University of California
Santa Barbara

Online Optimization and Learning for Sustainable Cyber-Human-Physical Systems

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

Doctor of Philosophy
in
Electrical and Computer Engineering

by

Nathaniel Surman Tucker

Committee in charge:

Professor Mahnoosh Alizadeh, Chair
Professor Ramtin Pedarsani
Professor Jason Marden
Professor Hamed Mohsenian-Rad
Professor Ranjit Deshmukh

September 2022

The Dissertation of Nathaniel Surman Tucker is approved.

Professor Ramtin Pedarsani

Professor Jason Marden

Professor Hamed Mohsenian-Rad

Professor Ranjit Deshmukh

Professor Mahnoosh Alizadeh, Committee Chair

June 2022

Online Optimization and Learning for Sustainable Cyber-Human-Physical Systems

Nathaniel Surman Tucker

To my mother, Michelle Surman

Acknowledgements

I truly believe that my PhD is both a personal achievement as well as a collaborative achievement. Over the last 5 years, I consider myself extremely lucky to have some amazing people supporting me.

First and foremost, thank you to my family. Cody- Although our lives have diverged after leaving Portland, I believe we are closer now than we have ever been. Your support, help, and advice in pivotal moments has absolutely helped me get to where I am today. I enjoy every conversation with you and I hope our future holds many more adventures together. Dave- You are the man. There is nothing else I need to say. You have given me everything you can and asked for nothing in return. Your actions, work ethic, and attitude are second to none and I hope to be as great as you someday. Michelle- You have been, without a doubt, the most consistent, supportive, and loving person in my life. Without you supporting me, I would not have completed my PhD. You singlehandedly dragged me out of some of the lowest points in my life. I owe you everything.

Thank you to my advisor, Mahnoosh. You have been a great mentor and friend over the last 5 years. There have been several rough patches, especially during the midst of the pandemic when I was struggling in my personal life and academic work. You were patient, supportive, and flexible in order to get me back to being the productive and happy person that I am. Two specific things I want to mention: 1) Writing with you has been one of my favorite experiences. From analyzing and discussing the value of every single word in our manuscripts to rewriting paragraphs multiple times, you have vastly improved my writing skills. 2) Your straightforward advice on critical decisions in my academic life as well as professional life has been invaluable to me. You have always given me your true opinions and led me in the right directions. I cannot imagine having completed a PhD under any other advisor. I hope we remain life-long friends.

Thank you to my committee. Ramtin- I admire you as an academic and as a person. Your course was one of my top 2 favorite courses I took at UCSB. Thank you for all of the chats, stories, and fun times. Jason- Your dynamic programming/optimal control course was my other favorite course I took at UCSB and it significantly helped me with my research. Thank you for always inviting me to join in your lab's events as well. Hamed and Ranjit- Thank you both for your input on my work. I greatly value both of your opinions and knowledge and I am grateful to have you on my committee.

Thank you to Gustavo Cezar. You supported me for multiple years and I thoroughly enjoyed the real-world problems that we worked on together.

Last, thank you to my friends and fellow ECE students. It is all of you who made my time at UCSB fun and enjoyable. Chatting about life, walking along the bluffs, and learning about each others' research were some of the small things that I consider the most memorable. So many of you are now some of my closest friends, and I hope to remain life-long friends with all of you.

Curriculum Vitæ

Nathaniel Surman Tucker

Education

2022	Ph.D. in Electrical and Computer Engineering (Expected), University of California, Santa Barbara.
2017	M.S. in Electrical Engineering, Santa Clara University.
2016	B.S. in Electrical Engineering and Computer Engineering, Santa Clara University.

Publications

Journal Publications

1. N. Tucker and M. Alizadeh, "An Online Scheduling Algorithm for a Community Energy Storage System," *IEEE Transactions on Smart Grid*, 2022, [1].
2. A. Moradipari, N. Tucker, and M. Alizadeh, "Mobility-Aware Electric Vehicle Fast Charging Load Models with Geographical Price Variations," *IEEE Transactions on Transportation Electrification*, 2020 [2].
3. N. Tucker, A. Moradipari, and M. Alizadeh, "Constrained Thompson Sampling for Real-Time Electricity Pricing with Grid Reliability Constraints," *IEEE Transactions on Smart Grid*, 2020, [3].
4. N. Tucker and M. Alizadeh, "An Online Admission Control Mechanism for Electric Vehicles at Public Parking Infrastructures," *IEEE Transactions on Smart Grid*, 2019, [4].

Proceedings and Refereed Conferences

1. N. Tucker and M. Alizadeh, "An Online Optimization Framework for Electric Vehicle Smart Charging at Large-Scale Facilities," *Submitted to IEEE SmartGridComm*, 2022.
2. N. Tucker, G. Cezar, and M. Alizadeh, "Real-Time Electric Vehicle Smart Charging at Workplaces: A Real-World Case Study," *IEEE Power and Energy Society General Meeting*, 2022, [5].
3. A. Moradipari, N. Tucker, T. Zhang, G. Cezar, and M. Alizadeh, "Mobility-Aware Smart Charging of Electric Bus Fleets," *IEEE Power and Energy Society General Meeting*, 2020, [6].
4. N. Tucker, B. Turan, and M. Alizadeh, "Online Charge Scheduling for Electric Vehicles in Autonomous Mobility on Demand Fleets," *IEEE Intelligent Transportation Systems Conference*, 2019, [7].

5. B. Turan, N. Tucker and M. Alizadeh, "Smart Charging Benefits in Autonomous Mobility on Demand Systems," *IEEE Intelligent Transportation Systems Conference*, 2019, [8].
6. N. Tucker, B. Ferguson, and M. Alizadeh, "An Online Pricing Mechanism for Electric Vehicle Parking Assignment and Charge Scheduling," *IEEE American Control Conference*, 2019, [9].
7. N. Tucker and M. Alizadeh, "Online Pricing Mechanisms for Electric Vehicle Management at Workplace Charging Facilities," *Allerton Conference on Communication, Control, and Computing*, 2018, [10].
8. N. Tucker and M. Khanbaghi, "Jump Linear Quadratic Control for Energy Management of a Nanogrid," *IEEE American Control Conference*, 2018, [11].

Awards/Honors

1. UCSB Electrical and Computer Engineering Department, Dissertation Fellowship, 2022
2. UCSB Electrical and Computer Engineering Department, Graduate Student Spotlight, 2021
3. UCSB Institute for Energy Efficiency, Excellence in Research Fellowship, 2019
4. UCSB Center for Control, Dynamical Systems, and Computation, Graduate Student Spotlight, 2019
5. Santa Clara University Electrical Engineering Department, Outstanding Graduate Student Researcher, 2017
6. Santa Clara University Electrical Engineering Department, Outstanding Graduating Senior, 2016

Abstract

Online Optimization and Learning for Sustainable Cyber-Human-Physical Systems

by

Nathaniel Surman Tucker

Due to the increasing connectivity of modern technology and extensive data availability, large-scale societal systems have massive potential to improve reliability, reduce operational costs, ensure safety, and decrease their carbon footprint. However, the increasing complexity of these systems also brings forth new challenges that traditional engineering fields are not well equipped to handle. Rather, interdisciplinary solutions from many fields including optimization, control theory, communications, signal processing, economics, power engineering, and transportation engineering are required. Specifically, my research has focused on the upcoming Cyber-Human-Physical-Systems (CHPS) within this domain (e.g., transportation systems, power grids, and smart cities with cooperative building communities). A Cyber-Human-Physical System is any real-world system that involves physical components that are controlled by both algorithms and human input. For example, these systems operate by way of physical components and infrastructure, utilize algorithms for the control and optimization of the infrastructure, and rely on human input. Additionally, these societal systems seek to maximize profit or social welfare while operating subject to constraints and subject to inherently stochastic environments.

This thesis is focused on **developing optimization frameworks and machine learning strategies** to improve the operation of these complex modern societal infrastructure systems in uncertain environments. Namely, 1) leveraging recent advancements in online optimization for system scheduling with applications to electric vehicle (EV)

charging and community energy storage (CES) management, 2) using the machine learning framework called Thompson Sampling for the design of effective price signals for an electricity aggregator passively learning customers' price sensitivities while running a load shaping program and providing theoretical safety guarantees on critical infrastructure constraints, and 3) optimizing real-world workplace EV charging in an online fashion and scheduling the charging/routing of a real electric bus fleet to minimize operational costs.

Contents

Chapter 1

Introduction

Across the world, societies are individually and collectively shifting away from traditional operating strategies to ones that prioritize sustainability. Two critical components of this paradigm shift are the electrification of the transportation sector and the modernization of the power grid.

In the transportation sector, there is currently a strong push to replace traditional internal-combustion-engine vehicles (ICEVs) with electric vehicles (EVs). The switch to EVs reduces the distributed CO2 emissions of ICEVs and allows for the usage of clean renewable energy. EVs have seen a drastic increase in their popularity and sales over the last 10 years. In 2012, there were a total of 120,000 EVs sold worldwide. In 2021, there were over 120,000 EVs sold each week[1]. While this increase is generally seen as a positive, there are new challenges arising from this large EV population. Specifically, the infrastructure needed to recharge these EVs must grow alongside the EV population and these charging infrastructure systems must be managed in order to maximize their benefits. For example, residential EV charging systems or large-scale EV fleet facilities must be optimized to maximize their effectiveness.

[1]Global EV Outlook, 2022, *International Energy Agency (IEA)*

Alongside the electrification of the transportation sector, power grid modernization is playing a key role in many societies' sustainable development goals. Traditionally, the power grid was set up as a one-way system, with remote generation plants providing power to our communities. However, recently there has been much development in distributed energy resources (DERs) that are in close proximity to loads and provide clean, renewable energy (e.g., solar and wind). While DERs are commonly seen as positive additions to the grid, the inherently stochastic nature of renewable energy sources brings forth new challenges. Currently, in order to mitigate the issues arising from intermittent renewable generation, energy storage systems are increasing in popularity to store clean energy to be used at a later time. Additionally, demand response (DR) programs are being used to shape how customers use electricity and increase the usage of renewable generation.

As these sustainable efforts continue to grow, so does the need for management strategies to make smart decisions to maximize the effectiveness of these systems. The goal of this thesis is to design, analyze, and test novel management strategies for these systems. Specifically, this thesis is focused on developing, analyzing, and testing optimization frameworks and machine learning strategies that aid in the electrification of the transportation sector and modernization of the power grid.

1.1 Motivations of this Thesis

Online (Real-Time) Optimization for System Scheduling

Many people rely on societal infrastructure systems in their day-to-day lives and can experience immense inconveniences due to poor management of these systems (e.g., losing air conditioning on a hot day because of a blackout, experiencing excessively long wait times for a ride-share because the vehicles in-service were routed poorly, or losing the opportunity to charge an electric vehicle (EV) at a workplace parking lot because all

the chargers were occupied). Because these societal systems are such a critical part of our modern world, designing effective management and scheduling algorithms for these large-scale systems is a difficult interdisciplinary endeavor. Modern large-scale societal infrastructure systems exist within our inherently stochastic world and must be managed in real-time; that is, they are subject to large amounts of uncertainty in their usage patterns as well as from exogenous factors that affect their operation. As such, the management algorithms of these systems must operate in real-time without knowledge of future events. This thesis considers several problem settings where real-time (online) optimization frameworks are required to design effective management and scheduling systems (e.g., scheduling and routing in ride-sharing/mobility-on-demand, strategic admission control and energy dispatch in electric vehicle charging equipped parking lots, and scheduling multiple users' charging and discharging profiles in a community energy storage system). Moreover, this thesis presents novel online scheduling systems in each of the aforementioned problem settings that make immediate decisions when new events occur, readily handle the inherent stochasticity of the users of the system, and always yield total welfare that is bound with respect to the offline optimal (i.e., bounding the total welfare of a schedule generated from the online management system to that of an omniscient offline schedule).

Safe Learning for Cyber-Human-Physical-Systems

The addition of humans into the control loop for large-scale societal systems yields many benefits but also results in numerous additional hurdles for reliable operation. As we know, day-to-day human activity is inherently stochastic; humans' responses to system signals can vary from day-to-day or even minute-to-minute, without any reason. This unknown and time-varying human response can lead to economic uncertainty for a system operator as well as reliability concerns for the system itself. Furthermore, the problem

of communicating with the users involved in these large-scale systems is a challenging one. Communication requires a framework and consistent user interaction; the framework could be costly and the user interaction could be unreliable. Moreover, many users might not be able to characterize their responses, and even if they could, they might not be willing to share this private information. With this in mind, it is evident that many large-scale societal systems with humans in the control loop require advanced methods to *passively* learn human responses only from historical data of past interactions. To this end, this thesis presents a novel machine learning strategy to learn electricity customers' preferences through a pricing mechanism designed to shape customers' electricity demand while upholding critical infrastructure constraints throughout the learning procedure.

Real-World Implementations

Additionally, this thesis presents several implementations of various optimization and control strategies for real-world systems. Specifically, the focus of the latter half of this thesis is on the following implementations:

- Real-Time (Online) Charge Scheduling via Model Predictive Control (MPC) for EVs at Google Campuses
- Route and Schedule Optimization via Mixed Integer Linear Programming (MILP) for Stanford's Marguerite Electric Shuttle System

1.2 Chapter Overviews

Chapter 1

Chapter 1 presents the motivation for this work, summary of main contributions, and chapter overviews.

Chapter 2

Chapter 2 presents results on online primal-dual optimization for electric vehicle smart charging.

Section 2.1 presents an online reservation system that allows electric vehicles (EVs) to park and charge at parking facilities equipped with electric vehicle supply equipment (EVSEs). We consider the case where EVs arrive in an online fashion and the facility coordinator must immediately make an admission or rejection decision as well as assign a specific irrevocable parking spot to each admitted EV. By means of strategic user admittance and smart charging, the objective of the facility coordinator is to maximize total user utility minus the operational costs of the facilities. We discuss an online pricing mechanism based on primal-dual methods for combinatorial auctions that functions as both an admission controller and a distributor of the facilities' limited charging resources. We analyze the online pricing mechanism's performance compared to the optimal offline solution and provide numerical results that validate the mechanism's performance for various test cases.

Section 2.2 presents an online charge scheduling strategy for fleets of autonomous-mobility-on-demand electric vehicles (AMoD EVs). We consider the case where vehicles complete trips and then enter a between-ride state throughout the day, with their information becoming available to the fleet operator in an online fashion. In the between-ride state, the vehicles must be scheduled for charging and then routed to their next passenger pick-up locations. Additionally, due to the unknown daily sequences of ride requests, the problem cannot be solved by any offline approach. As such, we study an online welfare maximization heuristic based on primal-dual methods that allocates limited fleet charging resources and rebalances the vehicles while avoiding congestion at charging facilities and pick-up locations. We discuss a competitive ratio result comparing the performance

of our online solution to the clairvoyant offline solution and provide numerical results highlighting the performance of our heuristic.

