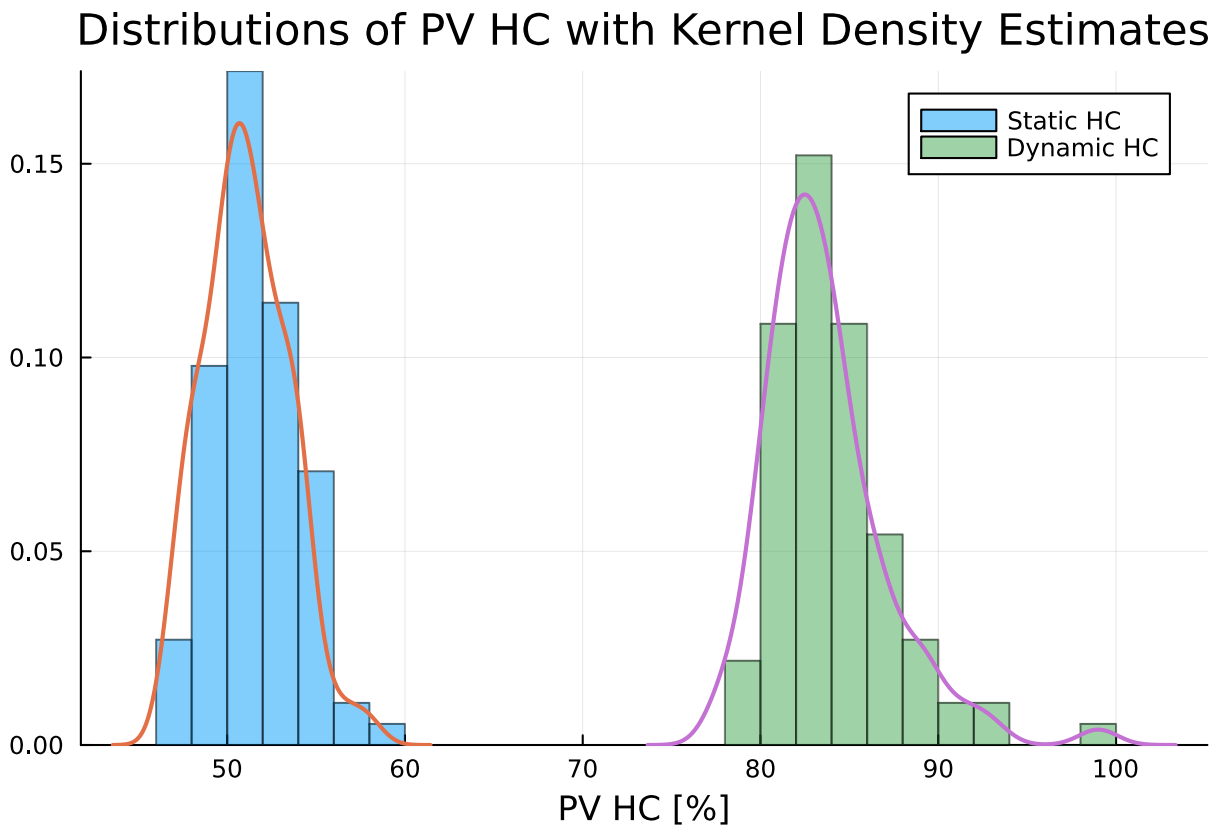

Figure 6.19: Stochastic iteration hosting capacity distribution.

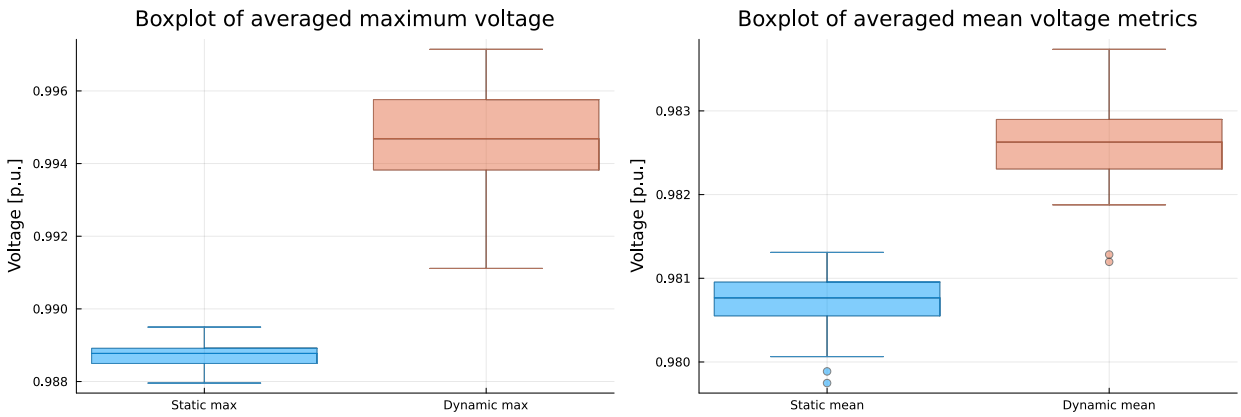

Figure 6.20: Stochastic iteration voltage distribution.

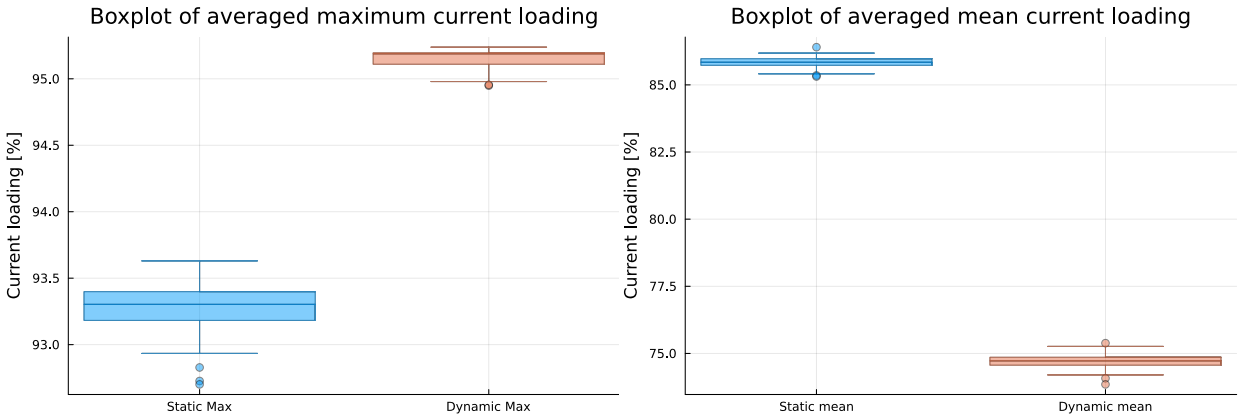

Figure 6.21: Stochastic iteration current distribution.

similar amounts of variation. This may be because the extreme voltages are more sensitive to changes in constraints. The maximum currents are slightly higher for the dynamic case but with lower uncertainty, relative to the static case. This is once again due to the higher PV injections enabled by dynamic coordination, which raises both the worst-case voltages and currents. However, similar to Fig. 6.14 and Fig. 6.13, the dynamic mean current is significantly lower and with similar uncertainty levels as the static case. This is likely because the flexibilities of the HPs, EVs, and BS in solving the ACOPF problem help reroute the power flows more efficiently and also minimize resistive losses, which is one of the objectives in Section 6.4.4. Furthermore, coordinated operation of BS and flexible loads co-located with solar can help reduce the net injections by that node into the grid.

6.11.4 2-SSP HCA results

Here we present selected results obtained by solving the two-stage stochastic program (2-SSP) on the IEEE 123-node test network with 10% BS penetration, 10% of homes with EVs and HPs, and by considering $N = 25$ scenarios. Fig. 6.22 shows the dynamically optimized locations and capacities of each of the four DER types. Interestingly, we see that the optimization

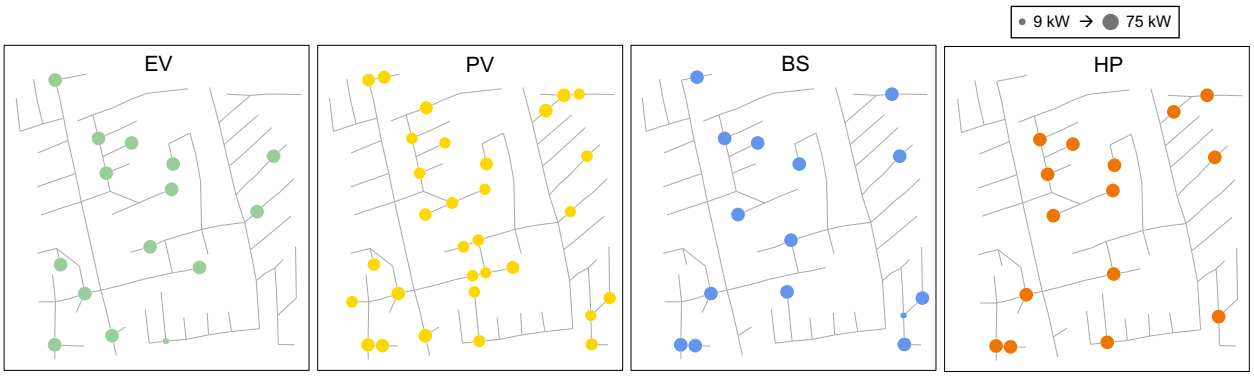

Figure 6.22: DER distribution in 2SSP.

solution results in a generally uniform distribution of capacities over nodes and also favors co-locating batteries with solar. Electrified loads (HP, EV) also tend to be co-located with batteries and/or solar. Note that this co-location was not enforced by design but is rather just an outcome of the optimization. An intuitive explanation for this phenomenon could be that co-location enables each node to be more self-sufficient (by using self-generation to meet self-demand), and thus reduces its net injections into the grid (or power drawn by it from the grid). This reduces the stress on the grid resulting from additional PV or demand and thus allows the network to host more DERs. The total PV HC for this case was determined to be 72%.

Finally, the accelerated version of the 2-SSP with scenario reduction + warm start enables rapid sensitivity studies. Fig. 6.23 shows the results for a sensitivity case study where we varied both the BS and HP penetrations in order to assess their effects on PV penetration. The left plot shows the results for a contrived worst-case version of the static case with very high renewable generation and low demand (including HP heating or cooling demand). In this case, we see that the PV HC is mainly sensitive to the BS penetration since moving horizontally causes significant PV changes. HP penetration has minimal effects since moving vertically doesn't significantly change the PV hosting capacity values. This is because the HP flexibility is not very helpful since demand is already low. The colored square indicates all the combinations of DER penetrations that were found to be feasible for each case, with combinations outside the box being infeasible for the power flow. The feasible region is thus very limited, only allowing up to about 40% of BS and HP penetration and around 35% PV.

The center plot shows results for a more realistic static case. The direction of the contour lines suggests that BS and HP penetration have opposing or competing effects on PV hosting capacity. As the BS penetration increases, we need to accordingly decrease HP penetration in order to retain the same PV hosting capacity. This creates a substitution relationship between BS and HP, which arises due to the grid constraints and insufficient hosting capacity. However, the feasible region is larger than the worst case, allowing higher DER penetrations. Finally, on the right, we see that the dynamic approach relieves limiting grid constraints and significantly expands the feasible space, allowing up to 100% BS and HP electrification and very high PV penetration levels. The contour lines suggest that BS and HPs have complementary effects on PV hosting capacity since higher levels of both BS and HP penetration together lead to higher PV hosting capacity. The smooth gradient indicates that both technologies work

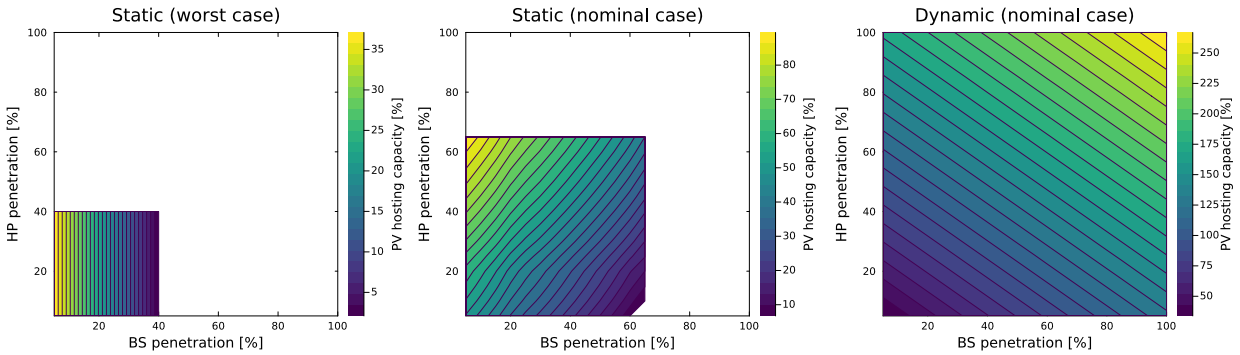

Figure 6.23: 2SSP sensitivity analysis.

together synergistically.

6.12 Conclusions

In this chapter, we showed that market-based distributed coordination of DERs can aid in both distribution grid planning and operation. Leveraging complementary effects among different types of DERs can boost grid HC, and such coordination can also improve network operation and efficiency by reducing current flows and avoiding voltage violations. Using a dynamic approach can integrate more DERs while minimizing new physical infrastructure or upgrades. This can also help circumvent long permitting and build times, clear DER connection queues, and accelerate grid decarbonization while also lowering costs for all stakeholders.

Chapter 7

Conclusion

In this chapter, we developed a suite of computational tools to optimally plan and operate future decarbonized power grids that are rich in renewables, batteries, and flexible electrified demand, such as electric vehicles, heat pumps, and data centers. The methods and algorithms are primarily designed for medium and low-voltage distribution grids with large numbers of small-scale distributed energy resources (DERs). However, these can be easily adapted to be applied to larger, utility-scale resources at the high transmission grid level as well.

The first three chapters of the thesis focus on a new hierarchical local retail electricity market structure for the distribution grid. Three different markets are proposed, each using distinct optimization algorithms and analytical tools. The tiered structure allows us to accurately account for all physical constraints as well as the constraints and preferences of disparate stakeholders. The uppermost primary market incorporates grid physics and power flow constraints into a network-wide optimization problem. The use of distributed optimization helps preserve data privacy and makes the method scalable for large networks with thousands of nodes. By leveraging local information and limited information from neighbors, we are able to successfully solve the problem while also achieving global optimality and feasibility. The middle secondary market layer utilizes decentralized optimization at each node and is coupled with the upper layer. This accounts for other factors like agent utility, flexibility maximization, reliability, cost, and budget constraints. Extensive simulations show that the primary and secondary markets together successfully increase reliability, reduce losses, and reduce operating costs and retail tariffs for consumers. The market is also generalized to different types of networks using different power flow models and then applied to provide grid services like voltage regulation. In addition to significantly improving voltage profiles, the framework also allows for deriving and decomposing prices for different grid services, in order to accurately charge or compensate prosumers.

The third chapter describes the lower consumer market layer, which is closest to the end users and consumers. This applies game theory and mechanism design to optimally aggregate DERs and extract flexibility. This is done via a repeated, two-stage Stackelberg game with incomplete information. We also model different types of DERs in detail to solve a multiperiod optimization problem. We prove the existence of unique Stackelberg and Nash equilibria among market operators and agents with closed-form, analytical solutions. The fourth chapter applies our distributed coordination framework to increase grid resilience by detecting and mitigating several types of small, medium, and large-scale attacks and outages. This is done

by coordinating DERs that are digitally connected via Internet of Things (IoT) networks. We compute various custom metrics for each agent's market commitment reliability and IoT trustability, which lead to overall resilience scores. The situational awareness provided by the market structure and resilience infrastructure, combined with these resilience scores, allows the grid operators to successfully deploy trustable DERs to counter generation or load alteration attacks. In addition to numerical simulations, we validate our algorithms and results using both industry-grade software as well as hardware-in-the-loop experiments.

The final chapter examines how we can coordinate DERs to increase the hosting capacity of distribution grids. This is essential to rapidly integrate new DERs and meet growing electricity demand in a reliable manner. We show that a dynamic approach based on real-time optimization and coordination significantly boosts the capacity of the grid to host DERs, compared to the static case. We leverage the inherent flexibility available in distributed generation, energy storage, and loads at the grid edge. The main factors that allow us to increase hosting capacity are the complementary relationships and synergies among different types of DERs, such as solar, batteries, electric vehicles, and heat pumps, that help relieve grid constraints.

This thesis makes several novel and meaningful contributions on both theoretical and applied fronts of distribution grid planning and operation. These methods will be useful to a wide range of users, including grid operators, utilities, DER aggregators, and individual consumers or prosumers. Simulation results confirm that we can thus improve the reliability, resilience, and affordability of low-carbon power systems. We also provide deeper insights and interpretations into how DER flexibility and coordination among different agents in the network is able to achieve such beneficial effects.

References

[1] M. F. Akorede, H. Hizam, and E. Pouresmaeil. "Distributed energy resources and benefits to the environment". In: *Renewable and Sustainable Energy Reviews* 14.2 (Feb. 2010), pp. 724–734. ISSN: 1364-0321. DOI: 10.1016/J.RSER.2009.10.025.