Chapter 3

Chapter 3 presents an online scheduling heuristic for Community Energy Storage (CES) systems. Specifically, a community energy storage (CES) system is studied that is shared by various electricity consumers who want to charge and discharge the CES throughout a given time span. We study the problem facing the manager of such a CES who must schedule the charging, discharging, and capacity reservations for numerous users. Moreover, we consider the case where requests to charge/discharge the CES arrive in an online fashion and the CES manager must immediately allocate charging power and energy capacity to fulfill the request or reject the request altogether. The objective of the CES manager is to maximize the total value gained by all of the users of the CES while accounting for the operational constraints of the CES. We discuss an algorithm titled CommunityEnergyScheduling that acts as a pricing mechanism based on online primal-dual optimization as a solution to the CES manager's problem. The online algorithm estimates the dual variables (prices) in real-time to allow for requests to be allocated or rejected immediately as they arrive. Furthermore, the proposed method promotes charging and discharging cancellations to reduce the CES's usage at popular times and is able to handle the inherent stochastic nature of the requests to charge/discharge stemming from randomness in users' net load patterns and weather uncertainties. Additionally, we are able to show that the algorithm is able to handle any adversarially chosen request sequence and will always yield total welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal welfare.

Chapter 4

Chapter 4 presents a safe learning framework for an electricity aggregator running a load shaping program. Specifically, the problem of an electricity aggregator attempting to learn customers' electricity usage models while implementing a load shaping program by means of broadcasting dispatch signals in real-time is studied. We adopt a multi-armed bandit problem formulation to account for the stochastic and unknown nature of customers' responses to dispatch signals. We propose a constrained Thompson sampling heuristic, Con-TS-RTP, as a solution to the load shaping problem of the electricity aggregator attempting to influence customers' usage to match various desired demand profiles (i.e., to reduce demand at peak hours, integrate more intermittent renewable generation, track a desired daily load profile, etc). The proposed Con-TS-RTP heuristic accounts for day-varying target load profiles (i.e., multiple target load profiles reflecting renewable forecasts and desired demand patterns) and takes into account the operational constraints of a distribution system to ensure that customers receive adequate service and to avoid potential grid failures. We provide a discussion on the regret bounds for our algorithm as well as a discussion on the operational reliability of the distribution system's constraints being upheld throughout the learning process.

Chapter 5

Chapter 5 presents work done to implement real-time optimization algorithms for EV smart charging at workplaces as well as routing/recharging for electric bus fleets.

Section 5.1 in Chapter 5 studies the joint route assignment and charge scheduling problem of a transit system dispatcher operating a fleet of electric buses in order to maximize solar energy integration and reduce energy costs. Specifically, we consider a complex bus transit system with preexisting routes, limited charging infrastructure,

limited number of electric buses, and time-varying electricity rates. We present a mixed integer linear program (MILP) that yields the minimal cost daily operation strategy for the fleet (i.e., route assignments and charging schedules using daily solar forecasts). We present numerical results from a real-world case study with Stanford University's Marguerite Shuttle (a large-scale electric bus fleet) to demonstrate the validity of our solution and highlight the significant cost savings compared to the status quo.

Section 5.2 in Chapter 5 presents a real-time smart charging algorithm for electric vehicles (EVs) at a workplace parking lot in order to minimize electricity cost from time-of-use electricity rates and demand charges while ensuring that the owners of the EVs receive adequate levels of charge. Notably, due to real-world constraints, our algorithm is agnostic to both the state-of-charge and the departure time of the EVs and uses scenario generation to account for each EV's unknown future departure time as well as certainty equivalent control to account for the unknown EV arrivals in the future. Real-world charging data from a Google campus in California allows us to build realistic models of charging demand for each day of the week. We then compare various results from our smart charging algorithm to the status quo for a two week period at a Google parking location.

Chapter 6

Chapter 6 presents conclusions of this thesis and future directions.

Chapter 2

Online Optimization for Electric Vehicle Smart Charging

2.1 Online Admission Control for Electric Vehicles at Large-Scale Facilities

2.1.1 Introduction

As of October 2018, one million plug-in electric vehicles (PEVs) have been sold in the United States [12]. Furthermore, sales have exceeded 20,000 units per month since May 2018 and these numbers are expected to continue trending upward beyond 2020 [12]. As such, coordinated charging strategies and charging infrastructure planning are paramount for ensuring the growing charging demand is satisfied in an environmentally responsible manner.

There has been a growing number of related papers that study EV smart charging

This work was supported in part by NSF under Grant 1847096, and in part by the California Energy Commission through SLAC. Solicitation: GFO-16-303. Agreement: EPC-16-057.

methods as well as infrastructure planning and investment analysis to encourage renewable energy usage in vehicle fleets, aggregate groupings, and parking facilities. For an overview, [13, 14, 15] provide in-depth reviews of smart charging technologies as well as societal and grid impacts. Investigations on the interactions between EV aggregations and the grid can be found in [16, 17]. Because smart charging has proven to benefit society, infrastructure investments must be made to support future charging implementations [18, 19, 20, 21, 22, 23]. Papers [23, 22, 20] study where to locate charging stations as well as how to effectively size the facilities. In [18], the authors study a planning framework for charging stations from the perspective of a social planner. Likewise, the authors in [19] study the design criteria for Fast Charging Stations (FCSs) based on mobility behaviors and paper [21] studies a planning scheme to maximize FCS usage and minimize infrastructure costs.

A critical but less studied problem is that coordinated charging at infrastructures can be heavily stunted if usage of the EVSEs is left uncontrolled [24]. Without EV routing within parking facilities, EVSEs at preferred locations (e.g., near an elevator) can become congested while other EVSEs are left empty. This limits the smart charging benefits as congested EVSEs are forced to charge one EV after another to satisfy charging demand. Similarly, without admission control the limited charging resources at facilities could be allocated to low priority users, (e.g., users with small charging demands, users with long sojourn times, or users who are willing to park elsewhere) therefore, occluding high priority users that arrive later in the day. As such, the focus of this chapter is to jointly perform admission control and smart charging, complementing previous work on coordinated charging and infrastructure planning.

Prior work in this area includes [25] where the authors investigate both First-Come-First-Serve (FCFS) and State-of-Charge (SoC) threshold policies for discerning which EVs are granted permission to use the EV charging infrastructure. Paper [26] studies an

online mechanism for the allocation of electricity to a population of EVs that have non-increasing marginal value for energy. Their setting allows for cancellation of reservations, which in our case is not allowed. In [27], an online algorithm for scheduling deferrable charging requests to balance the total value of vehicle owners and the total cost for providing charging service is studied, but they also allow for revocation of previously allocated resources. Paper [28] investigates an online auction that allows EV users to submit bids on their charging demand to the charging station and then the mechanism makes corresponding electricity allocation and pricing decisions. In this approach, users are expected to update their bids while parked at the charging station instead of only submitting one initial bid. Additionally, [29] studied a consensus approach for an online setting where selfish EVs compete for a limited amount of energy. An intelligent parking lot energy management system is studied in [30] to manage the scheduling of EVs to maximize charging for all EVs. Moreover, paper [31] formulates and analyzes a market model for deadline-differentiated pricing of deferrable electric power services; however, it does not focus on high levels of congestion or adversarial user valuations. Paper [32] presents a menu-based pricing scheme for allocating charge time within a facility and can lead to an efficient alternative approach to the EVSE reservation problem. The mechanism we present in this chapter focuses more on the congestion within facilities due to limited number of EVSEs and high demand with the objective of admitting highest priority users.

In most previous work, charging facilities are assumed to have traditional Single-Output-Single-Cable (SOSC) EVSEs. Recently, a more versatile charger has been gaining popularity: the Single-Output-Multiple-Cable (SOMC) EVSE which allows multiple EVs to be connected to the same charger, but only one EV receives charge at a time [33]. SOMC EVSEs can improve facility operations by allowing more flexibility in charge scheduling and decreasing idle plug-in time from traditional SOSC chargers. Further-

more, SOMC chargers eliminate the need for users to remove their vehicles once their charging session is complete. In any SOMC facility charging strategy, these idle EVs need to be accounted for; if not, the revenue of the facility will be reduced (our solution accounts for the times when EVs are charging *and* when they are idle). Utilizing SOMC EVSES, the authors of [33] study infrastructure investments, the authors of [24] study centralized online assignment methods such as Next-Fit and Worst-Fit for arriving EVs at a parking facility, and the authors of [10] study multiple online pricing heuristics for EV to EVSE allocation to increase smart charging capabilities.

In this work, we present an online pricing mechanism that functions as both an admission controller for parking facility access and a resource manager that optimizes smart charging strategies for vehicles admitted to the facility. The work presented in this section complements existing literature in the area and the main contributions are as follows:

1. The online mechanism readily accommodates multiple locations, multiple limited resources, operational costs, and renewable generation integration.

2. The online mechanism does not rely on fractional allocations or rounding methods to produce integer allocations in a computationally feasible manner and it never revokes previously made reservations.

3. The online mechanism readily handles the inherent stochasticity of the EVSE reservation problem including unknown sojourn times, unknown energy requests, and unknown user valuation functions.

4. The online mechanism is robust to adversarially chosen arrival sequences and always yields social welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal.

2.1.2 System Model

System Structure and User Characteristics

In this section, we describe our model for the EVSE reservation problem and user characteristics. We consider an EVSE reservation system that controls L dispersed parking facilities. Each parking facility $l \in L$ is equipped with M_l SOMC EVSEs each with C_l cables (i.e., each facility can park $M_l C_l$ EVs at any given time $t = 1, \ldots, T$ but only charge M_l). In addition to the cable constraints, each EVSE has a maximum power output constraint denoted by E_l that limits the amount of energy the EVSE can deliver in one time slot. To supply the EVSEs with electricity, each parking facility can procure energy from two sources: a rooftop solar generation system or the local distribution grid. We denote the available solar energy at facility l at time t with the variable $s_l(t) \in [0, S_l]$ where S_l is the maximum rating for facility l's rooftop system. Additionally, we denote $\pi_l(t)$ as the per unit price of electricity from the grid. Due to physical limits of the local transformer, we constrain facility l to procure no more than $G_l(t)$ units of energy from the grid at each time slot.

Each day, N EV owners submit requests to park and charge at various facilities. Each EV owner (user) is characterized by a set of attributes. Suppose user n wants to park and charge her EV. When she submits her reservation request at time t_n, she commits to arrive at one of her desired parking facilities $\{l_n\}$ at time t_n^- and to depart at t_n^+. Furthermore, user n receives value $\{v_{nl}\}$ if her EV receives h_n units of energy from facility l, meaning users have preferences for different facilities. With the aforementioned nomenclature, each arrival can be characterized by user 'type':

$$\theta_n = \{t_n^-, t_n^+, h_n, \{l_n\}, \{v_{nl}\}\} \in \Theta, \tag{2.1}$$

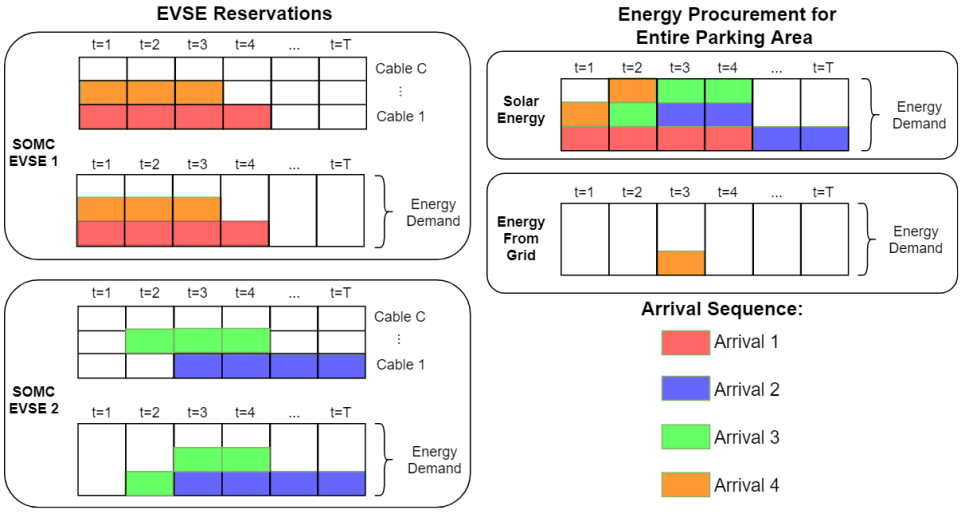

Figure 2.1: Example reservation schedule.

where Θ is the type space of all possible users. Fig. 2.1 presents an example allocation sequence with 4 arrivals and 2 EVSEs. Specifically, Fig. 2.1 showcases the fact that there are limited charging resources within the parking facility that need to be allocated to the arrivals. Each arriving EV needs an EVSE cable, an EVSE energy schedule, and the facility needs an energy procurement schedule from i) behind-the-meter solar, ii) the local distribution grid, or iii) a combination of solar and grid energy. The arriving vehicles enter the facility one-by-one and utilize the limited resources during their stay, affecting how future arrivals are allocated as seen in Fig. 2.1.

Offline Problem Formulation

To request a reservation, user n submits her user type θ_n to the EVSE reservation system. The EVSE reservation system creates a set of possible schedules that will fulfill user n's requirements. Namely, each possible schedule, or *option*, contains a cable reservation for the entire parking duration and the charge schedules that sum up to her desired charge amount. The reservation system generates these options for each facility within user n's desired facility set and then the option that yields the highest utility to

the user is selected.

We denote the set of options (potential schedules) as $\mathcal{O}_n$. Each option $o \in \mathcal{O}_n$ corresponds to a facility l_n, a cable reservation $c_{no}^{ml}(t)$, and a charge schedule $e_{no}^{ml}(t)$. The cable reservation $c_{no}^{ml}(t)$ takes values 0 or 1 depending if user n is assigned a cable from EVSE m at facility l at time t in option o. Similarly, $e_{no}^{ml}(t)$ takes values from a discrete set corresponding to the energy delivered to user n's EV. Through $e_{no}^{ml}(t)$, the EVSE reservation system is able to customize when each EV will receive charge and when it will be idle as well as the rate of charge. The set of feasible options for user n can be written as:

$$\{t_n^-, t_n^+, \{c_{no}^{ml}(t)\}, \{e_{no}^{ml}(t)\}, \{l_n\}, \{v_{nl}\}\}. \tag{2.2}$$

When deciding whether or not to admit user n and which option to allocate, the reservation system sets the binary variable x_{no}^{ml} to 1 if option o is chosen at EVSE m at facility l. Additionally, the reservation system computes payments $\hat{p}_{no}^{ml}$ for each option that the user pays if accepted. If a user is not admitted into any parking facilities, she receives zero value and parks in an auxiliary lot without EVSEs.