[2] R. Haider, D. D'Achiardi, V. Venkataramanan, A. Srivastava, A. Bose, and A. M. Annaswamy. "Reinventing the utility for distributed energy resources: A proposal for retail electricity markets". In: *Advances in Applied Energy* 2 (2021), p. 100026. ISSN: 2666-7924. DOI: https://doi.org/10.1016/j.adapen.2021.100026. URL: https://www.sciencedirect.com/science/article/pii/S2666792421000196.

[3] Y. K. Renani, M. Ehsan, and M. Shahidehpour. "Optimal Transactive Market Operations With Distribution System Operators". en. In: *IEEE Transactions on Smart Grid* 9.6 (Nov. 2018), pp. 6692–6701. ISSN: 1949-3053, 1949-3061. DOI: 10.1109/TSG.2017.2718546. URL: https://ieeexplore.ieee.org/document/7955039/ (visited on 02/03/2021).

[4] Y. Wang, Y. He, Q. Duan, H. Wang, and P. Bie. "Distribution Locational Marginal Pricing for Active Network Management Considering High Penetration of Distributed Energy Resources". In: *2019 IEEE 3rd Conference on Energy Internet and Energy System Integration (EI2)*. Nov. 2019, pp. 648–653. DOI: 10.1109/EI247390.2019.9061809.

[5] J. Driesen and F. Katiraei. "Design for distributed energy resources". In: *IEEE Power and Energy Magazine* 6.3 (May 2008), pp. 30–39. ISSN: 15407977. DOI: 10.1109/MPE.2008.918703.

[6] P. D. Ferreira, P. M. Carvalho, L. A. Ferreira, and M. D. Ilic. "Distributed energy resources integration challenges in low-voltage networks: Voltage control limitations and risk of cascading". In: *IEEE Transactions on Sustainable Energy* 4.1 (2013), pp. 82–88. ISSN: 19493029. DOI: 10.1109/TSTE.2012.2201512.

[7] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei. "A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems". In: *IEEE Transactions on Smart Grid* 8.6 (Nov. 2017). Conference Name: IEEE Transactions on Smart Grid, pp. 2941–2962. ISSN: 1949-3061. DOI: 10.1109/TSG.2017.2720471.

[8] J. J. Romvary, G. Ferro, R. Haider, and A. M. Annaswamy. "A Proximal Atomic Coordination Algorithm for Distributed Optimization". In: *IEEE Transactions on Automatic Control* (2021).

[9] R. Haider, S. Baros, Y. Wasa, J. Romvary, K. Uchida, and A. M. Annaswamy. "Toward a Retail Market for Distribution Grids". In: *IEEE Transactions on Smart Grid* 11.6 (2020), pp. 4891–4905. DOI: 10.1109/TSG.2020.2996565.

[10] R. Haider, G. Ferro, M. Robba, and A. M. Annaswamy. "Flattening the Duck Curve: A Case for Distributed Decision Making". In: (Nov. 2021). DOI: 10.48550/arxiv.2111.06361. URL: https://arxiv.org/abs/2111.06361v2.

[11] T. Chen, Q. Alsafasfeh, H. Pourbabak, and W. Su. "The Next-Generation U.S. Retail Electricity Market with Customers and Prosumers—A Bibliographical Survey". en. In: *Energies* 11.1 (Jan. 2018). Number: 1 Publisher: Multidisciplinary Digital Publishing Institute, p. 8. DOI: 10.3390/en11010008. URL: https://www.mdpi.com/1996-1073/11/1/8 (visited on 02/22/2021).

[12] H. Kim and M. Thottan. "A two-stage market model for microgrid power transactions via aggregators". In: *Bell Labs Technical Journal* 16.3 (2011), pp. 101–107. DOI: 10.1002/bltj.20524.

[13] Y. Liu, P. Bansal, R. Daziano, and S. Samaranayake. "A framework to integrate mode choice in the design of mobility-on-demand systems". en. In: *Transportation Research Part C: Emerging Technologies* 105 (Aug. 2019), pp. 648–665. ISSN: 0968-090X. DOI: 10.1016/j.trc.2018.09.022. URL: http://www.sciencedirect.com/science/article/pii/S0968090X18313718 (visited on 11/07/2019).

[14] P. Olivella-Rosell, P. Lloret-Gallego, Í. Munné-Collado, R. Villafafila-Robles, A. Sumper, S. Ø. Ottessen, J. Rajasekharan, and B. A. Bremdal. "Local Flexibility Market Design for Aggregators Providing Multiple Flexibility Services at Distribution Network Level". en. In: *Energies* 11.4 (Apr. 2018). Number: 4 Publisher: Multidisciplinary Digital Publishing Institute, p. 822. DOI: 10.3390/en11040822. URL: https://www.mdpi.com/1996-1073/11/4/822 (visited on 03/08/2021).

[15] S. Burger, J. P. Chaves-Ávila, C. Batlle, and I. J. Pérez-Arriaga. "A review of the value of aggregators in electricity systems". In: *Renewable and Sustainable Energy Reviews* 77 (Sept. 2017), pp. 395–405. ISSN: 1364-0321. DOI: 10.1016/J.RSER.2017.04.014.

[16] L. Kristov, P. De Martini, and J. D. Taft. "A Tale of Two Visions: Designing a Decentralized Transactive Electric System". In: *IEEE Power and Energy Magazine* 14.3 (2016), pp. 63–69. DOI: 10.1109/MPE.2016.2524964.

[17] A. M. Annaswamy, Y. Guan, H. E. Tseng, H. Zhou, T. Phan, and D. Yanakiev. "Transactive Control in Smart Cities". In: *Proceedings of the IEEE* 106.4 (Apr. 2018). Conference Name: Proceedings of the IEEE, pp. 518–537. ISSN: 1558-2256. DOI: 10.1109/JPROC.2018.2790841.

[18] A. Papalexopoulos. "The Evolution of the Multitier Hierarchical Energy Market Structure: The emergence of the transactive energy model". In: *IEEE Electrification Magazine* 9.3 (Sept. 2021), pp. 37–45. ISSN: 23255889. DOI: 10.1109/MELE.2021.3093598.

[19] S. Li, J. Lian, A. J. Conejo, and W. Zhang. "Transactive Energy Systems: The Market-Based Coordination of Distributed Energy Resources". In: *IEEE Control Systems* 40.4 (Aug. 2020), pp. 26–52. ISSN: 1941000X. DOI: 10.1109/MCS.2020.2990514.

[20] Y. Wang, W. Saad, Z. Han, H. V. Poor, and T. Basar. "A Game-Theoretic Approach to Energy Trading in the Smart Grid". en. In: *IEEE Transactions on Smart Grid* 5.3 (May 2014), pp. 1439–1450. ISSN: 1949-3053, 1949-3061. DOI: 10.1109/TSG.2013.2284664. URL: http://ieeexplore.ieee.org/document/6798766/ (visited on 02/03/2021).

[21] M. Motalleb, A. Annaswamy, and R. Ghorbani. "A real-time demand response market through a repeated incomplete-information game". en. In: *Energy* 143 (Jan. 2018), pp. 424–438. ISSN: 0360-5442. DOI: 10.1016/j.energy.2017.10.129. URL: https://www.sciencedirect.com/science/article/pii/S0360544217318340 (visited on 11/22/2021).

[22] Z. Liang and W. Su. "Game theory based bidding strategy for prosumers in a distribution system with a retail electricity market". en. In: *IET Smart Grid* 1.3 (2018). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1049/iet-stg.2018.0048, pp. 104–111. ISSN: 2515-2947. DOI: 10.1049/iet-stg.2018.0048. URL: https://onlinelibrary.wiley.com/doi/abs/10.1049/iet-stg.2018.0048 (visited on 10/19/2021).

[23] Z. M. Fadlullah, Y. Nozaki, A. Takeuchi, and N. Kato. "A survey of game theoretic approaches in smart grid". In: *2011 International Conference on Wireless Communications and Signal Processing (WCSP)*. Nov. 2011, pp. 1–4. DOI: 10.1109/WCSP.2011.6096962.

[24] O. Karaca and M. Kamgarpour. "Game theoretic analysis of electricity market auction mechanisms". In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. Dec. 2017, pp. 6211–6216. DOI: 10.1109/CDC.2017.8264596.

[25] D. Csercsik. "Strategic bidding via the interplay of minimum income condition orders in day-ahead power exchanges". In: *Energy Economics* 95 (Mar. 2021), p. 105126. ISSN: 0140-9883. DOI: 10.1016/J.ENECO.2021.105126.

[26] M. S. Ibrahim, W. Dong, and Q. Yang. "Machine learning driven smart electric power systems: Current trends and new perspectives". en. In: *Applied Energy* 272 (Aug. 2020), p. 115237. ISSN: 03062619. DOI: 10.1016/j.apenergy.2020.115237. URL: https://linkinghub.elsevier.com/retrieve/pii/S0306261920307492 (visited on 02/03/2021).

[27] X. Pan, T. Zhao, M. Chen, and S. Zhang. "DeepOPF: A deep neural network approach for security-constrained DC optimal power flow". In: *IEEE Transactions on Power Systems* 36.3 (May 2021), pp. 1725–1735. ISSN: 15580679. DOI: 10.1109/TPWRS.2020.3026379.

[28] M. Chatzos, F. Fioretto, T. W. K. Mak, and P. V. Hentenryck. "High-Fidelity Machine Learning Approximations of Large-Scale Optimal Power Flow". In.

[29] S. Liu, C. Wu, and H. Zhu. "Graph Neural Networks for Learning Real-Time Prices in Electricity Market". In.

[30] V. Venkataramanan, S. Kaza, and A. M. Annaswamy. "DER Forecast using Privacy Preserving Federated Learning". In: *IEEE Internet of Things Journal* (2022). ISSN: 23274662. DOI: 10.1109/JIOT.2022.3157299.

[31] X. Chen, G. Qu, Y. Tang, S. Low, and N. Li. "Reinforcement Learning for Selective Key Applications in Power Systems: Recent Advances and Future Challenges". In.

[32] W. W. Hogan. "Independent system operator: Pricing and flexibility in a competitive electricity market". In: *Center for Business and Government, JF Kennedy School of Government, Harvard University, MA* (1998).

[33] Y. Tohidi, M. Farrokhseresht, and M. Gibescu. "A Review on Coordination Schemes Between Local and Central Electricity Markets". In: *2018 15th International Conference on the European Energy Market (EEM)*. 2018, pp. 1–5. DOI: 10.1109/EEM.2018.8470004.

[34] F. Hvelplund. "Renewable energy and the need for local energy markets". In: *Energy* 31.13 (2006), pp. 2293–2302.

[35] T. Chen, Q. Alsafasfeh, H. Pourbabak, and W. Su. "The next-generation US retail electricity market with customers and prosumers—A bibliographical survey". In: *Energies* 11.1 (2018), p. 8.

[36] R. Bray and B. Woodman. "Unlocking Local Energy Markets". In: *Proceedings of Energy for a Net Zero Society – BIEE's Oxford Conference* (2018).

[37] C. E. S. Center. *Net Metering and Net Billing; available at https://cleanenergysolutions.org/instruments/net-metering-net-billing*. 2021.

[38] F. Lezama, J. Soares, P. Hernandez-Leal, M. Kaisers, T. Pinto, and Z. Vale. "Local Energy Markets: Paving the Path Toward Fully Transactive Energy Systems". In: *IEEE Transactions on Power Systems* 34.5 (2019), pp. 4081–4088. DOI: 10.1109/TPWRS.2018.2833959.

[39] P. Olivella-Rosell, E. Bullich-Massagué, M. Aragüés-Peñalba, A. Sumper, S. Ø. Ottesen, J.-A. Vidal-Clos, and R. Villafáfila-Robles. "Optimization problem for meeting distribution system operator requests in local flexibility markets with distributed energy resources". In: *Applied energy* 210 (2018), pp. 881–895.

[40] S. D. Manshadi and M. E. Khodayar. "A hierarchical electricity market structure for the smart grid paradigm". In: *IEEE Transactions on Smart Grid* 7.4 (2015), pp. 1866–1875.

[41] S. Bjarghov, M. Löschenbrand, A. I. Saif, R. A. Pedrero, C. Pfeiffer, S. K. Khadem, M. Rabelhofer, F. Revheim, and H. Farahmand. "Developments and Challenges in Local Electricity Markets: A Comprehensive Review". In: *IEEE Access* 9 (2021), pp. 58910–58943.

[42] T. Sousa, T. Soares, P. Pinson, F. Moret, T. Baroche, and E. Sorin. "Peer-to-peer and community-based markets: A comprehensive review". In: *Renewable and Sustainable Energy Reviews* 104 (2019), pp. 367–378.

[43] Utility Dive. *ConEd virtual power plant shows how New York's REV is reforming utility practices; available at https://www.utilitydive.com/news/coned-virtual-power-plant-shows-how-new-yorks-rev-is-reforming-utility-pra/421053/*. 2016.

[44] E. Mengelkamp, J. Gärttner, K. Rock, S. Kessler, L. Orsini, and C. Weinhardt. "Designing microgrid energy markets: A case study: The Brooklyn Microgrid". In: *Applied Energy* 210 (2018), pp. 870–880.

[45] D. Thomas, J. Kazempour, A. Papakonstantinou, P. Pinson, O. Deblecker, and C. S. Ioakimidis. "A local market mechanism for physical storage rights". In: *IEEE Transactions on Power Systems* 35.4 (2020), pp. 3087–3099.

[46] A. Lüth, J. Weibezahn, and J. M. Zepter. "On Distributional Effects in Local Electricity Market Designs—Evidence from a German Case Study". In: *Energies* 13.8 (2020), p. 1993.

[47] T. Pinto, Z. Vale, and S. Widergren, eds. *Local Electricity Markets*. Academic Press, 2021, pp. i–iii. ISBN: 978-0-12-820074-2. DOI: https://doi.org/10.1016/B978-0-12-820074-2.00025-3. URL: https://www.sciencedirect.com/science/article/pii/B9780128200742000253.

[48] R. Marler and J. Arora. "Survey of multi-objective optimization methods for engineering". en. In: *Structural and Multidisciplinary Optimization* 26.6 (Apr. 2004), pp. 369–395. ISSN: 1615-147X, 1615-1488. DOI: 10.1007/s00158-003-0368-6. URL: http://link.springer.com/10.1007/s00158-003-0368-6 (visited on 06/22/2021).

[49] N. Gunantara. "A review of multi-objective optimization: Methods and its applications". In: *Cogent Engineering* 5.1 (2018), p. 1502242.

[50] O. Grodzevich and O. Romanko. "Normalization and Other Topics in MultiObjective Optimization". In.

[51] R. T. Makler and J. S. Arora. "Function-transformation methods for multi-objective optimization". In: *http://dx.doi.org.libproxy.mit.edu/10.1080/03052150500114289* 37.6 (Sept. 2006), pp. 551–570. ISSN: 0305215X. DOI: 10.1080/03052150500114289. URL: https://www-tandfonline-com.libproxy.mit.edu/doi/abs/10.1080/03052150500114289.

[52] K. Miettinen and M. M. Mäkelä. "On scalarizing functions in multiobjective optimization". In: *OR Spectrum 2002 24:2* 24.2 (2002), pp. 193–213. ISSN: 1436-6304. DOI: 10.1007/S00291-001-0092-9. URL: https://link-springer-com.libproxy.mit.edu/article/10.1007/s00291-001-0092-9.

[53] K. Deb and K. Deb. "Multi-objective Optimization". In: *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques, Second Edition* (Jan. 2014), pp. 403–449. DOI: 10.1007/978-1-4614-6940-7{_}15. URL: https://link-springer-com.libproxy.mit.edu/chapter/10.1007/978-1-4614-6940-7_15.

[54] A. Safari, P. Salyani, and M. Hajiloo. "Reactive power pricing in power markets: a comprehensive review". In: *International Journal of Ambient Energy* 41.13 (Nov. 2020). Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01430750.2018.1517675, pp. 1548–1558. ISSN: 0143-0750. DOI: 10.1080/01430750.2018.1517675. URL: https://doi.org/10.1080/01430750.2018.1517675 (visited on 09/27/2021).

[55] Lazard. "Levelized Cost of Energy Analysis - Version 12.0". In: (2018). URL: https://www.lazard.com/media/450784/lazards-levelized-cost-of-energy-version-120-vfinal.pdf.

[56] F. Camera. "Renewable Power Generation Costs in 2018". In: *International Renewable Energy Agency (IRENA): Abu Dhabi, UAE* (2019).

[57] T. R. Nudell, M. Brignone, M. Robba, A. Bonfiglio, G. Ferro, F. Delfino, and A. M. Annaswamy. "Distributed control for polygeneration microgrids: A Dynamic Market Mechanism approach". In: *Control Engineering Practice* 121 (Apr. 2022), p. 105052. ISSN: 0967-0661. DOI: 10.1016/J.CONENGPRAC.2021.105052. URL: https://linkinghub.elsevier.com/retrieve/pii/S0967066121002999.

[58] H. Hao, Y. Lin, A. S. Kowli, P. Barooah, and S. Meyn. "Ancillary Service to the grid through control of fans in commercial Building HVAC systems". In: *IEEE Transactions on Smart Grid* 5.4 (2014), pp. 2066–2074. ISSN: 19493053. DOI: 10.1109/TSG.2014.2322604.