The EVSE reservation system keeps track of the allocated resources throughout the day. The variables $y_c^{ml}(t)$ and $y_e^{ml}(t)$ correspond to the allocated cables and energy, respectively, at EVSE m, facility l, at time t. Each facility also has to procure the energy needed by all the EVSEs within; therefore, the total energy needed by facility l at time t is denoted as $y_g^l(t)$. Equations (2.23)-(2.25) detail how each resource demand

is calculated:

$$y_c^{ml}(t) = \sum_{\mathcal{N},\mathcal{O}_n} c_{no}^{ml}(t) x_{no}^{ml}, \tag{2.3}$$

$$y_e^{ml}(t) = \sum_{\mathcal{N},\mathcal{O}_n} e_{no}^{ml}(t) x_{no}^{ml}, \tag{2.4}$$

$$y_g^{l}(t) = \sum_{\mathcal{N},\mathcal{O}_n,\mathcal{M}_l} e_{no}^{ml}(t) x_{no}^{ml}. \tag{2.5}$$

The energy procurement, $y_g^l(t)$, determines the operational cost of facility l:

$$f_g^l(y_g^l(t)) = \begin{cases} 0 & y_g^l(t) \in [0, s_l(t)\big) \\ \pi_l(t)(y_g^l(t) - s_l(t)) & y_g^l(t) \in [s_l(t), s_l(t) + G_l(t)] \\ +\infty & y_g^l(t) > s_l(t) + G_l(t). \end{cases} \tag{2.6}$$

The operational cost of the facility is zero while solar energy is available. Once the demand, $y_g^l(t)$, exceeds the available solar, energy is purchased from the grid. Once the demand exceeds the sum of available behind-the-meter solar energy and the transformer limit, no more energy can be procured.

With the system variables and equations defined, we can write the offline social welfare

maximization problem (assuming all users' information is known beforehand):

$$\max_{x} \sum_{\mathcal{N},\mathcal{O}_n,\mathcal{L},\mathcal{M}_l} v_{nl} x_{no}^{ml} - \sum_{\mathcal{T},\mathcal{L}} f_g^l(y_g^l(t)) \tag{2.7a}$$

subject to:

$$\sum_{\mathcal{O}_n,\mathcal{L},\mathcal{M}_l} x_{no}^{ml} \leq 1, \quad \forall\, n \tag{2.7b}$$

$$x_{no}^{ml} \in \{0,1\}, \quad \forall\, n,o,l,m \tag{2.7c}$$

$$y_c^{ml}(t) \leq C_l, \quad \forall\, l,m,t \tag{2.7d}$$

$$y_e^{ml}(t) \leq E_l, \quad \forall\, l,m,t \tag{2.7e}$$

and (2.23), (2.24), (2.25).

Moreover, the objective (3.5a) is to maximize the total *social welfare* of the system. This includes the utility gained by arrivals using the system minus the operational costs of the facilities (we note that users who are not admitted receive utility equal to zero). Constraints (2.30b)-(2.30e) respectively ensure at most one option is selected per user, the assignment variable is an integer, the cable demand does not exceed capacity, and the energy demand does not exceed capacity. Equations (2.23)-(2.25) sum the resource demands.

Temporarily relaxing the integrality constraint (2.30c) on x_{no}^{ml} allows us to find the Fenchel dual problem with dual variables u_n, $p_c^{ml}(t)$, $p_e^{ml}(t)$, and $p_g^l(t)$ [34]. In the following, the Fenchel conjugate of a function is given as:

$$f^*(p(t)) = \sup_{y(t)\geq 0} \big\{p(t)y(t) - f(y(t))\big\}. \tag{2.8}$$

Accordingly, the Fenchel dual of (3.5a)-(2.30e) can be written:

$$\min_{u,p} \sum_{\mathcal{N}} u_n + \sum_{\mathcal{T},\mathcal{L}} f_g^{l*}(p_g^l(t)) + \sum_{\mathcal{T},\mathcal{L},\mathcal{M}_l} \Big(f_c^{ml*}(p_c^{ml}(t)) + f_e^{ml*}(p_e^{ml}(t)) \Big) \tag{2.9a}$$

subject to:

$$u_n \geq \ v_{nl} - \sum_{\mathcal{T}} \Big(c_{no}^{ml}(t) p_c^{ml}(t) + e_{no}^{ml}(t) \big(p_e^{ml}(t) + p_g^l(t) \big) \Big) \quad \forall\, n, o, l, m \tag{2.9b}$$

$$u_n \geq \ 0, \qquad \forall\, n \tag{2.9c}$$

$$p_c^{ml}(t),\ p_e^{ml}(t),\ p_g^l(t) \geq \ 0, \quad \forall\, l, m, t, \tag{2.9d}$$

where $f^*(p(t))$ is the Fenchel conjugate for the limited resources' dual variables. The Fenchel conjugates for the capacity constraints can be written as:

$$f_c^{ml*}(p_c^{ml}(t)) = p_c^{ml}(t) C_l, \qquad p_c^{ml}(t) \geq 0 \tag{2.10}$$

$$f_e^{ml*}(p_e^{ml}(t)) = p_e^{ml}(t) E_l, \qquad p_e^{ml}(t) \geq 0. \tag{2.11}$$

Additionally, the Fenchel conjugate for the energy procurement operational cost function can be written as:

$$f_g^{l*}(p_g^l(t)) = \begin{cases} s_l(t) p_g^l(t), & p_g^l(t) < \pi_l(t) \\ (s_l(t) + G_l(t)) p_g^l(t) - G_l(t) \pi_l(t) & p_g^l(t) \geq \pi_l(t). \end{cases} \tag{2.12}$$

Admittance, Rejection, and Allocation Decisions

To determine how the EVSE reservation system decides whether or not to admit a user as well as which option to select if admitted, we make use of the Fenchel dual (2.32a)-(2.32d). Specifically, we examine the KKT conditions for constraint (2.32b). If a user is denied in the offline problem, u_n will be 0; otherwise, if a user is admitted, u_n will be positive. As such, the EVSE reservation system solves the following equation to determine user n's acceptance and her resource allocation:

$$u_n = \max\Big\{0, \max_{\mathcal{O}_n,\mathcal{L},\mathcal{M}_l}\big\{v_{nl} - \sum_{t\in[t_n^-,t_n^+]}\big(c_{no}^{ml}(t)p_c^{ml}(t) + e_{no}^{ml}(t)(p_e^{ml}(t) + p_g^l(t))\big)\big\}\Big\}. \quad (2.13)$$

We note that u_n corresponds to user n's utility from the EVSE reservation system. If admitted, the cable reservation and charge schedule chosen for user n correspond to the option o, EVSE m, and facility l that maximize the second term in equation (2.38). Furthermore, the dual variables $p_c^{ml}(t)$, $p_e^{ml}(t)$, and $p_g^l(t)$ correspond to prices that users must pay for cables, energy, and energy procurement. As such, the total payments corresponding to user n's different options are calculated as:

$$\hat{p}_{no}^{ml} = \sum_{\mathcal{T}}\Big(c_{no}^{ml}(t)p_c^{ml}(t) + e_{no}^{ml}(t)(p_e^{ml}(t) + p_g^l(t))\Big). \quad (2.14)$$

The EVSE reservation system is allocating options that maximize each user's utility with respect to the current marginal prices. Additionally, users receive non-negative utility for participating in the EVSE reservation system; therefore, we satisfy individual rationality constraints.

We would like to note that our proposed mechanism can also be used without any actual payments if users do not have the option of choosing their type (i.e., their types are preassigned). In a company, if users are assigned valuations (e.g., CEO has a high

value and regular employee has lower value, or someone with high charge level has lower value), then the prices do not have to be economic incentives. Rather they are used as dual variables that guide each user's allocation without any monetary transfer (i.e., each employee does not actually have to pay to use the infrastructure, but different employees have different valuations and the "shadow prices" allow for quick allocations).

The optimization problems presented in (3.5a)-(2.30e) and (2.32a)-(2.32d) assume complete knowledge of the arrivals beforehand. In practice, this is not the case; rather, users arrive and depart throughout the day. As such, the solution needs to be an online mechanism that can immediately allocate an arrival without knowledge of the future sequence of arrivals. Additionally, once a user has parked her car within a charging facility, she should not be asked to prematurely move her EV before her departure time. As such, the online mechanism should never revoke previous allocations. In the following, we discuss an online allocation mechanism that solves the EVSE reservation problem and meets the aforementioned design goals.

2.1.3 Online Allocation Mechanism

Online Marginal Prices

It is evident that the EVSE reservation problem requires an online solution. In many online problems, approximate dynamic programming (ADP) heuristics have good performance given accurate statistics even with large state-spaces [35, 36, 37, 38]. However, performance guarantees can be very hard to obtain for multi-stage decision making problems with complex action spaces over long time periods, and in our case, nonstationary arrival patterns and variable forecasts prohibit many traditional ADP techniques. As such, we present an online pricing mechanism that calculates the marginal prices on EVSE cables, energy, and generation based on a pricing heuristic, for which we provide

performance guarantees. Specifically, our EVSE reservation system updates the prices $p(t)$ heuristically as the amounts of allocated resources $y(t)$ evolve, but only based on past observations. The pricing scheme has two major goals: (1) to make sure that the marginal gain in welfare from an allocation is greater than the operational cost incurred to serve the allocation, and (2) to filter out low value users early to ensure there are adequate resources for higher value users later on. The structure of the marginal price functions we use is similar to that of [39], where the authors present a pricing framework for cloud-computing systems utilizing data centers with limited computation resources and server costs under an adversarial setting. For the limited number of cables at each EVSE, the proposed marginal payment function is given by:

$$p_c^{ml}(y_c^{ml}(t)) = \Big(\frac{L_c}{2\sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l})}\Big)\Big(\frac{2\sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l})U_c}{L_c}\Big)^{\frac{y_c^{ml}(t)}{C_l}}, \tag{2.15}$$

where $y_c^{ml}(t)$ is the current demand for the cables at EVSE m at location l at time t. Additionally, L_c and U_c are the lower and upper bounds on users' valuation per cable per unit of time, respectively:

$$L_c = \min_{\mathcal{N},\mathcal{O}_n,\mathcal{L},\mathcal{M}_l} \frac{v_{nl}}{\sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l}) \sum_{t\in[t_n^-,t_n^+]} c_{no}^{ml}(t)}, \tag{2.16a}$$

$$U_c = \max_{\mathcal{N},\mathcal{O}_n,\mathcal{L},\mathcal{M}_l,\mathcal{T}} \frac{v_{nl}}{c_{no}^{ml}(t)}, \quad c_{no}^{ml}(t) \neq 0. \tag{2.16b}$$

The pricing function for the EVSE energy units is the same as (2.39) with the exponent changed to E_l instead of C_l. Likewise, calculate L_e and U_e using $e_{no}^{nl}(t)$ in (2.41a) and (2.41b). Additionally, for the energy procurement resource, L_g and U_g are the same as L_e and U_e, respectively.

To explain this pricing function, set $y_c^{ml}(t) = 0$ and (2.39) outputs a price low enough that any user will be accepted (subject to L_c). Moreover, the pricing function (2.39)

yields low initial values to allow reservations early on. As more arrivals are admitted into the reservation system, congestion begins to affect the shared resources. To combat congestion and filter our low value arrivals, as $y_c^{ml}(t)$ increases, the prices from (2.39) increase exponentially. When $y_c^{ml}(t)$ is equal to the capacity of the limited resource, the marginal price is set high enough to reject all future arrivals to ensure that no resource will ever be overallocated (we assume L_c and U_c are known).

Designing a pricing function for energy procurement at each facility is more complicated than the cable pricing. Here, the cost to procure energy is piecewise linear and depends on the current solar generation and the transformer capacity. As such, we propose the pricing function:

$$p_g^l(y_g^l(t)) = \begin{cases} \left(\frac{L_g}{2R}\right)\left(\frac{2R\pi_l(t)}{L_g}\right)^{\frac{y_g^l(t)}{s_l(t)}}, & y_g^l(t) < s_l(t), \\ \left(\frac{L_g-\pi_l(t)}{2R}\right)\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)^{\frac{y_g^l(t)}{s_l(t)+G_l(t)}} + \pi_l(t), & y_g^l(t) \geq s_l(t), \end{cases} \tag{2.17}$$

$$\text{where } R = \sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l}).$$

Equation (2.43) is similar to the pricing function for the EVSE cables and energy; however, because procuring energy from the grid has non-zero cost, we need to ensure each user's payment is greater than the electricity cost needed to charge their vehicle. Additionally, when a facility's energy demand is less than the available solar, the marginal energy procurement price is reduced below the cost of electricity $\pi_l(t)$ to promote solar consumption.

Proposed Algorithm and Performance Guarantees

The admittance, allocation, and price update procedure for the EVSE reservation system is presented in Algorithm OnlineParkNCharge. When arrival n submits her

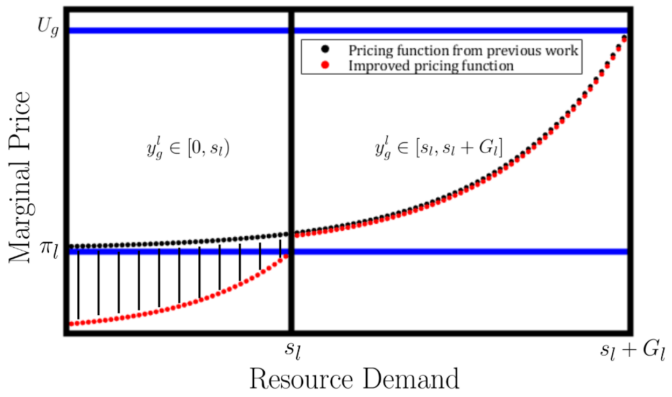

Figure 2.2: Pricing function for energy procurement. Shaded area: increase in users' utilities from updated pricing function.

request, the system generates the feasible options $\mathcal{O}_n$ that fulfill her demands. Then, the system accepts or rejects user n depending on her potential utility gain due to her valuation and the current resource prices (line 7). We note that line 9 requires solving an integer constrained maximization problem. This is not computationally burdensome as the optimization is solved for each individual vehicle at the time of arrival, with the potential utilities for each option can be calculated quickly via multiplication and addition. Then, any sorting method can be used to find the highest utility option. The algorithm updates the primal variables x_{no}^{ml} after each acceptance and rejection. The total resource demands are updated in line 11 if user n is accepted into the system. Similarly, the marginal resource prices are updated accordingly in line 12.