[59] M. J. Garcia, T. R. Nudell, and A. M. Annaswamy. "A Dynamic Regulation Market Mechanism for improved financial settlements in wholesale electricity markets". In: *2017 American Control Conference (ACC)*. ISSN: 2378-5861. May 2017, pp. 1425–1430. DOI: 10.23919/ACC.2017.7963153.

[60] E. G. Kardakos, C. K. Simoglou, and A. G. Bakirtzis. "Optimal bidding strategy in transmission-constrained electricity markets". en. In: *Electric Power Systems Research* 109 (Apr. 2014), pp. 141–149. ISSN: 03787796. DOI: 10.1016/j.epsr.2013.12.014. URL: https://linkinghub.elsevier.com/retrieve/pii/S037877961300343X (visited on 02/03/2021).

[61] M. D. Somma, G. Graditi, and P. Siano. "Optimal Bidding Strategy for a DER Aggregator in the Day-Ahead Market in the Presence of Demand Flexibility". In: *IEEE Transactions on Industrial Electronics* 66.2 (Feb. 2019). Conference Name: IEEE Transactions on Industrial Electronics, pp. 1509–1519. ISSN: 1557-9948. DOI: 10.1109/TIE.2018.2829677.

[62] T. Athay, R. Podmore, and S. Virmani. "A practical method for the direct analysis of transient stability". In: *IEEE Transactions on Power Apparatus and Systems* 2 (1979), pp. 573–584.

[63] U. C. Bureau. *Quick Facts: Boston City, Massachusetts*. available at https://www.census.gov/quickfacts/fact/table/bostoncitymassachusetts/PST120222 (accessed 23 April 2024). 2023.

[64] W. H. Kersting. *Radial distribution test feeders*. 1991.

[65] F. E. Postigo Marcos, C. Mateo Domingo, T. Gomez San Roman, B. Palmintier, B.-M. Hodge, V. Krishnan, F. de Cuadra García, and B. Mather. "A review of power distribution test feeders in the United States and the need for synthetic representative networks". In: *Energies* 10.11 (2017), p. 1896.

[66] D. Olsen, M. Sohn, M. A. Piette, and S. Kiliccote. *Demand Response Availability Profiles for California in the Year 2020*. en. Tech. rep. LBNL–1004414, 1341727. Nov. 2014, LBNL–1004414, 1341727. DOI: 10.2172/1341727. URL: http://www.osti.gov/servlets/purl/1341727/ (visited on 09/18/2021).

[67] ISO-NE. "Pricing Reports". In: (2021). URL: https://www.iso-ne.com/isoexpress/web/reports/pricing/-/tree/lmps-rt-five-minute-final.

[68] Federal Energy Regulatory Commission. *Payment for Reactive Power- Commission Staff Report AD14-7*. Apr. 22, 2014. URL: https://www.ferc.gov/sites/default/files/2020-05/04-11-14-reactive-power.pdf.

[69] C. Eid, J. Reneses Guillén, P. Frías Marín, and R. Hakvoort. "The economic effect of electricity net-metering with solar PV: Consequences for network cost recovery, cross subsidies and policy objectives". In: *Energy Policy* 75 (Dec. 2014), pp. 244–254. ISSN: 0301-4215. DOI: 10.1016/J.ENPOL.2014.09.011.

[70] A. Picciariello, C. Vergara, J. Reneses, P. Frías, and L. Söder. "Electricity distribution tariffs and distributed generation: Quantifying cross-subsidies from consumers to prosumers". In: *Utilities Policy* 37 (Dec. 2015), pp. 23–33. ISSN: 0957-1787. DOI: 10.1016/J.JUP.2015.09.007.

[71] F. Bu, Y. Yuan, Z. Wang, K. Dehghanpour, and A. Kimber. "A Time-Series Distribution Test System Based on Real Utility Data". en. In: *2019 North American Power Symposium (NAPS)*. Wichita, KS, USA: IEEE, Oct. 2019, pp. 1–6. ISBN: 978-1-72810-407-2. DOI: 10.1109/NAPS46351.2019.8999982. URL: https://ieeexplore.ieee.org/document/8999982/ (visited on 02/03/2021).

[72] Y. Tohidi, M. Farrokhseresht, and M. Gibescu. "A review on coordination schemes between local and central electricity markets". In: *2018 15th International Conference on the European Energy Market (EEM)*. IEEE. 2018, pp. 1–5.

[73] V. J. Nair, V. Venkataramanan, R. Haider, and A. M. Annaswamy. "A hierarchical local electricity market for a der-rich grid edge". In: *IEEE Transactions on Smart Grid* 14.2 (2022), pp. 1353–1366.

[74] S. Karagiannopoulos, C. Mylonas, P. Aristidou, and G. Hug. "Active distribution grids providing voltage support: The swiss case". In: *IEEE Transactions on Smart Grid* 12.1 (2020), pp. 268–278.

[75] D. B. Arnold, M. Sankur, R. Dobbe, K. Brady, D. S. Callaway, and A. Von Meier. "Optimal dispatch of reactive power for voltage regulation and balancing in unbalanced distribution systems". In: *2016 IEEE Power and Energy Society General Meeting (PESGM)*. IEEE. 2016, pp. 1–5.

[76] R. Haider, S. Baros, Y. Wasa, J. Romvary, K. Uchida, and A. M. Annaswamy. "Toward a retail market for distribution grids". In: *IEEE Transactions on Smart Grid* 11.6 (2020), pp. 4891–4905.

[77] D. Holmberg, F. Omar, and T. Roth. "Comparison of Two Approaches for Managing Building Flexibility Using Energy Market Prices". In: (2022).

[78] R. Haider, D. D'Achiardi, V. Venkataramanan, A. Srivastava, A. Bose, and A. M. Annaswamy. "Reinventing the utility for distributed energy resources: A proposal for retail electricity markets". In: *Advances in Applied Energy* 2 (2021), p. 100026.

[79] A. Potter, R. Haider, G. Ferro, M. Robba, and A. M. Annaswamy. "A reactive power market for the future grid". In: *Advances in Applied Energy* 9 (2023), p. 100114.

[80] L. Bai, J. Wang, C. Wang, C. Chen, and F. Li. "Distribution locational marginal pricing (DLMP) for congestion management and voltage support". In: *IEEE Transactions on Power Systems* 33.4 (2017), pp. 4061–4073.

[81] A. Papavasiliou. "Analysis of distribution locational marginal prices". In: *IEEE Transactions on Smart Grid* 9.5 (2017), pp. 4872–4882.

[82] G. Ferro, M. Robba, D. D'Achiardi, R. Haider, and A. M. Annaswamy. "A distributed approach to the Optimal Power Flow problem for unbalanced and mesh networks". In: *IFAC-PapersOnLine* 53.2 (2020), pp. 13287–13292.

[83] H. Mausser. "Normalization and other topics in multi-objective optimization". In: *Fields-MITACS Industrial Problems Workshop*. Citeseer. 2006, p. 89.

[84] A. Reuther, J. Kepner, C. Byun, S. Samsi, W. Arcand, D. Bestor, B. Bergeron, V. Gadepally, M. Houle, M. Hubbell, et al. "Interactive supercomputing on 40,000 cores for machine learning and data analysis". In: *2018 IEEE High Performance extreme Computing Conference (HPEC)*. IEEE. 2018, pp. 1–6.

[85] "Operational Coordination across Bulk Power, Distribution and Customer Systems Prepared for the Electricity Advisory Committee". In: (2019).

[86] *FERC Order No. 2222: Fact Sheet | Federal Energy Regulatory Commission*. URL: https://www.ferc.gov/media/ferc-order-no-2222-fact-sheet.

[87] *Performance-Based Regulation (PBR)*. URL: https://www.utc.wa.gov/performancebased.

[88] V. J. Nair and A. Annaswamy. "A game-theoretic, market-based approach to extract flexibility from distributed energy resources". In: *IFAC-PapersOnLine* 58.30 (2024), pp. 163–168.

[89] R. Haider, D. D'Achiardi, V. Venkataramanan, A. Srivastava, A. Bose, and A. M. Annaswamy. "Reinventing the utility for distributed energy resources: A proposal for retail electricity markets". In: *Advances in Applied Energy* 2 (May 2021), p. 100026. ISSN: 2666-7924. DOI: 10.1016/J.ADAPEN.2021.100026.

[90] W. Saad, Z. Han, H. V. Poor, and T. Basar. "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications". In: *IEEE signal processing magazine* 29.5 (2012), pp. 86–105.

[91] Z. M. Fadlullah, Y. Nozaki, A. Takeuchi, and N. Kato. "A survey of game theoretic approaches in smart grid". In: *2011 International Conference on Wireless Communications and Signal Processing (WCSP)*. IEEE. 2011, pp. 1–4.

[92] C. Eid, P. Codani, Y. Perez, J. Reneses, and R. Hakvoort. "Managing electric flexibility from Distributed Energy Resources: A review of incentives for market design". In: *Renewable and Sustainable Energy Reviews* 64 (Oct. 2016), pp. 237–247. ISSN: 1364-0321. DOI: 10.1016/J.RSER.2016.06.008.

[93] S. Maharjan, Q. Zhu, Y. Zhang, S. Gjessing, and T. Başsar. "Dependable demand response management in the smart grid: A stackelberg game approach". In: *IEEE Transactions on Smart Grid* 4.1 (2013), pp. 120–132. ISSN: 19493053. DOI: 10.1109/TSG.2012.2223766.

[94] B. Turdybek, M. Tostado-Véliz, S. A. Mansouri, A. Rezaee Jordehi, and F. Jurado. "A local electricity market mechanism for flexibility provision in industrial parks involving Heterogenous flexible loads". In: *Applied Energy* 359 (Apr. 2024), p. 122748. ISSN: 0306-2619. DOI: 10.1016/J.APENERGY.2024.122748.

[95] H. Saeian, T. Niknam, M. Zare, and J. Aghaei. "Coordinated optimal bidding strategies methods of aggregated microgrids: A game theory-based demand side management under an electricity market environment". In: *Energy* 245 (Apr. 2022), p. 123205. ISSN: 0360-5442. DOI: 10.1016/J.ENERGY.2022.123205.

[96] N. Li, L. Chen, and S. H. Low. "Optimal demand response based on utility maximization in power networks". In: *IEEE Power and Energy Society General Meeting* (2011). ISSN: 19449925. DOI: 10.1109/PES.2011.6039082.

[97] K. Anoh, S. Maharjan, A. Ikpehai, Y. Zhang, and B. Adebisi. "Energy Peer-to-Peer Trading in Virtual Microgrids in Smart Grids: A Game-Theoretic Approach". In: *IEEE Transactions on Smart Grid* 11.2 (Mar. 2020), pp. 1264–1275. ISSN: 19493061. DOI: 10.1109/TSG.2019.2934830.

[98] E. Nekouei, T. Alpcan, and D. Chattopadhyay. "Game-theoretic frameworks for demand response in electricity markets". In: *IEEE Transactions on Smart Grid* 6.2 (Mar. 2015), pp. 748–758. ISSN: 19493053. DOI: 10.1109/TSG.2014.2367494.

[99] L. Zhao, W. Zhang, H. Hao, and K. Kalsi. "A Geometric Approach to Aggregate Flexibility Modeling of Thermostatically Controlled Loads". In: *IEEE Transactions on Power Systems* 32.6 (Nov. 2017), pp. 4721–4731. ISSN: 08858950. DOI: 10.1109/TPWRS.2017.2674699.

[100] V. J. Nair, V. Venkataramanan, R. Haider, and A. M. Annaswamy. "A Hierarchical Local Electricity Market for a DER-rich Grid Edge". In: *IEEE Transactions on Smart Grid* (2022). ISSN: 19493061. DOI: 10.1109/TSG.2022.3174036.

[101] I. L. Glicksberg. "A further generalization of the Kakutani fixed theorem, with application to Nash equilibrium points". In: *Proceedings of the American Mathematical Society* 3.1 (1952), pp. 170–174.

[102] M. C. Cohen, A. N. Elmachtoub, and X. Lei. "Price Discrimination with Fairness Constraints". In: *Management Science* 68.12 (Dec. 2022), pp. 8536–8552. ISSN: 15265501. DOI: 10.1287/MNSC.2022.4317/ASSET/IMAGES/LARGE/MNSC.2022.4317F3.JPEG. URL: https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2022.4317.

[103] NASEM. *Enhancing the Resilience of the Nation's Electricity System*. Washington, DC: The National Academies Press, 2017. ISBN: 978-0-309-46307-2. DOI: 10.17226/24836. URL: https://www.nap.edu/catalog/24836/enhancing-the-resilience-of-the-nations-electricity-system.

[104] *CHERNOVITE's PIPEDREAM Malware Targeting Industrial Control Systems (ICS)*. available at https://www.dragos.com/blog/industry-news/chernovite-pipedream-malware-targeting-industrial-control-systems/ (accessed 10 April 2024). 2022.

[105] T. Nguyen, S. Wang, M. Alhazmi, M. Nazemi, A. Estebsari, and P. Dehghanian. "Electric power grid resilience to cyber adversaries: State of the art". In: *IEEE Access* 8 (2020), pp. 87592–87608.

[106] D. E. Whitehead, K. Owens, D. Gammel, and J. Smith. "Ukraine cyber-induced power outage: Analysis and practical mitigation strategies". In: *Proceedings of the 2017 70th Annual Conference for Protective Relay Engineers (CPRE)*. IEEE. 2017, pp. 1–8.

[107] NASEM. *The Role of Net Metering in the Evolving Electricity System*. Washington, DC: The National Academies Press, 2023. ISBN: 978-0-309-69331-8. DOI: 10.17226/26704. URL: https://nap.nationalacademies.org/catalog/26704/the-role-of-net-metering-in-the-evolving-electricity-system.

[108] S. M. Dibaji, M. Pirani, D. B. Flamholz, A. M. Annaswamy, K. H. Johansson, and A. Chakrabortty. "A systems and control perspective of CPS security". In: *Annual Reviews in Control* 47 (2019), pp. 394–411.

[109] C. D. Brummitt, R. M. D'Souza, and E. A. Leicht. "Suppressing cascades of load in interdependent networks". In: *Proceedings of the National Academy of Sciences* 109.12 (2012), E680–E689.

[110] Z. Li, M. Shahidehpour, F. Aminifar, A. Alabdulwahab, and Y. Al-Turki. "Networked microgrids for enhancing the power system resilience". In: *Proceedings of the IEEE* 105.7 (2017), pp. 1289–1310.

[111] Y. Li, X. Wei, Y. Li, Z. Dong, and M. Shahidehpour. "Detection of false data injection attacks in smart grid: A secure federated deep learning approach". In: *IEEE Transactions on Smart Grid* 13.6 (2022), pp. 4862–4872.

[112] T. Shekari, A. A. Cardenas, and R. Beyah. "MaDIoT 2.0: Modern High-Wattage IoT Botnet Attacks and Defenses". In: *Proceesdings of the 31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 3539–3556. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/shekari.

[113] P. Eder-Neuhauser, T. Zseby, J. Fabini, and G. Vormayr. "Cyber attack models for smart grid environments". In: *Sustainable Energy, Grids and Networks* 12 (2017), pp. 10–29.

[114] H. He and J. Yan. "Cyber-physical attacks and defences in the smart grid: a survey". In: *IET Cyber-Physical Systems: Theory & Applications* 1.1 (2016), pp. 13–27.

[115] V. J. Nair, P. Srivastava, and A. Annaswamy. "Enhancing power grid resilience to cyber-physical attacks using distributed retail electricity markets". In: *Proceedings of the 2024 IEEE/ACM International Conference on International Conference on Cyber-Physical Systems (ICCPS)*. IEEE. 2024.

[116] V. J. Nair and A. Annaswamy. "Local retail electricity markets for distribution grid services". In: *2023 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE. 2023, pp. 32–39.

[117] S. S. Amiri, M. Rahmani, and J. D. McDonald. "An updated review on distribution management systems within a smart grid structure". In: *2021 11th Smart Grid Conference (SGC)*. IEEE. 2021, pp. 1–5.

[118] K. A. Horowitz, Z. Peterson, M. H. Coddington, F. Ding, B. O. Sigrin, D. Saleem, S. E. Baldwin, B. Lydic, S. C. Stanfield, N. Enbar, et al. *An overview of distributed energy resource (DER) interconnection: Current practices and emerging solutions*. Tech. rep. 2019.

[119] P. Srivastava, R. Haider, V. J. Nair, V. Venkataramanan, A. M. Annaswamy, and A. K. Srivastava. "Voltage regulation in distribution grids: A survey". In: *Annual Reviews in Control* 55 (2023), pp. 165–181. ISSN: 1367-5788. DOI: https://doi.org/10.1016/j.arcontrol.2023.03.008. URL: https://www.sciencedirect.com/science/article/pii/S1367578823000123.

[120] K. Zetter. *Inside the Cunning, Unprecedented Hack of Ukraine's Power Grid*. available at https://www.wired.com/2016/03/inside-cunningunprecedented-hack-ukraines-power-grid/ (accessed 11 April 2024). July 2018.

[121] J. Kennedy. *Dragonfly: Western energy sector targeted by sophisticated attack group*. available at https://www.symantec.com/blogs/threatintelligence/dragonfly-energy-sector-cyber-attacks (accessed 10 April 2024). 2017.