Next, we compare the total social welfare resulting from the online solution to the optimal offline solution. Specifically, an online mechanism is said to be α-competitive when the ratio of social welfare from the optimal offline solution to the social welfare from the mechanism is bounded by $\alpha \geq 1$. We extend a competitive ratio performance guarantee from [39] in Proposition 2.1.1. In the following, to ensure no user purchases too large of a fraction of the total available resource, we assume each user's resource demands are much smaller than the capacity limits.

Algorithm 1 OnlineParkNCharge

Input: $\mathcal{L}, \mathcal{M}_l, C_l, E_l, G_l, S_l, \pi_l, L_{c,e,g}, U_{c,e,g}$

Output: x, p

Define $f_g^l(y_g^l(t))$ according to (2.28).
Define the pricing functions $p(y(t))$ according to (2.39) and (2.43) for cables, energy, and generation.
Initialize $x_{no}^{ml} = 0$, $y^{ml}(t) = 0$, $u_n = 0$.
Initialize prices $p(0)$ according to (2.39) and (2.43).
Repeat for all N **users:**
User n submits θ_n, generate feasible charging options.
Update dual variable u_n according to (2.38).
if $u_n > 0$ **then**
    $(o^\star, m^\star, l^\star) = \arg\max_{\mathcal{L},\mathcal{M}_l,\mathcal{O}_n} \left\{ v_{nl} - \sum_{t\in[t_n^-,t_n^+]} \left(c_{no}^{ml}(t) p_c^{ml}(t) + e_{no}^{ml}(t)(p_e^{ml}(t) + p_g^l(t)) \right) \right\}$
    $\hat{p}_{no^\star}^{m^\star l^\star} = \sum_{t\in[t_n^-,t_n^+]} \left(c_{no^\star}^{m^\star l^\star}(t) p_c^{m^\star l^\star}(t) + e_{no^\star}^{m^\star l^\star}(t)(p_e^{m^\star l^\star}(t) + p_g^{l^\star}(t)) \right)$
    $x_{no^\star}^{m^\star l^\star} = 1$ and $x_{no}^{ml} = 0$ for all $(o, l, m) \neq (o^\star, l^\star, m^\star)$
    Update total demand $y(t)$ for cables, energy, and generation according to (2.23)-(2.25).
    Update marginal prices $p(t)$ according to (2.39) and (2.43).
else
    $x_{no}^{ml} = 0, \quad \forall\ \mathcal{L}, \mathcal{M}_l$ and $\mathcal{O}_n$.
end if
if $\exists o^\star, m^\star, l^\star$ and $x_{no^\star}^{m^\star l^\star} = 1$ **then**
    Accept user n and allocate cables and energy in parking location $l^\star$ at EVSE $m^\star$.
    Charge user n at $\hat{p}_{no^\star}^{m^\star l^\star}$.
else
    Send user n to auxiliary parking.
end if

Proposition 2.1.1 *The marginal pricing function* (2.43) *is* α_1*-competitive in social welfare when selling limited resources with the piecewise linear operational cost in* (2.28) *where*

$$\alpha_1 = 2 \max_{\mathcal{L},\mathcal{T}} \Big\{ \ln \Big(\frac{2 \sum_{\mathcal{L}} M_l (C_l + E_l + \frac{1}{M_l})(U_g - \pi_l(t))}{L_g - \pi_l(t)} \Big) \Big\}$$

with the assumption $\sum_{\mathcal{L}} M_l (C_l + E_l + \frac{1}{M_l}) \geq \lceil \frac{eL_g}{2 \max_{\mathcal{L},\mathcal{T}} \pi_l(t)} \rceil$.

Proof. The proof is in Appendix A.1.

Corollary 2.1.1 *If the final demand for energy procurement* $y_g^l(t)$ *for a given day is less than the available solar* $s_l(t)$*, the marginal pricing function* (2.43) *is* α_2*-competitive (*$\alpha_2 < \alpha_1$*) in social welfare when selling limited resources with the piecewise linear operational cost in* (2.28) *where*

$$\alpha_2 = 2 \max_{\mathcal{L},\mathcal{T}} \Big\{ \ln \Big(\frac{2 \sum_{\mathcal{L}} M_l (C_l + E_l + \frac{1}{M_l})(\pi_l(t))}{L_g} \Big) \Big\}.$$

In the previous proposition, the pricing function (2.43) relies on complete knowledge of the solar generation $s_l(t)$. If the system has inaccurate solar irradiation forecasts, the solar generation could be overestimated and resources are over-allocated resulting in infeasible solutions, which our online solution should avoid at all costs; or solar generation is underestimated and prices are set too high and the system rejects users that should otherwise be accepted. We analyze the case where we have a forecast of the solar generation each day in terms of a confidence interval. We do not assume a specific solar irradiance forecasting method; rather, we make use of a confidence interval for the potential solar each day as yearly solar irradiance recordings can provide minimum and maximum bounds for any given day. Additionally, our method assumes that these confidence regions are tightening as the day progresses. In this work, we assume that the solar

forecast for a future time t increases in accuracy as the current time $t_{current}$ approaches t. Specifically, the solar forecast takes the following form:

$$s_l(t) \in [\underline{s}_l(t, t_{current}), \overline{s}_l(t, t_{current})], \tag{2.18}$$

for $t = 1, \ldots, T$ and $1 \leq t_{current} \leq t$. Here, $s_l(t)$ is the actual solar generation at time t and the terms $\underline{s}_l(t, t_{current})$ and $\overline{s}_l(t, t_{current})$ are lower and upper bounds given by the forecast, respectively, at an earlier time $t_{current}$. We assume that the forecast is improving, specifically $\overline{s}_l(t, t_{current})$ is non-increasing and $\underline{s}_l(t, t_{current})$ is non-decreasing as $t_{current}$ approaches t.

To account for the dependence of the solar forecast on the current time, the marginal pricing function (2.43) is now written as $p_g^l(y_g^l(t), t_{current})$. To avoid possible infeasible allocations associated with overestimation of solar availability, we analyze the performance of pricing function (2.43) that conservatively uses the underestimate of the solar generation, $\underline{s}_l(t, t_{current})$, in Proposition 2.1.2. Fig. 2.3 shows how the pricing function (2.43) changes as the solar forecast improves.

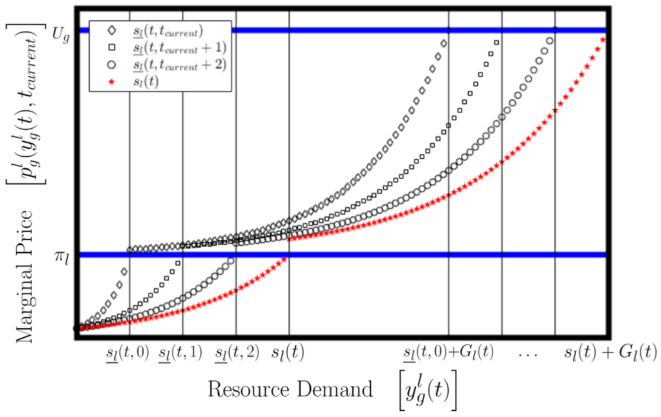

Figure 2.3: Pricing function with solar forecast $\underline{s}_l(t, t_{current})$.

Proposition 2.1.2 *The marginal pricing function* (2.43) *with an underestimate of solar*

generation, $\underline{s}_l(t, t_{current})$*, is* $\alpha_3 = 2\max_{\mathcal{L},\mathcal{T}}\left\{\alpha_g^l(t)\right\}$ *competitive in social welfare when selling limited resources with the operational cost in* (2.28) *where*

$$\alpha_g^l(t) = \begin{cases} \max\left\{\frac{\overline{s}_l(t,1)}{\underline{s}_l(t,1)}\ln(\frac{2R\pi_l(t)}{L_g}), \frac{\overline{s}_l(t,1)+G_l(t)}{\underline{s}_l(t,1)+G_l(t)}\ln(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)})\right\}, & \underline{s}_l(t,1) \neq 0, \\ \frac{\overline{s}_l(t,1)+G_l(t)}{G_l(t)}\ln(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}), & \underline{s}_l(t,1) = 0, \end{cases} \tag{2.19}$$

with the assumption $R \geq \lceil \frac{eL_g}{2\max_{\mathcal{L},\mathcal{T}}\pi_l(t)} \rceil$.

Proof. The proof is in Appendix A.2.

We would like to note the significance of Proposition 1 and Proposition 2 in the following. Namely, our competitive ratio results ensure the social welfare generated by the approximate online solution (that runs in real-time) cannot deviate too far from the social welfare generated by the oracle offline solution. The results in Proposition 1 and Proposition 2 ensure the online system, which acts without knowledge of future arrivals, performs within a constant factor of the offline/oracle system. Furthermore, the competitive ratios are worst case bounds on performance. That is, if this pricing scheme is used in a real scenario, even the social welfare generated with respect to an adversarially chosen arrival sequence is within the constant α of the optimal oracle solution.

2.1.4 Experimental Evaluation

The Case of Variable Arrival Patterns

In this section, we present a comparison of our online pricing mechanism against an online certainty equivalent controller (CEC) for a downtown parking facility to show the performance our mechanism under different arrival statistics. CEC is an approximate dynamic programming (ADP) technique that replaces all future uncertain quantities with some typical values, more specifically, the expected values. In this case, we assume

that the facility has arrival patterns following the distributions in Figure 2.4. In this

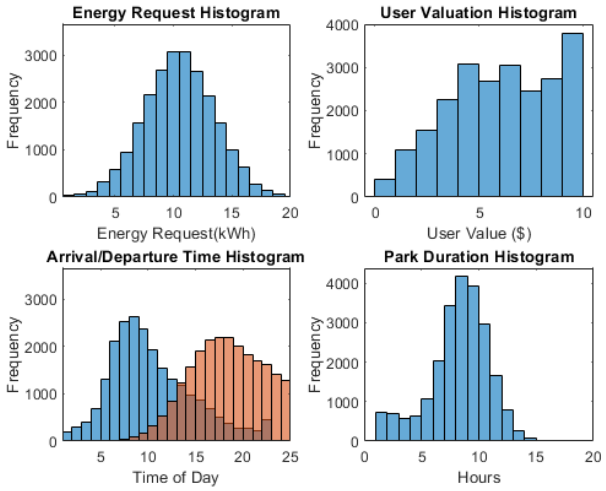

Figure 2.4: Top Left: Energy requests. Top Right: User valuations. Bottom Left: Arrival times (blue) and departure times (red). Bottom Right: Length of stay.

example, there is 1 parking facility with 5 EVSEs and 4 cables per EVSE (i.e., there are 20 parking slots available in each time period). The facility can purchase energy from the Los Angeles grid at a cost of \$0.127/kWh. Lastly, the facility has a 32 kW rooftop solar generation system that follows a production curve from an LA location in January 2018 [40]. We assume standard crystalline silicon panels with 14% system loss due to shading, wiring, connections, mismatch, and degradation. We simulated the CEC and pricing mechanism for 2 different arrival count distributions (as shown in Fig. 2.5). To demonstrate the value of adverserial solutions like ours in situations when the future is hard to predict, the distribution exhibits a larger variance under the second scenario.

From our results in Figure 2.5, it is evident that in the case of higher variance arrival patterns, our pricing mechanism outperforms a controller that is dependant on expected statistics. If a parking facility does not have consistent arrival statistics each day, our pricing mechanism performs better because it accounts for worst-case arrival patterns. Additionally, in the unimodal and bimodal arrival count cases, the minimum daily social welfare achieved by our pricing mechanism is significantly higher than the minimum of

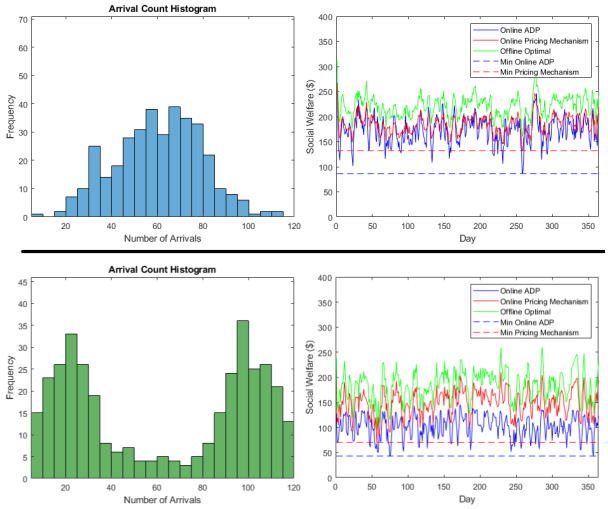

Figure 2.5: Top: Social welfare for facility with unimodal arrival count distribution. Bottom: Social welfare for facility with bimodal arrival count distribution.

the CEC ADP.

Comparison with First-Come-First-Serve Strategy

In this section, we present a comparison of our online pricing mechanism against the first-come-first-serve (FCFS) strategy that is commonplace in many EVSE equipped charging facilities. Specifically, we highlight the performance of our mechanism over varying demand levels to show the effectiveness of our mechanism when the infrastructure becomes congested. In the FCFS strategy, an arriving EV selects the closest available parking spot and begins charging immediately (without any controller directing them). In this test case, we assume the arrivals' energy requests, valuations, and durations follow the same statistics as in Fig. 2.4; however, in this simulation, we directly control the number of arrivals each day (to highlight different demand levels) and we limit the arrival times to 8:00am-10:00am and limit the departure times to after 10:00am (thus showcasing the performance of FCFS and our mechanism when large quantities of vehicles arrive in a short time period each morning). We assume the parking infrastructure has 15 EVSEs

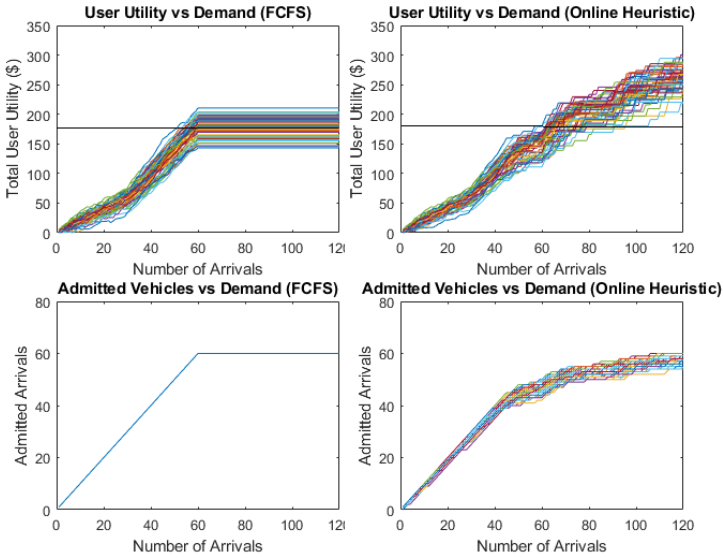

Figure 2.6: Top Left: FCFS's total user utility vs. number of daily arrivals (50 scenarios simulated). Top Right: Our online heuristic's total user utility vs. number of daily arrivals (50 scenarios simulated). Bottom Left: FCFS's number of admitted EVs vs. demand level (identical result for all 50 scenarios). Bottom Right: Our online heuristic's number of admitted EVs vs. demand level (50 scenarios simulated).

with 4 cables each, yielding 60 parking spots total.

We compare the total user utility yielded from the FCFS strategy to our online heuristic with demand increasing from 1 to 120 arrivals each day. As shown in Fig. 2.6, the total user utility increases steadily as the number of EVs entering the system increases. However, when the demand for the EVSEs is high, our online mechanism is able to filter out low value arrivals to admit higher value arrivals instead, and yield higher total utility. It is worth noting that the FCFS strategy yields similar total utility if the demand for charging is low; this is because FCFS admits all arrivals as long as there are open parking spots.

Additionally, in Fig. 2.7, we show the fraction of behind-the-meter solar used by our mechanism and FCFS for a day when there are 100 arriving EVs. From the plot, it is evident that our online mechanism is able to utilize significantly more solar energy, thus eliminating the need to send large amounts of excess energy back to the distribution

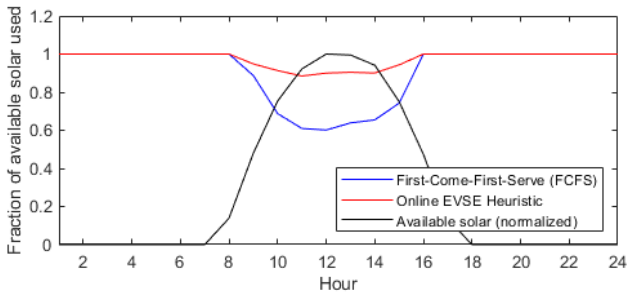

Figure 2.7: Comparison of behind-the-meter solar energy usage for a day with 100 arrivals.

grid. Specifically, our mechanism is able to schedule charging to time slots when there is available solar, while FCFS is not able to schedule charging times. The results in this section show that our mechanism outperforms the commonplace FCFS strategy in congested facilities in addition to better utilizing behind-the-meter solar.

Multi-Facility Test Case

In this section, we present a multi-facility example located in downtown Los Angeles. Specifically, we look at 6 parking facilities with varying rooftop generation amounts. We assume standard crystalline silicon panels with 14% system loss. Facility 6 has a 75kW solar generation system, facility 5 has 60kW, facility 4 has 45kW, facility 3 has 30kW, facility 2 has 15kW, and facility 1 does not have any solar generation. We examined a 20 day period with 600 arrivals each day. Each arrival has valuation in $[\$1, \$10]$, energy request in [1, 20] kWh, and stay length in [1, 8] hours. Each of the 6 facilities has 8 SOMC EVSEs each equipped with 4 cables. Furthermore, each facility purchases electricity from the Los Angeles grid at \$0.127/kWh. We examine the performance of the system with transformer capacity limits of 75kVA.

Figure 2.8 shows the total user utility, social welfare, and electricity cost for 20 days. An observation worth noting is the total user utility from our updated pricing function is always larger than that of a solar agnostic pricing framework. This is due to setting

lower prices on the electricity generation resource when there is solar available as seen in Figure 2.2 (our previous work was agnostic to the free solar generation). Moreover, over the 20 day period, our updated pricing mechanism admits 387 arrivals on average while our previous work only admits 369 arrivals on average. As such, our updated mechanism is favorable for users of the system as prices are lower and more users are admitted. Additionally, our improved mechanism is able to utilize more solar, reducing reliance on the local grid.

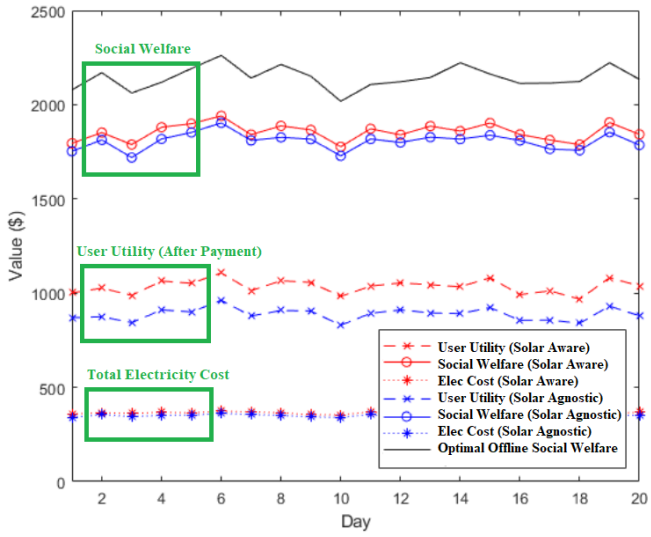

Figure 2.8: Multi-facility example. Red: solar aware pricing. Blue: solar agnostic pricing.

Importance of Accurate Departure Times

In this section, we discuss the effect of inaccurate departure time reporting. For example, consider the case when a user reports that she will exit the system by 4:00pm; however, she gets delayed and cannot remove her EV until 5:00pm or later. This affects the reservation system because there might be a reservation for the EVSE at that timeslot. To avoid these reservation collisions due to delayed departures, we analyze the performance of the system with extra hours added to each arrival's stay length as a

"buffer" to prevent double-allocations. We examined the same test case as Section 2.1.4 with buffer sizes of 1 and 2 hours added to each arrival's stay length. Adding 1 and 2 hour departure buffers yielded average social welfare losses of 16% and 29%, respectively.

Infrastructure Recommendation

In this section, we demonstrate the importance of infrastructure planning in order to maximize the smart charging capabilities of a parking facility. A facility with too few EVSEs will limit the users' utilities as well as the smart charging potential. Conversely, installing too many EVSEs results in idle chargers. As such, we perform a cost-benefit analysis to determine the number of cables at each SOMC EVSE as well as the number of EVSEs that should be installed at a facility in order to maximize social welfare over an extended period. Specifically, we are simulating the same system as described in Sections 2.1.4 and 2.1.4; however, we have increased the duration to 2 years (730 days). Furthermore, we are including initial and recurring costs relating to a parking structure equipped with SOMC EVSEs. These costs include EVSE unit costs, installation costs, electricity consumption, maintenance, and networking costs. In the following, we are looking to choose the constraint variables C_l and M_l (number of cables per EVSE and number of EVSEs, respectively) that maximize users' utilities minus the aforementioned investment and operational costs. We use I_C to denote the EVSE unit investment per cable and I_M to denote the installation cost per SOMC EVSE. Additionally, $I_{m,n}$ represents a recurring infrastructure maintenance and networking cost per EVSE. As such, the infrastructure investment cost can be written as:

$$\sum_{\mathcal{L}} \left(I_C C_l + I_M\right) M_l + T \sum_{\mathcal{L}} I_{m,n} M_l. \tag{2.20}$$

In (2.20), the first term is the initial investment cost for the SOMC EVSE hardware and

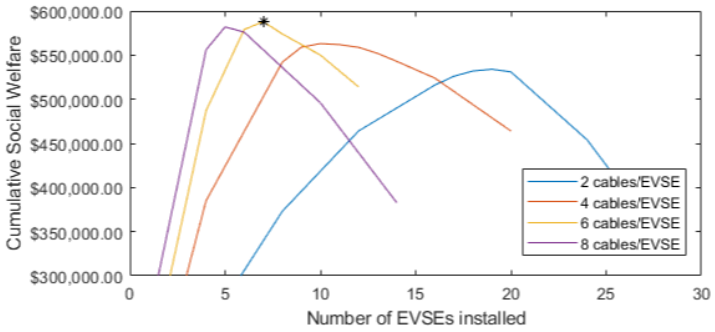

Figure 2.9: Cost-benefit analysis over 2 years.

installation and the second term represents the recurring maintenance and networking cost. In the following, we assumed each EVSE was a pedestal mounted unit with an installation cost of $\$3,308$ [41]. Additionally, each extra cable for each SOMC increased the cost of the EVSE by $\$3,343$ [41]. Last, the recurring networking and maintenance fee was assumed to be \$75 per month [41]. Figure 2.9 shows the comparison of total social welfare generated across the entire time period for various levels of infrastructure investments. For this test case, the best result occurred when each location had 7 SOMC EVSEs each equipped with 6 cables. With this level of infrastructure, the system did not yield positive social welfare until the second year of operation. As seen in Fig. 2.9, it is clear that sizing a facility for the given use case is critical. Smart charging strategies require a sufficient number of EVSEs to yield maximal benefits; however, welfare decreases if extra EVSEs are purchased and underutilized.

2.1.5 Conclusion

In this work, we presented an online pricing mechanism as a solution to the EVSE reservation problem. The online mechanism functions as both an admission controller and a distributor of the facilities' limited charging resources. The work presented in this section complements existing literature in the area and the important characteristics

are as follows. First, the mechanism readily accommodates multiple locations, multiple limited resources, operational costs, and variable arrival patterns. The mechanism does not rely on fractional allocations or rounding methods to produce integer allocations in a computationally feasible manner and it never revokes previously made reservations. Moreover, our online mechanism readily handles the inherent stochasticity of the EVSE reservation problem including unknown sojourn times, unknown energy requests, and unknown user valuation functions. The online mechanism can handle adversarially chosen arrival sequences and still generate social welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal. We discussed a competitive ratio as a performance guarantee for the online mechanism compared to the oracle offline solution and provided numerical results showing the efficacy of the mechanism.

2.2 Online Scheduling for Charging and Routing in Electric Vehicle Fleets

2.2.1 Introduction

Three developing technologies in the transportation sector have the potential to revolutionize the paradigm of *personal urban mobility*: autonomous vehicles (self-driving or driverless vehicles), mobility-on-demand (car-sharing or ride-sharing), and plug-in electric vehicles [42, 43]. These technologies have independently garnered much research and experimentation; however, literature addressing the potential synergies is still emerging [44]. Consequently, we consider the welfare maximization problem for a fleet dispatcher who operates a large number of Autonomous-Mobility-on-Demand Electric Vehicles (AMoD EVs). Because of the real-time requirements of AMoD systems, we propose a novel online solution for optimizing the charging and rebalancing processes of a fleet of AMoD EVs.

Regarding AMoD fleets, much work has been done focusing on matching riders with vehicles, routing vehicles to destinations, and rebalancing the vehicles throughout a set of pick-up/drop-off locations [45, 46, 47, 48, 49, 7]. There is also work in the area of coordinated charging for fleets of AMoD EVs. Paper [50] gives an overview of managing AMoD fleets and energy services in future smart cities. The authors of [51] utilize a model predictive approach to optimize charge scheduling and routing in an AMoD system. Paper [52] presents a study of the operations of a AMoD fleet including the implications of vehicle and charging infrastructure decisions. Furthermore, [53] studies the implications of pricing schemes on an AMoD fleet. Work has also been done in congestion aware [54] and model predictive routing methods [55] for AMoD systems. Additionally, [56, 57] study interactions between AMoD systems with the power grid and public transit.

Regarding charging strategies for large populations of EVs, papers [13, 14, 15] provide

in-depth reviews and studies of smart charging technologies. An important but less studied issue is that the benefits of smart charging can be severely limited if usage of the shared Electric Vehicle Supply Equipment (EVSEs) is uncontrolled [24]. Moreover, without access control and allocation strategies within charging facilities, EVSEs can become congested while other EVSEs are left empty. This limits the smart charging benefits as congested EVSEs are forced to charge one EV after another to satisfy demand. To address this issue, papers including [26, 28, 31, 32] have studied smart charging, admission control, and resource allocation for facilities equipped with EVSEs.

In this section, we aim to complement both the recent work in smart charging and AMoD fleet routing with the objective of optimizing AMoD fleet charging and rebalancing processes in an online fashion. Specifically, we study an online heuristic that schedules fleet charging, allocates limited fleet resources, and rebalances the vehicles while avoiding congestion at charging facilities and pick-up locations. Moreover, our methodology does not rely on statistics of the daily sessions (unlike model predictive approaches). The work presented here complements existing literature on smart charging and fleet management and the main contributions are as follows:

- The online heuristic makes decisions for multiple charging facilities, multiple pick-up/drop-off regions, constrained fleet charging resources, operational costs, renewable generation integration, and rebalancing.
- The online heuristic does not rely on fractional allocations or rounding methods to produce integer assignments in a computationally feasible manner.
- The online heuristic readily handles the inherent stochasticity of the fleet scheduling problem without requiring statistics on the future inputs of the system.
- The online heuristic accounts for worst case (i.e., adversarially chosen between-ride

sequences) and always yields welfare within a factor of $\frac{1}{\alpha}$ of the optimal offline.

Organization: Section 2.2.2 describes the AMoD EV fleet charge scheduling problem and system model. Section 2.1.3 presents the online heuristic that updates the dual variables when solving the online problem and discusses the online heuristic's properties and performance guarantees. Section 2.1.4 discusses numerical results showing the performance of our heuristic.

2.2.2 Preliminaries

Fleet Objective

In this section, we describe the charge scheduling problem for a fleet of AMoD EVs that are servicing customers within a city. The objective of the fleet dispatcher is to maximize profit by optimizing the between-ride schedules of the EVs to exploit cheaper time-of-use electricity rates and behind-the-meter solar generation in addition to efficiently distributing the vehicles throughout the area.

We consider an area of operation consisting of a set of regions $\mathcal{D} = \{1, \ldots, D\}$ where each region $d \in \mathcal{D}$ can be viewed as a destination for a vehicle to pick up its next customer. Each region d has a maximum capacity $\Omega_d(t)$ that limits the number of AMoD vehicles in the area (e.g., due to municipal restrictions, congestion mitigation, or ride demand forecasts). Within this area of operation, we consider J consecutive between-ride sessions occurring within a time span $t = 1, \ldots, T$. Each between-ride session $j \in \mathcal{J}$ begins at time t_j^- when a vehicle in the fleet completes a previous ride. Information for the jth between-ride session could be revealed earlier when the EV picks up the passenger rather than at drop-off; however, due to unknown traffic conditions and travel times, the session information is not available until t_j^- (drop-off). At this time, the fleet dispatcher must decide what the vehicle should do next. Namely, a schedule must be created for the

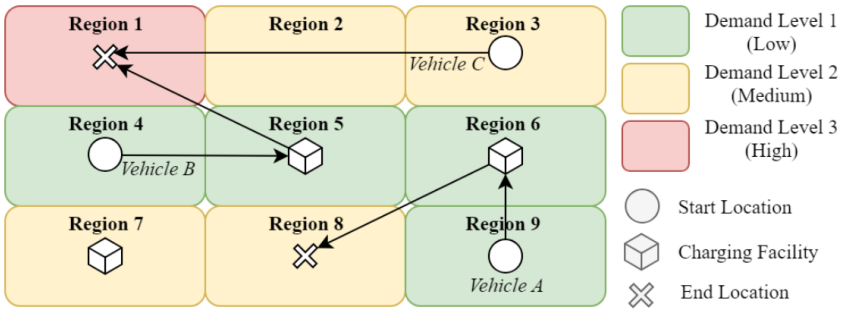

Figure 2.10: Example between-ride schedules for 3 vehicles.

between-ride session that includes a charging facility, a desired charge amount, and the next customer pickup destination (if the EV has sufficient battery level, the schedule may skip charging altogether). A simple example for three vehicles' between-ride schedules is portrayed in Fig. 2.10.