[122] R. M. Lee, M. J. Assante, and T. Conway. *ICS defense use case: Analysis of the cyber attack on the Ukrainian power grid*. Tech. rep. Electricity Information Sharing and Analysis Center, SANS ICS, 2016.

[123] M. Zeller. "Myth or reality – does the Aurora vulnerability pose a risk to my generator?" In: *64th Ann. Conf. for Protective Relay Engineers*. 2011, pp. 130–136.

[124] S. Soltan, P. Mittal, and H. V. Poor. "BlackIoT: IoT Botnet of High Wattage Devices Can Disrupt the Power Grid". In: *Proceedings of the 27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 15–32. ISBN: 978-1-939133-04-5. URL: https://www.usenix.org/conference/usenixsecurity18/presentation/soltan.

[125] B. Huang, A. A. Cardenas, and R. Baldick. "Not Everything is Dark and Gloomy: Power Grid Protections Against IoT Demand Attacks". In: *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1115–1132. ISBN: 978-1-939133-06-9. URL: https://www.usenix.org/conference/usenixsecurity19/presentation/huang.

[126] V. Venkataramanan, S. Kaza, and A. M. Annaswamy. "DER Forecast Using Privacy-Preserving Federated Learning". In: *IEEE Internet of Things Journal* 10.3 (2022), pp. 2046–2055.

[127] P. S. Sarker, S. K. Sadanandan, and A. K. Srivastava. "Resiliency Metrics for Monitoring and Analysis of Cyber-Power Distribution System With IoTs". In: *IEEE Internet of Things Journal* 10.9 (May 2023), pp. 7469–7479. ISSN: 23274662. DOI: 10.1109/JIOT.2022.3183180.

[128] C. M. Schneider, A. A. Moreira, J. S. Andrade Jr, S. Havlin, and H. J. Herrmann. "Mitigation of malicious attacks on networks". In: *Proceedings of the National Academy of Sciences* 108.10 (2011), pp. 3838–3841.

[129] D. K. Molzahn and I. A. Hiskens. "A Survey of Relaxations and Approximations of the Power Flow Equations". In: *Foundations and Trends R in Electric Energy Systems* 4.2 (2019), pp. 1–221. DOI: 10.1561/3100000012. URL: http://dx.doi.org/10.1561/3100000012.

[130] G. Ferro, M. Robba, D. D'Achiardi, R. Haider, and A. M. Annaswamy. "A distributed approach to the Optimal Power Flow problem for unbalanced and mesh networks". In: *IFAC-PapersOnLine* 53.2 (Jan. 2020), pp. 13287–13292. ISSN: 2405-8963. DOI: 10.1016/J.IFACOL.2020.12.159.

[131] H. Booth, D. Rike, and G. Witte. *The national vulnerability database (nvd): Overview*. Tech. rep. 2013.

[132] J. J. Romvary, G. Ferro, R. Haider, and A. M. Annaswamy. "A Proximal Atomic Coordination Algorithm for Distributed Optimization". In: *IEEE Transactions on Automatic Control* 67.2 (Feb. 2022), pp. 646–661. ISSN: 15582523. DOI: 10.1109/TAC.2021.3053907.

[133] G. Ferro, M. Robba, R. Haider, and A. M. Annaswamy. "A Distributed-Optimization-Based Architecture for Management of Interconnected Energy Hubs". In: *IEEE Transactions on Control of Network Systems* 9.4 (Dec. 2022), pp. 1704–1716. ISSN: 23255870. DOI: 10.1109/TCNS.2022.3165022.

[134] A. Aljanaby, E. Abuelrub, and M. Odeh. "A Survey of Distributed Query Optimization." In: *Int. Arab J. Inf. Technol.* 2.1 (2005), pp. 48–57.

[135] R. Haider, S. Baros, Y. Wasa, J. Romvary, K. Uchida, and A. M. Annaswamy. "Toward a Retail Market for Distribution Grids". In: *IEEE Transactions on Smart Grid* 11.6 (Nov. 2020), pp. 4891–4905. ISSN: 19493061. DOI: 10.1109/TSG.2020.2996565.

[136] R. Meyur, A. Vullikanti, S. Swarup, H. S. Mortveit, V. Centeno, A. Phadke, H. V. Poor, and M. V. Marathe. "Ensembles of realistic power distribution networks". In: *Proceedings of the National Academy of Sciences* 119.42 (2022), e2205772119.

[137] *First cyberattack on solar, wind assets revealed widespread grid weaknesses*. available at https://www.utilitydive.com/news/first-cyber-attack-on-solar-wind-assets-revealed-widespread-grid-weaknesse/ (accessed 14 April 2024). 2019.

[138] P. Kundur. *Power system stability*. Vol. 10. CRC Press New York, 2007, pp. 7–1.

[139] T. D. Hardy, B. Palmintier, P. L. Top, D. Krishnamurthy, and J. C. Fuller. "HELICS: A Co-Simulation Framework for Scalable Multi-Domain Modeling and Analysis". In: *IEEE Access* 12 (2024), pp. 24325–24347. DOI: 10.1109/ACCESS.2024.3363615. URL: https://ieeexplore.ieee.org/document/10424422.

[140] D. P. Chassin, K. Schneider, and C. Gerkensmeyer. "GridLAB-D: An open-source power systems modeling and simulation environment". In: *Proceedings of the 2008 IEEE/PES Transmission and Distribution Conference and Exposition*. Apr. 2008, pp. 1–5. DOI: 10.1109/TDC.2008.4517260.

[141] J. Kurtz and R. Hovsapian. *ARIES: Advanced Research on Integrated Energy Systems Research Plan*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2021.

[142] D. Olsen, M. Sohn, M. A. Piette, and S. Kiliccote. *Demand Response Availability Profiles for California in the Year 2020*. Tech. rep. Berkeley, CA (United States): Lawrence Berkeley National Laboratory (LBNL), Nov. 2014. DOI: 10.2172/1341727. URL: http://www.osti.gov/servlets/purl/1341727/.

[143] J. Seuss, M. J. Reno, R. J. Broderick, and S. Grijalva. "Improving distribution network PV hosting capacity via smart inverter reactive power support". In: *IEEE Power and Energy Society General Meeting* 2015-September (Sept. 2015). ISSN: 19449933. DOI: 10.1109/PESGM.2015.7286523.

[144] S. M. Ismael, S. H. Abdel Aleem, A. Y. Abdelaziz, and A. F. Zobaa. "State-of-the-art of hosting capacity in modern power systems with distributed generation". In: *Renewable Energy* 130 (Jan. 2019), pp. 1002–1020. ISSN: 0960-1481. DOI: 10.1016/J.RENENE.2018.07.008.

[145] E. Mulenga, M. H. Bollen, and N. Etherden. *A review of hosting capacity quantification methods for photovoltaics in low-voltage distribution grids*. Feb. 2020. DOI: 10.1016/j.ijepes.2019.105445.

[146] M. Rylander, L. Rogers, and J. Smith. "Distribution feeder hosting capacity: What matters when planning for DER?" In: *Elect. Power Res. Inst., Inc.* (2015).

[147] A. Makhalfih, M. J. Hossain, C. MacAna, and H. R. Pota. "Distributed Energy Resources Hosting Capacity Assessment - An Industry Perspective". In: *2023 IEEE International Conference on Energy Technologies for Future Grids, ETFG 2023* (2023). DOI: 10.1109/ETFG55873.2023.10407889.

[148] A. K. Jain, K. Horowitz, F. Ding, N. Gensollen, B. Mather, and B. Palmintier. "Quasi-Static Time-Series PV Hosting Capacity Methodology and Metrics". In: *2019 IEEE Power and Energy Society Innovative Smart Grid Technologies Conference, ISGT 2019* (Feb. 2019). DOI: 10.1109/ISGT.2019.8791569.

[149] H. H. Mousa, K. Mahmoud, and M. Lehtonen. "A Comprehensive Review on Recent Developments of Hosting Capacity Estimation and Optimization for Active Distribution Networks". In: *IEEE Access* 12 (2024), pp. 18545–18593. ISSN: 21693536. DOI: 10.1109/ACCESS.2024.3359431.

[150] X. Geng, L. Tong, A. Bhattacharya, B. Mallick, and L. Xie. "Probabilistic Hosting Capacity Analysis via Bayesian Optimization". In: *IEEE Power and Energy Society General Meeting* 2021-July (2021). ISSN: 19449933. DOI: 10.1109/PESGM46819.2021.9637907.