Charging Model

There is a set of regions $\mathcal{F} \subseteq \mathcal{D}$ that have charging facilities with Electric-Vehicle-Supply-Equipment (EVSEs) where the EVs can replenish their batteries. Each charging facility $f \in \mathcal{F}$ is equipped with M_f Single-Output-Multiple-Cable (SOMC) EVSEs each with C_f output cables meaning each facility can fit up to $M_f C_f$ EVs simultaneously (but only charge M_f simultaneously) [33, 10, 9, 4]. Additionally, due to power delivery limitations of the EVSE hardware, each EVSE can deliver up to E_f units of energy in one time slot. Each facility can procure energy from two sources: 1) an on-site solar generation system and 2) the local power distribution grid. The available solar energy at facility f at time t is denoted as $\delta_f(t) \in [0, \Delta_f]$ where Δ_f is the maximum power rating of the on-site system at facility f. Energy can also be purchased from the local distribution grid at price $\pi_f(t)$ per unit. Due to the local transformer's operational limits, facility f is constrained to purchase less than $\mu_f(t)$ units of energy from the grid in each time slot.

Operational Cost Model

Since we are solely focused on optimizing a given number of between-ride sessions and are not explicitly modeling the rest of the AMoD system operation problem, we need to penalize vehicles for remaining out-of-service (i.e., charging or traveling without customers) for extended periods of time. This is also essential to compare the quality of our online solution to that of the clairvoyant offline solution. As such, we assume the fleet dispatcher incurs a virtual cost, $\phi(t)$, per out-of-service vehicle at each time t. Moreover, $\phi(t)$ is permitted to be time-varying to give the fleet dispatcher control over the quantity of active fleet vehicles throughout a given day. A large $\phi(t)$ could be used to ensure that less vehicles are out-of-service during peak demand periods.

Offline and Online Problem

In this work, we first formulate the charge scheduling problem as an offline optimization and then relate the offline problem to the online problem. In the offline case, we assume the fleet dispatcher is clairvoyant and knows the entire sequence of J events over the time span $t = 1, \ldots, T$. As such, the offline fleet dispatcher can create the optimal schedules for the between-ride sessions and can achieve maximal profit. However, the reality is that the dispatcher does not know the customers' desired destinations, pick-up to drop-off travel times, traffic conditions, or electricity grid conditions in advance. Instead, the between-ride sessions are revealed one-by-one throughout the time span meaning that an online solution method is required for real world implementation. Moreover, the charge scheduling problem presents challenges not easily overcome in many online solutions; namely, the lack of accurate statistics on the between-ride sessions as there are many exogenous factors that directly affect when and where between-ride sessions begin. These factors include ride demand, road congestion, weather, popular events, and

construction delays, to name a few. As such, we present an online solution that can account for adversarially chosen sequences of between-ride sessions and still yield profit that is within a constant factor of the clairvoyant offline solution.

Between-Ride Schedules

In this section, we describe the parameters associated with each between-ride schedule. Each between-ride session $j \in \mathcal{J}$ begins at time t_j^- when an AMoD EV drops off a passenger. At this time, a set of feasible between-ride schedules is generated based on the vehicle's current battery level and location. We denote the set of feasible between-ride schedules for session j as $\mathcal{S}_j$. Each schedule $s \in \mathcal{S}_j$ consists of the following components: 1) t_j^-: time when all between-ride schedules start for session j; 2) d_j^-: start location of all schedules; 3) t_{js}^+: end time for schedule s; 4) $d_{js}^+(t)$: binary (0,1) indicator function for when the EV reaches its end destination d_{js}^+ in schedule s of session j; 5) $o_{js}(t)$: out-of-service indicator that is set to 1 for $t \in [t_j^-, t_{js}^+]$ and 0 otherwise; 6) v_{js}: value of schedule s to the fleet dispatcher given by

$$v_{js} = V_d(SoC_{js}^+) + v_{d_{js}}. \tag{2.21}$$

Here, SoC_{js}^+ represents the *State of Charge* of the EV when it reaches its next pickup location. The function $V_d(\cdot)$ calculates the fleet dispatcher's valuation of the final energy level of the EV in schedule s at location d (maintaining sufficient energy levels in the EVs is critical to provide uninterrupted rides to customers). We note that the fleet dispatcher is free to choose $V_d(\cdot)$ for their desired operational objectives. The variable $v_{d_{js}}$ represents the profit that the fleet dispatcher receives for picking up a customer at destination d_{js}. The destinations present different values to the fleet dispatcher due to exogenous factors such as ride demand, location, weather, etc; 7) $c_{js}^{mf}(t)$: binary (0,1) cable reservation

variable that is set to 1 if the EV in session j will use a charging cable from EVSE m at facility f during time step t in schedule s; 8) $e_{js}^{mf}(t)$: energy delivered to EV at time t in session j at EVSE m at facility f in schedule s. Through $e_{js}^{mf}(t)$, the fleet dispatcher is able to customize when the EV will receive charge, when it will be idle, and the rate of charge (from a discrete set). This effectively allows the fleet dispatcher to *smart charge* the EVs (e.g., exploit cheaper time-of-use electricity rates or behind-the-meter solar). The feasible schedules for between-ride session j can be written as:

$$\{t_j^-, \{o_{js}(t)\}, \{c_{js}^{mf}(t)\}, \{e_{js}^{mf}(t)\}, d_j^-, \{d_{js}^+(t)\}, t_{js}^+, \{v_{js}\}\}. \tag{2.22}$$

Furthermore, the fleet dispatcher sets the variable x_{js} to 1 if schedule s is selected and 0 otherwise. If no schedule is desirable, the dispatcher routes the vehicle to the central depot where the vehicle will wait to be assigned later on.

Cost Model and Offline Scheduling Problem

In order to facilitate charge scheduling and vehicle routing, the fleet dispatcher maintains a total count for each shared resource across all assigned schedules. The variables $y_c^{mf}(t)$ and $y_e^{mf}(t)$ correspond to the total allocated cables and energy, respectively, at facility f at EVSE m at time t. Additionally, each facility has to procure the energy needed by all the EVSEs within. The total energy procurement at facility f at time t is denoted as $y_g^f(t)$. The resource demands at charging facilities are calculated in equations

(2.23)-(2.25):

$$y_c^{mf}(t) = \sum_{\mathcal{J},\mathcal{S}_j} c_{js}^{mf}(t) x_{js}, \tag{2.23}$$

$$y_e^{mf}(t) = \sum_{\mathcal{J},\mathcal{S}_j} e_{js}^{mf}(t) x_{js}, \tag{2.24}$$

$$y_g^{f}(t) = \sum_{\mathcal{J},\mathcal{S}_j,\mathcal{M}_f} e_{js}^{mf}(t) x_{js}. \tag{2.25}$$

Similarly, the fleet dispatcher counts the vehicles in the between-ride state and the number of vehicles committed to destinations in the variables $y_o(t)$ and $y_d(t)$, respectively:

$$y_o(t) = \sum_{\mathcal{J},\mathcal{S}_j} o_{js}(t) x_{js}, \tag{2.26}$$

$$y_d(t) = \sum_{\mathcal{J},\mathcal{S}_j} d_{js}^{+}(t) x_{js}. \tag{2.27}$$

Due to the number of vehicles allocated to each resource, the fleet dispatcher incurs various costs to serve the between-ride schedules. The energy procurement, $y_g^f(t)$, determines the generation cost of facility f:

$$G_f(y_g^f(t)) = \begin{cases} 0 & y_g^f(t) \in [0, \delta_f(t)] \\ \pi_f(t)(y_g^f(t) - \delta_f(t)) & y_g^f(t) \in \big(\delta_f(t), \delta_f(t) + \mu_f(t)\big] \\ +\infty & y_g^f(t) > \delta_f(t) + \mu_f(t). \end{cases} \tag{2.28}$$

Namely, electricity is free while solar is available, else the facility purchases from the grid at a rate of $\pi_f(t)$ per unit.

Additionally, vehicles that are charging and traveling to their next pickup destination are not able to serve customers. As described in Section 2.2.2, the fleet dispatcher incurs

a penalty proportional to the number of out-of-service vehicles:

$$O(y_o(t)) = \begin{cases} \phi(t) y_o(t), & y_o(t) \leq I(t) \\ +\infty & y_o(t) > I(t), \end{cases} \tag{2.29}$$

where $I(t)$ is the maximum number of out-of-service vehicles that the fleet dispatcher allows at time t. If the AMoD fleet dispatcher has full knowledge of all between-ride sessions $j \in \mathcal{J}$ over the entire time span $t = 1, \ldots, T$, the primal offline optimization is as follows:

$$\max_{x} \sum_{\mathcal{J}, \mathcal{S}_j} v_{js} x_{js} - \sum_{\mathcal{T}, \mathcal{F}} G_f(y_g^f(t)) - \sum_{\mathcal{T}} O(y_o(t)) \tag{2.30a}$$

subject to:

$$\sum_{\mathcal{S}_j} x_{js} \leq 1, \quad \forall\, j \tag{2.30b}$$

$$x_{js} \in \{0, 1\}, \quad \forall\, j, s \tag{2.30c}$$

$$y_c^{mf}(t) \leq C_f, \quad \forall\, f, m, t \tag{2.30d}$$

$$y_e^{mf}(t) \leq E_f, \quad \forall\, f, m, t \tag{2.30e}$$

$$y_d(t) \leq \Omega_d(t), \quad \forall\, d, t \tag{2.30f}$$

and (2.23), (2.24), (2.25), (2.26), (2.27).

In (3.5a), the objective maximizes fleet dispatcher's utility by distributing the AMoD EVs throughout all regions $d \in \mathcal{D}$ while minimizing the operational costs due to charging facilities and out-of-service vehicles. Specifically, the last term of (3.5a) limits the duration of the between-ride sessions to increase earnings and decrease the need to use extra vehicles. Constraint (2.30b) ensures only one schedule is chosen per between-ride session, (2.30c) is an integral constraint on the decision variable, (2.30d) ensures the cable limit

is not exceeded at each EVSE, (2.30e) ensures the EVSE energy limits are not exceeded, and (2.30f) enforces the vehicle limit in each region. By temporarily relaxing the integral constraint (2.30c), the problem can be examined in the dual domain (however, we note that our competitive ratio results are with respect to integer allocations). Specifically, we make use of Fenchel duality with dual variables u_j, $p_c^{mf}(t)$, $p_e^{mf}(t)$, $p_g^f(t)$, $p_d(t)$, and $p_o(t)$ [34]. In the following, the Fenchel conjugate of a function $f(y(t))$ is given as:

$$f^*(p(t)) = \sup_{y(t)\geq 0} \big\{p(t)y(t) - f(y(t))\big\}. \tag{2.31}$$

As such, the offline Fenchel dual of (3.5a)-(2.30f) is as follows:

$$\begin{aligned} \min_{u,p} & \sum_{\mathcal{J}} u_j + \sum_{\mathcal{T},\mathcal{D}} R_d^*(p_d(t)) \\ & + \sum_{\mathcal{T},\mathcal{F}} G_f^*(p_g^f(t)) + \sum_{\mathcal{T}} O^*(p_o(t)) \\ & + \sum_{\mathcal{T},\mathcal{F},\mathcal{M}_f} \Big(K_c^{mf*}(p_c^{mf}(t)) + K_e^{mf*}(p_e^{mf}(t))\Big) \end{aligned} \tag{2.32a}$$

subject to:

$$\begin{aligned} u_j \geq\ v_{js} & - d_{js}^+(t_{js}^+)p_d(t_{js}^+) - \sum_{\mathcal{T}} \Big(c_{js}^{mf}(t)p_c^{mf}(t) \\ & + e_{js}^{mf}(t)\big[p_e^{mf}(t) + p_g^f(t)\big] + o_{js}(t)p_o(t)\Big) \\ & \forall\ j,s,f,m,d \end{aligned} \tag{2.32b}$$

$$u_j \geq\ 0, \quad \forall\ j \tag{2.32c}$$

$$\begin{aligned} p_c^{mf}(t),\ p_e^{mf}(t),\ p_g^f(t),\ p_d(t),\ p_o(t) \geq\ 0, \\ \forall\ f,m,d,t, \end{aligned} \tag{2.32d}$$

where $R_d^*(p_d(t))$, $G_f^*(p_g^f(t))$, $O^*(p_o(t))$, $K_c^{mf*}(p_c^{mf}(t))$, and $K_e^{mf*}(p_e^{mf}(t))$ are the Fenchel conjugates for the regional vehicle limit, the facility generation cost, the out-of-service cost, the EVSE cable constraint, and the EVSE energy constraint, respectively. The Fenchel conjugates for the cable and energy constraints, respectively, are as follows:

$$K_c^{mf*}(p_c^{mf}(t)) = p_c^{mf}(t)C_f, \qquad p_c^{mf}(t) \geq 0, \tag{2.33}$$

$$K_e^{mf*}(p_e^{mf}(t)) = p_e^{mf}(t)E_f, \qquad p_e^{mf}(t) \geq 0. \tag{2.34}$$

The Fenchel conjugate for the energy procurement cost function at each facility can be written as:

$$G_f^*(p_g^f(t)) = \begin{cases} \delta_f(t)p_g^f(t), & p_g^f(t) < \pi_f(t) \\ (\delta_f(t) + \mu_f(t))p_g^f(t) - \mu_f(t)\pi_f(t) & p_g^f(t) \geq \pi_f(t). \end{cases} \tag{2.35}$$

The Fenchel conjugate for the regional vehicle limit is:

$$R_d^*(p_d(t)) = p_d(t)\Omega_d(t), \qquad p_d(t) \geq 0. \tag{2.36}$$

Lastly, the Fenchel conjugate for the penalty for the out-of-service vehicles is as follows:

$$O^*(p_o(t)) = \begin{cases} 0, & p_o(t) < \phi(t) \\ (p_o(t) - \phi(t))I(t), & p_o(t) \geq \phi(t). \end{cases} \tag{2.37}$$

Scheduling Decisions

In the offline case, let us examine the Fenchel dual (2.32a)-(2.32d) when choosing charging schedules for the between-ride sessions. Specifically, examining the KKT con-

ditions for constraint (2.32b) reveals the optimal scheduling decisions. If a between-ride session yields $u_j \leq 0$, then the AMoD EV is not needed to serve customers at that time or electricity prices are too high; therefore, u_j is set to 0 and the vehicle is routed to the central depot. Alternatively, when the fleet dispatcher wants vehicles to charge and serve customers, u_j will be positive. If $u_j > 0$, the optimal schedule for session j can be found by finding the schedule $s \in \mathcal{S}_j$ that results in the maximal u_j:

$$u_j = \max_{s \in \mathcal{S}_j} \Big\{ v_{js} - p_d(t_{js}^+) d_{js}^+(t_{js}^+) - \sum_{t \in [t_j^-, t_{js}^+]} \Big(o_{js}(t) p_o(t) + c_{js}^{mf}(t) p_c^{mf}(t) + e_{js}^{mf}(t) \big[p_e^{mf}(t) + p_g^f(t) \big] \Big) \Big\}. \tag{2.38}$$

We note that u_j corresponds to the utility gained from between-ride session j for the fleet dispatcher. Furthermore, the optimization problems in (3.5a)-(2.30f) and (2.32a)-(2.32d) require full knowledge of the between-ride sessions beforehand. As discussed in Section 2.2.2, the fleet dispatcher must utilize an online solution that can schedule vehicles without any knowledge of the future (i.e., without knowledge of the optimal dual variables). In the following, we discuss such an online heuristic for the between-ride charge scheduling problem.

2.2.3 Online Scheduling Heuristic

Online Scheduling

Because the between-ride sessions are revealed throughout the time span, it is apparent that an online solution method is required. We consider a heuristic that updates the dual variables in an online fashion as between-ride sessions are revealed and then solves equation (2.38) for each session. The online scheduling heuristic updates the dual variables for each resource based only on the amount of resource that has been allocated

up to the current time (i.e., only utilizing the resource allocation counts). The online heuristic serves two main purposes: 1) it ensures that each between-ride schedule yields more value to the fleet dispatcher than the operational cost pertaining to the schedule, and 2) if the demand for rides is low enough or electricity prices are high enough such that no schedule nets positive utility, vehicles are sent back to the central depot. This eliminates further costs from vehicles circulating without serving riders. The underlying framework for the heuristic we use is akin to that of [39], where the authors present an auction for allocating computing resource bundles at data centers for the purpose of cloud computing and virtual machine provisioning.

In our online scheduling heuristic, we make use of specialized functions proposed in [39] that approximate the optimal dual variables throughout the time span. These dual variable update functions increase slowly at first, but then increase rapidly as the number of allocated resources approach the capacity limits. Furthermore, when the amount of allocated resource is at the capacity limit, the update functions output dual variables high enough such that no schedule will yield positive utility in (2.38), thus enforcing the hard capacity limits. The updating function for the dual variable associated with the SOMC EVSE cables at charging facilities is written as follows:

$$p_c^{mf}(y_c^{mf}(t)) = \Big(\frac{L_c}{2\Psi}\Big)\Big(\frac{2\Psi U_c}{L_c}\Big)^{\frac{y_c^{mf}(t)}{C_f}}, \tag{2.39}$$

where Ψ is the total number of shared resources within the fleet system:

$$\Psi = 2\sum_{\mathcal{F}} M_f + D + F + 1. \tag{2.40}$$

Furthermore, L_c and U_c correspond to the minimum and maximum value per cable per time unit, respectively. The online scheduling heuristic requires knowledge of L_c and U_c

beforehand to set the initial values and to ensure capacity limits are not breached:

$$L_c = \min_{\mathcal{J},\mathcal{S}_j,\mathcal{F},\mathcal{M}_f} \frac{v_{js}}{\Psi \sum_{t\in[t_j^-,t_{js}^+]} c_{js}^{mf}(t)}, \tag{2.41a}$$

$$U_c = \max_{\mathcal{J},\mathcal{S}_j,\mathcal{F},\mathcal{M}_f,\mathcal{T}} \frac{v_{js}}{c_{js}^{mf}(t)}, \quad c_{js}^{mf}(t) \neq 0. \tag{2.41b}$$

The EVSE charging power, facility generation, out-of-service cost, and destination vehicle limit require similar lower and upper bounds on valuations: L_e, U_e, L_g, U_g, L_o, U_o, L_d, and U_d, respectively. These are calculated as in equations (2.41a) and (2.41b) with the corresponding variables to replace $c_{js}^{mf}(t)$.

The dual variable update function for the dual variable associated with the SOMC EVSE energy limitations at charging facilities is as follows:

$$p_e^{mf}(y_e^{mf}(t)) = \Big(\frac{L_e}{2\Psi}\Big)\Big(\frac{2\Psi U_e}{L_e}\Big)^{\frac{y_e^{mf}(t)}{E_f}}. \tag{2.42}$$

The update function of the dual variable for the piecewise linear generation cost at each charging facility is more complex than (2.39) and (2.42). It has to account for the free solar generation as well as the linear price to procure energy from the local distribution grid. As such, we propose the following dual variable update function:

$$p_g^f(y_g^f(t)) = \begin{cases} \Big(\frac{L_g}{2\Psi}\Big)\Big(\frac{2\Psi\pi_f(t)}{L_g}\Big)^{\frac{y_g^f(t)}{\delta_f(t)}}, & y_g^f(t) < \delta_f(t), \\ \Big(\frac{L_g-\pi_f(t)}{2\Psi}\Big)\Big(\frac{2\Psi(U_g-\pi_f(t))}{L_g-\pi_f(t)}\Big)^{\frac{y_g^f(t)}{\delta_f(t)+\mu_f(t)}} + \pi_f(t), & y_g^f(t) \geq \delta_f(t). \end{cases} \tag{2.43}$$

The heuristic dual variable update function for the vehicle limits at region d can be written as follows:

$$p_d(y_d(t)) = \Big(\frac{L_d}{2\Psi}\Big)\Big(\frac{2\Psi U_d}{L_d}\Big)^{\frac{y_d(t)}{\Omega_d(t)}}. \tag{2.44}$$

Last, the penalty from vehicles in the out-of-service state also requires a heuristic update function for the dual variable $p_o(y_o(t))$ which can be written as:

$$p_o(y_o(t)) = \Big(\frac{L_o - \phi(t)}{2\Psi}\Big)\Big(\frac{2\Psi(U_o - \phi(t))}{L_o - \phi(t)}\Big)^{\frac{y_o(t)}{I(t)}} + \phi(t). \tag{2.45}$$

With the 5 aforementioned dual variable update functions (2.39), (2.42), (2.43), (2.44), and (2.45), we now have surrogate functions to use in place of the optimal dual variables in order to solve equation (2.38) in an online fashion (i.e., at the inception of each between-ride session).

Procedure and Performance Guarantees

The step-by-step procedure for scheduling between-ride sessions for a fleet of AMoD EVs is outlined in Algorithm ONLINEAMODSCHEDULING. Namely, at the inception of between-ride session j, the fleet dispatcher generates feasible schedules $\mathcal{S}_j$ and then the best schedule, s^*, is chosen in line 8 which makes use of the heuristically updated dual variables. After every between-ride schedule selection, the fleet dispatcher updates the dual variables accounting for the total amounts of allocated resources. We can compare the total welfare generated from our online solution to that of the clairvoyant offline solution in the form of a competitive ratio. An online heuristic is described as α-competitive when the ratio of welfare from the clairvoyant offline solution to the welfare from the online heuristic is bounded by $\alpha \geq 1$. For the between-ride charge scheduling problem, we extend a previous competitive ratio result from [39]. In this work, we assume that each between-ride session utilizes only a small amount of the available resources, thus ensuring that the allocation of one schedule does not adversely affect too many future sessions.

Algorithm 2 ONLINEAMODSCHEDULING

Input: $I, J, \mathcal{F}, \mathcal{M}_f, C_f, E_f, \mu_f, \delta_f, \pi_f,$
$\mathcal{D}, \Omega_d, \phi, \Psi, \{L, U\}_{c,e,g,d,o}$

Output: x, p

Define $G_f(\cdot)$ and $O(\cdot)$ according to (2.28) and (2.29).
Define and initialize the dual update functions (2.39), (2.42), (2.43), (2.44), and (2.45).
Initialize $x_{js} = 0$, $y(t) = 0$, $u_j = 0$.
At the inception of between-ride session j:
Generate feasible charging/pickup schedules $\mathcal{S}_j$.
Update dual variable u_j according to (2.38).
if $u_j > 0$ **then**
    $s^\star = \arg\max_{\mathcal{S}_j} \{v_{js} - p_d(t^+_{js}) d^+_{js}(t^+_{js}) - \sum_{t\in[t^-_j, t^+_{js}]} [o_{js}(t) p_o(t) + c^{mf}_{js}(t) p^{mf}_c(t) + e^{mf}_{js}(t)[p^{mf}_e(t) + p^f_g(t)]]\}$
    $x_{js^\star} = 1$ and $x_{js} = 0$ for all $s \neq s^\star$
    Update demand $y(t)$ for cables, energy, generation, destination, and out-of-service according to (2.23)-(2.27).
    Update dual variables $p(t)$ according to (2.39), (2.42), (2.43), (2.44), and (2.45).
else
    Send the AMoD EV to central depot to re-enter system later and set $x_{js} = 0, \quad \forall\, s \in \mathcal{S}_j$.
end if

Theorem 2.2.1 *The online heuristic* ONLINEAMODSCHEDULING *in Algorithm 2 is α-competitive in welfare across all fleet resources for the fleet dispatcher over J between-ride sessions where $\alpha = \max\{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5\}$.*

Proof: From the Lemmas (in online Appendix [8]), we have welfare guarantees $\alpha_1, \ldots, \alpha_5$ for each of the shared resources. To find the $\alpha \geq 1$ for the entire system, we take the maximum over $\alpha_1, \ldots, \alpha_5$ to yield the bound that accounts for all resources. ■

2.2.4 Numerical Results

In this section, we discuss numerical results showing the performance of our online heuristic. In the following simulation, electricity prices and solar generation data were sourced from actual California ISO data in the Bay Area from 2018 [58],[59]. We simulated for a fleet operating in San Jose, CA with $D = 46$ regions and $F = 8$ charging

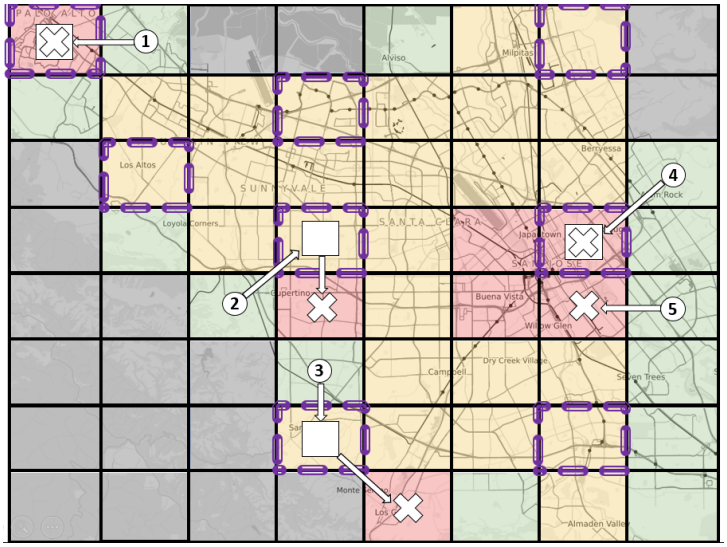

Figure 2.11: San Jose, CA ride demand split into $D = 46$ regions. Ride demand high to low: Red, Yellow, Green. Purple Outline: Charging facility available. Gray Regions: No demand. Also Shown: Most popular between-ride routes for 5 different starting points. Circle: Between-ride session start location. Square: Charging session at facility. Cross: Between-ride end location.

facilities. Each charging facility is identical with $M_f = 10$ and $C_f = 4$. Each of the 8 facilities has on-site solar with a maximum generation of 256 kWh. Likewise, each facility can purchase energy from the grid up to 256 kWh per time step. We set the penalty for out-of-service vehicles equal to 2× the highest grid electricity price to penalize lengthy between-ride durations. Valuations for each of the regions were either \$15, \$10, or \$5 (red, yellow, green, in Fig. 2.11, respectively) based on daily ride demand [60]. The regions have AMoD vehicle limits set to 40 (separate from facility capacity if there is a facility in the region). Each vehicle entered the system with either 25%, 50%, or 75% battery level (50 kWh batteries) and were allowed make charge requests in increments of 12.5 kWh. The EVSEs at the charging facilities were limited to deliver either 0 kWh per time slot or 5 kWh per time slot. Furthermore, $V(100\%) = \$10, V(75\%) = \$7.5, V(50\%) = \$5$, and $V(25\%) = \$2.5$. We also included a linear penalty (\$2 per region traveled) to the schedule valuations to devalue long between-ride routes.

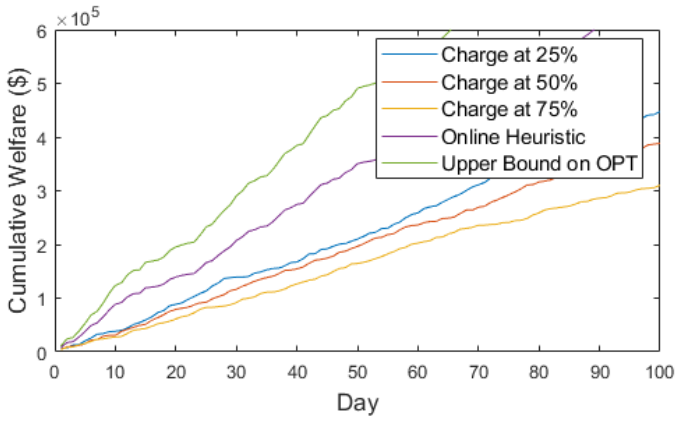

Figure 2.12: Fleet dispatcher welfare across 100 days.

In Figure 2.11 we show the most popular between-ride routes for vehicles starting at 5 of the regions (randomly chosen). In Figure 2.12, we compare the welfare generated by our online heuristic to 3 conservative online methods and an upper bound on the optimal solution. Specifically, we compare our heuristic to 3 threshold policies where each EV automatically charges if it is below 75%, 50%, or 25% at the nearest facility and then routes itself to the closest high value destination that is not at the AMoD vehicle limit. As seen in Figure 2.12, our heuristic is able to consistently outperform these threshold strategies. This is primarily due to shifting charging to time slots when there is available solar or cheaper time-of-use electricity rates (the threshold strategies charge once the vehicle is plugged in and do not stop until fully charged). Because the state-space of any system such as this grows exponentially to intractable sizes, we instead present an upper bound on the optimal offline solution. Accordingly, we calculated the upper bound by relaxing the capacity limits at each facility and destination to allow for all sessions to select their utility maximizing schedule without being constrained by capacity limits.