[151] M. S. S. Abad, J. Ma, D. Zhang, A. S. Ahmadyar, and H. Marzooghi. "Probabilistic Assessment of Hosting Capacity in Radial Distribution Systems". In: *IEEE Transactions on Sustainable Energy* 9.4 (Oct. 2018), pp. 1935–1947. ISSN: 19493029. DOI: 10.1109/TSTE.2018.2819201.

[152] H. Al-Saadi, R. Zivanovic, and S. F. Al-Sarawi. "Probabilistic Hosting Capacity for Active Distribution Networks". In: *IEEE Transactions on Industrial Informatics* 13.5 (Oct. 2017), pp. 2519–2532. ISSN: 15513203. DOI: 10.1109/TII.2017.2698505.

[153] K. Lee, P. Zhao, A. Bhattacharya, B. K. Mallick, and L. Xie. "An Active Learning-Based Approach for Hosting Capacity Analysis in Distribution Systems". In: *IEEE Transactions on Smart Grid* 15.1 (Jan. 2024), pp. 617–626. ISSN: 19493061. DOI: 10.1109/TSG.2023.3290414.

[154] S. Zhan, T. Gu, W. van den Akker, W. Brus, A. van der Molen, and J. Morren. "Towards Congestion Management in Distribution Networks: a Dutch Case Study on Increasing Heat Pump Hosting Capacity". In: *IET Conference Proceedings*. Vol. 2022. 27. Institution of Engineering and Technology, 2022, pp. 364–369. ISBN: 9781839538513. DOI: 10.1049/icp.2023.0128.

[155] U. E. I. Administration. *Electric Power Annual - Residential Energy Consumption*. Accessed: 2025-02-09. 2024. URL: https://www.eia.gov/electricity/annual/.

[156] I. Power and E. Society. *IEEE 123-Bus Test Feeder*. Accessed: 2025-02-10. 2025. URL: https://site.ieee.org/pes-testfeeders/resources/.

[157] D. K. Molzahn, F. Dörfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei. "A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems". In: *IEEE Transactions on Smart Grid* 8.6 (Nov. 2017), pp. 2941–2962. ISSN: 19493053. DOI: 10.1109/TSG.2017.2720471.

[158] B. M. Sanandaji, T. L. Vincent, and K. Poolla. "Ramping Rate Flexibility of Residential HVAC Loads". In: *IEEE Transactions on Sustainable Energy* 7.2 (Apr. 2016), pp. 865–874. ISSN: 19493029. DOI: 10.1109/TSTE.2015.2497236.

[159] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "Aggregate flexibility of thermostatically controlled loads". In: *IEEE Transactions on Power Systems* 30.1 (Jan. 2015), pp. 189–198. ISSN: 08858950. DOI: 10.1109/TPWRS.2014.2328865.

[160] S. Samsi et al. "The MIT Supercloud Dataset". In: *2021 IEEE High Performance Extreme Computing Conference, HPEC 2021* (2021). DOI: 10.1109/HPEC49654.2021.9622850.

[161] V. Sakalkar et al. "Data center power oversubscription with a medium voltage power plane and priority-aware capping". In: *International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS*. Association for Computing Machinery, Mar. 2020, pp. 497–511. ISBN: 9781450371025. DOI: 10.1145/3373376.3378533.

[162] R. Torquato, D. Salles, C. O. Pereira, P. C. M. Meira, and W. Freitas. "A Comprehensive Assessment of PV Hosting Capacity on Low-Voltage Distribution Systems". In: *IEEE Transactions on Power Delivery* 33.2 (Apr. 2018), pp. 1002–1012. ISSN: 08858977. DOI: 10.1109/TPWRD.2018.2798707.

[163] G. Castillo Jr. "Using Electric Vehicles for Grid Services: Capacity Available and Applications for Electric Utility Commercialization". In: (2022). URL: https://dspace.mit.edu/handle/1721.1/146643.

[164] A. J. Conejo, M. Carrión, and J. M. Morales. *Decision Making Under Uncertainty in Electricity Markets*. Tech. rep. URL: www.springer.com/series/6161.

[165] B. Mohammadi-Ivatloo, H. Zareipour, N. Amjady, and M. Ehsan. "Application of information-gap decision theory to risk-constrained self-scheduling of GenCos". In: *IEEE Transactions on Power Systems* 28.2 (2013), pp. 1093–1102. ISSN: 08858950. DOI: 10.1109/TPWRS.2012.2212727.

[166] Z. Tan, G. Wang, L. Ju, Q. Tan, and W. Yang. "Application of CVaR risk aversion approach in the dynamical scheduling optimization model for virtual power plant connected with wind-photovoltaic-energy storage system with uncertainties and demand response". In: *Energy* 124 (Apr. 2017), pp. 198–213. ISSN: 0360-5442. DOI: 10.1016/J.ENERGY.2017.02.063.

[167] O. Dowson and L. Kapelevich. "SDDP.jl: A Julia Package for Stochastic Dual Dynamic Programming". In: *https://doi-org.libproxy.mit.edu/10.1287/ijoc.2020.0987* 33.1 (Aug. 2020), pp. 27–33. ISSN: 15265528. DOI: 10.1287/IJOC.2020.0987. URL: https://pubsonline-informs-org.libproxy.mit.edu/doi/abs/10.1287/ijoc.2020.0987.

[168] J. L. Pulsipher, W. Zhang, T. J. Hongisto, and V. M. Zavala. "A unifying modeling abstraction for infinite-dimensional optimization". In: *Computers and Chemical Engineering* 156 (Jan. 2022). ISSN: 00981354. DOI: 10.1016/j.compchemeng.2021.107567.

[169] Z. Zhou, J. Zhang, P. Liu, Z. Li, M. C. Georgiadis, and E. N. Pistikopoulos. "A two-stage stochastic programming model for the optimal design of distributed energy systems". In: *Applied Energy* 103 (2013), pp. 135–144. ISSN: 03062619. DOI: 10.1016/j.apenergy.2012.09.019.

[170] J. F. Benders. "Partitioning procedures for solving mixed-variables programming problems". In: *Numerische Mathematik* 4.1 (1962), pp. 238–252. DOI: 10.1007/BF01386316.

[171] R. M. V. Slyke and R. J.-B. Wets. "L-shaped linear programs with applications to optimal control and stochastic programming". In: *SIAM Journal on Applied Mathematics* 17.4 (1969), pp. 638–663. DOI: 10.1137/0117061.

[172] G. Laporte and F. V. Louveaux. "The integer L-shaped method for stochastic integer programs with complete recourse". In: *Operations Research Letters* 13.3 (1993), pp. 133–142. DOI: 10.1016/0167-6377(93)90038-8.

[173] A. M. Geoffrion. "Generalized Benders decomposition". In: *Journal of Optimization Theory and Applications* 10.4 (1972), pp. 237–260. DOI: 10.1007/BF00934810.

[174] S. N. Parragh, F. Tricoire, and W. Gutjahr. "A branch-and-Benders-cut algorithm for a bi-objective stochastic facility location problem". In: *OR Spectrum* 42.3 (2020), pp. 737–765. DOI: 10.1007/s00291-020-00616-7. URL: https://link.springer.com/article/10.1007/s00291-020-00616-7.

[175] R. T. Rockafellar and R. J.-B. Wets. "A new scenario decomposition method for large-scale stochastic optimization". In: *Operations Research* 43.3 (1995), pp. 477–490. DOI: 10.1287/opre.43.3.477. URL: https://pubsonline.informs.org/doi/abs/10.1287/opre.43.3.477.

[176] J. Zou, S. Ahmed, and X. A. Sun. "Stochastic dual dynamic integer programming". In: *Mathematical Programming* 175.1-2 (May 2019), pp. 461–502. ISSN: 14364646. DOI: 10.1007/S10107-018-1249-5/TABLES/5. URL: https://link-springer-com.libproxy.mit.edu/article/10.1007/s10107-018-1249-5.

[177] H. Dai, Y. Xue, Z. Syed, D. Schuurmans, and B. Dai. *NEURAL STOCHASTIC DUAL DYNAMIC PROGRAMMING*. Tech. rep.

[178] V. Guigues, M. A. Lejeune, and W. Tekaya. "Regularized stochastic dual dynamic programming for convex nonlinear optimization problems". In: *Optimization and Engineering* 21.3 (Sept. 2020), pp. 1133–1165. ISSN: 15732924. DOI: 10.1007/S11081-020-09511-0/FIGURES/6. URL: https://link.springer.com/article/10.1007/s11081-020-09511-0.

[179] S. Zhang and X. A. Sun. "Stochastic dual dynamic programming for multistage stochastic mixed-integer nonlinear optimization". In: *Mathematical Programming* 196.1-2 (Nov. 2022), pp. 935–985. ISSN: 14364646. DOI: 10.1007/S10107-022-01875-8/FIGURES/1. URL: https://link.springer.com/article/10.1007/s10107-022-01875-8.

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