2.2.5 Conclusion

In this section, we studied an online heuristic that provides an approximate solution

to the AMoD fleet charge scheduling problem. The online heuristic makes fleet decisions for multiple regions, constrained resources, operational costs, renewable generation integration, and rebalancing. The online heuristic does not rely on fractional allocations or rounding methods to produce integer assignments in a computationally feasible manner. Additionally, online heuristic readily handles the inherent stochasticity of the fleet charge scheduling problem including unknown start locations, unknown energy levels, and unknown ride length/destination statistics. Last, the online heuristic accounts for worst case scenarios (i.e., adversarially chosen between-ride sequences) and always yields welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal. We discussed a competitive ratio (with proofs in the online Appendix) and presented simulation results highlighting the performance of the heuristic.

Chapter 3

Online Scheduling for a Community Energy Storage System

3.1 Introduction

Due to the increasing integration of distributed renewable generation in modern power grids, there is growing interest towards implementing distributed *energy storage* (ES) systems in close proximity to energy consumers [61, 62, 63]. Implementing ES near consumers enables various positive outcomes stemming from increased opportunities in demand-side management, e.g., CO_2 emission reduction from peak load shaving, increasing the amount of locally-consumed energy from nearby renewable distributed generation, or electricity cost reduction from shifting electricity purchases to off-peak hours [64]. Additionally, the concept of *energy communities* is on the rise. Specifically, these are groups of residential and commercial consumers/prosumers that cooperate and take advantage of shared resources (e.g., energy storage systems [65, 11]) and make use of

This work was supported in part by NSF under Grant 1847096 and in part by the UCSB Institute for Energy Efficiency (IEE).

each others' excess renewable generation. Recently, energy communities have garnered much research interest in various areas, including, but not limited to, peer-to-peer energy trading [66],[67], blockchain based energy transactions [68],[69], real-time optimization for energy management [70], [10], and game-theoretic market designs [71].

To maximize the utility gained from distributed ES implementations and energy communities, the concept of *community energy storage* (CES) is increasing in popularity [72, 73]. Specifically, a CES is a modular ES implemented within an energy consumption area (e.g., neighborhood, shopping center, etc.) in combination with renewable distributed generation in the area. CES systems are larger than single-consumer ES systems and have larger technical and economic benefits than single-consumer systems due to diversity in load profiles, removing the need for personal investments by individual consumers, as well as economies of scale [72, 74, 75]. Recently, there has been much work focusing on optimizing the design [76] and the basic operation [77, 78] of CES. A comprehensive review of different aspects of modern CES can be found in [79].

While it is evident that CES has great potential to positively impact energy consumers, the effectiveness of a CES system can be severely limited if it is not operated well. Namely, because there are multiple users who want to take advantage of a CES, there must be a smart management system in place to schedule the users' charging and discharging of the CES. If there is no smart management system in place, the CES might be underutilized or overutilized at various times. For example, all the users of the CES might choose to charge and discharge at similar times, (i.e., charging the CES with excess solar generation midday and discharging in the early evening) which limits the number of users who are able to make use of the CES and potentially leaves the CES underutilized at all other time periods. Additionally, any CES management system also has to deal with large amounts of uncertainty. Currently, one of the major technical challenges for future CES implementations is the requirement to handle uncertainty [80] in

the charging, discharging, and storage demands of the users. The users of a CES have inherently stochastic electricity demand and their desired usage of the CES is unknown and time-varying as users' net load patterns can vary significantly from day to day and weather can affect distributed renewable generation. With this in mind, it is clear that future CES implementations require advanced scheduling algorithms in order to operate effectively (i.e., maximize value gained by the system) under uncertain usage patterns.

Main Contributions

The work presented in this chapter considers the problem of a CES manager attempting to schedule the charging and discharging of a CES for a group of users. Our proposed solution allows for the users to request temporal charging and discharging profiles from the CES in real-time (as they learn about their needs) and the CES manager is able to immediately accept or deny a request and, if accepted, select the profile that maximizes the users utility. Additionally, due to the fact that our solution handles charging and discharging profiles instead of pure capacity requests, our heuristic is able to promote diverse charging and discharging patterns via dynamically updated prices to exploit charging/discharging cancellations and increase the CES's utilization. For example, a charge/discharge cancellation occurs when user A commits to charging the CES at a given time and user B commits to discharging the CES at the same time, thus effectively cancelling each other's power usage of the CES at that time and allowing other users access to charge/discharge at that time slot. Furthermore, we present a theoretical guarantee on the performance of our heuristic which operates in real-time without knowledge of future requests (i.e., our algorithm does not require a forecast of future requests). We are able to bound the worst case performance of our *online* solution in relation to the *offline* optimal solution (i.e., if the CES manager had known the entire sequence of CES requests beforehand) in the form of a competitive ratio. We note that this is a worst case

performance guarantee that holds for any *adversarially* chosen CES request sequence.

The main contributions of this work are as follows:

- *Temporal User Flexibility:* The proposed online heuristic allows users to submit requests for temporal charging and discharging profiles from the CES in real-time instead of committing to long-term capacity reservations far in advance. Furthermore, the CES scheduling heuristic will immediately accept or deny the request.

- *Charging and Discharging Cancellation:* As stated previously, the proposed scheduling heuristic deals with temporal charging and discharging profiles instead of capacity reservations. This allows for the heuristic's dynamically updated prices to promote diverse charging and discharging schedules of the users to take advantage of concurrent charging and discharging requests cancelling each other out, hence increasing efficiency.

- *Upholding CES Constraints:* The proposed online heuristic makes use of dynamically updated prices that are designed to ensure that the CES constraints (e.g., maximum charging power, maximum discharging power, maximum capacity) are met at all times.

- *Unknown Nature of Future Requests:* The proposed online heuristic readily handles the inherent uncertainty of the CES scheduling problem including unknown request times, unknown charging/capacity requests, and unknown valuations without the need of a future model. Specifically, we develop an online primal-dual optimization framework (an overview of primal-dual approaches for solving large-scale optimization problems can be found in [81]) that is able to provide a worst-case performance guarantee for any future request sequence (even adversarially selected sequences). The developed online optimization framework is akin to algorithmic posted pricing

mechanisms for online combinatorial auctions.

- *Theoretical Worst Case Performance Guarantee:* The online heuristic is robust to adversarially chosen request sequences and always yields social welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal (i.e., if the CES manager had known the entire sequence of CES requests beforehand).

Related Works

A number of recent studies have proposed methodologies for optimizing shared ES at the end-user side. Specifically, [82] presents a game-theoretic approach to managing a shared ES where users are competing for limited capacity and [83] presents a coalition game formulation for the sizing, operation, and cost allocation of a shared ES with multiple investors. Additionally, [84] presents a Nash bargaining based benefits sharing model for energy cooperation between users and a CES and is focused on the presence of 'cheaters' within the system, attempting to gain additional benefits by providing dishonest information. Centralized control of such a shared ES is studied in [85], but the solution method does not scale with the number of participants and is approximated instead. The authors of [86] present a reinforcement learning approach to manage the operation of an ES under uncertain conditions stemming from wind generation. In [87], a stochastic optimization is formulated to manage the operation of multiple shared ES systems and the performance of their proposed control policy is compared to the deterministic optimal solution via numerical experiments; however, there is no theoretical performance guarantee (i.e., bounding the gap between the cost of the deterministic optimal solution and the cost of the proposed policy). Papers [88] and [89] also study shared ES strategies, and both make use of models that disallow users to increase or decrease their allotted capacity in real-time. Similarly, [90] presents a business model for

a shared ES that promotes diverse charging/discharging schedules, but the users' capacity reservations are constrained to remain constant across days, thus limiting flexibility. Additionally, [91] studies a posted price mechanism for energy customers arriving in an arbitrary manner and choosing to either purchase a certain amount of energy based on the posted price, or leave without buying. The mechanism has similarities to the one in this chapter; however, [91] focuses on the case of transactive electric vehicle charging rather than scheduling the charging/discharging of a CES.

There are two papers closest to our work. First, [92] presents a distributed combinatorial auction approach to schedule capacity, charging, and discharging power for a shared ES. In this work, the solution method is allowed to violate the ES's total capacity limit and the over-capacity energy must be purchased from the local grid. Second, [93] presents a pricing mechanism to sell 'virtualized' portions of a shared ES each day. In this work, the prices are selected to be constant for each optimization period, which is simple to implement but limits the ability to promote diverse charging/discharging patterns from the users in real-time. Different from [92] and [93], our goal is to present a scheduling heuristic that never violates CES constraints (i.e., does not allocate more capacity than the CES has available and then purchase the over capacity power from the local grid) and makes use of dynamically updated prices that increase and decrease depending on the current utilization of the CES (i.e., dynamically increases prices at times when utilization is high to discourage usage and decreases prices at times when utilization is low to promote usage).

3.2 System Model

CES Manager's Objective

In this section, we describe the problem facing the manager of a community energy storage (CES) system attempting to optimize the energy storage (charging and discharging) schedules for a group of diverse users. Specifically, the objective of the CES manager is to maximize the total value gained by all of the users by optimizing the usage of the system and incentivizing diverse user schedules in order to maximize the benefits delivered by the capacity-limited CES.

In the following, we consider a singular CES that is co-located with potential users of the system in a neighborhood, shopping center, or business park (we note that this work can be readily extended to account for numerous energy storage systems throughout a given area). We assume that each user of the CES has the physical infrastructure in place to charge and discharge the CES at any time and each user may or may not be equipped with behind-the-meter renewable generation. Additionally, we assume that each user has the ability to communicate with the CES manager to submit requests to charge and discharge.

Over the time span $t = 1, \ldots, T$, the CES manager receives $n = 1, \ldots, N$ requests to use the shared energy storage system. We note that N is a priori unknown to the CES manager as the CES users are inherently stochastic agents and the CES manager does not know how many requests will be submitted in the time span. In this work, each request n to use the CES is in the form of a temporal charging and discharging profile. Specifically, users submit potential schedules for charging the CES, storing the charged power for a duration, and then discharging the CES at a future time as described in Section 3.2.

The job of the CES manager is to either accept and allocate storage capacity to each

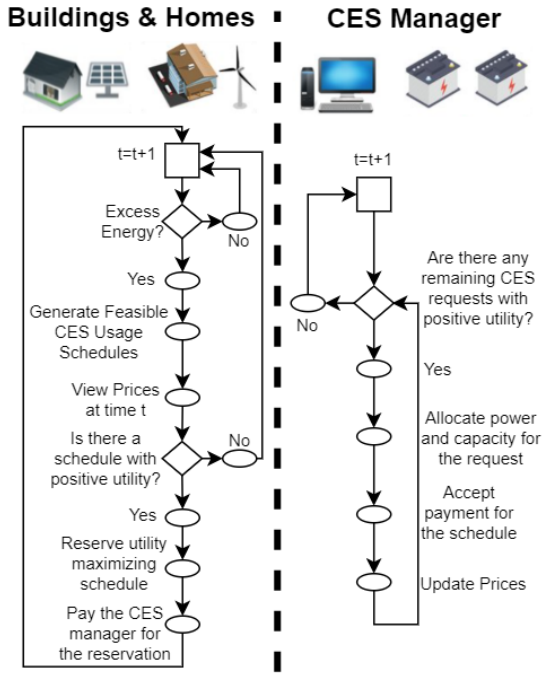

Figure 3.1: Left: System interactions for CES users (buildings and homes). Right: System interactions for the CES Manager.

request n, or to deny the request. Furthermore, due to the stochastic nature of distributed renewable generation and unknown factors affecting users' power consumption, we assume that the users cannot submit their charging/discharging requests far in advance and the CES manager cannot create the usage schedule ahead of time. Rather, users submit charging/discharging requests to the CES manager at random times throughout the time span and the CES manager must make the scheduling decisions immediately, so the users can begin charging and discharging the CES. This means that the allocation algorithm must work in real-time and without knowledge of future requests.

Charging and Discharging Schedules

Each CES request begins when a user would like to store some energy (e.g., from cheap electricity rates or from excess renewable generation) and make use of it at a later time. A user submitting a storage request may benefit from multiple possible

charging/discharging schedules, each providing a different value to the user. The user provides the CES manager with the list of such schedules and the value she associated to each of them. In the end, only a single schedule may be accepted for the user's charging and discharging of the CES. For example, if a user has excess solar generation available from 1:00pm-2:00pm and wants to charge the CES, then they could potentially benefit from discharging the power at numerous time periods later in the day, e.g., discharging 5:00pm-6:00pm, 5:30pm-6:30pm, or 6:00pm-7:00pm, etc., each providing different value to the user. Let us define the mathematical notation associated with each request. At time t_n^-, when the user submits a CES usage request, a set of potential charging and discharging schedules, $\mathcal{S}_n$, is created for request n. Each potential charge/discharge schedule $s \in \mathcal{S}_n$ has the following parameters:

1. t_n^-: The start time for all potential charge/discharge schedules for request n.

2. t_{ns}^+: The end time for charge/discharge schedule s for request n (Note that the potential charge/discharge schedules need not share an end time).

3. $i_{nsc}(t)$: The CES charging power profile for request n in feasible schedule s at time t. Positive values of $i_{nsc}(t)$ denote that the user is charging the CES and negative values of $i_{nsc}(t)$ denote that the user is discharging from the CES. Note that $i_{nsc}(t)|_{t=1,\ldots,T}$ describes the complete power profile across the entire time span that is to be *charged into* and *discharged from* the CES by user n in schedule s.

4. $i_{nse}(t)|_{t=1,\ldots,T}$: The CES capacity that must be reserved for request n in feasible schedule s across the time span in order to serve the schedule's charging profile.

5. v_{ns}: The value of potential schedule s to the user who submitted request n. This value is described in detail at the end of this subsection.

Example: Consider the following simple example with a user submitting a request to charge the CES in the time period 8:00am-9:00am at 5kW and then discharge from the CES in the time period 10:00am-11:00am at 5kW and values this schedule at \$0.50. Furthermore, let us assume that $t = 0$ corresponds to 8:00am, $t = 1$ corresponds to 9:00am, $t = 2$ corresponds to 10:00am, and $t = 3$ corresponds to 11:00am (Note that 1 hour intervals are for simplicity of the example and an implementation would use smaller intervals, 1min, 5min, etc.). As such, the requested schedule's parameters are as follows:

1. Start time: $t_n^- = 0$
2. End time: $t_n^+ = 3$
3. Charging profile (kW): $i_{nc}(t)|_{t=0,1,2,3} = 5, 0, -5, 0$
4. CES capacity profile (kWh): $i_{ne}(t)|_{t=0,1,2,3} = 5, 5, 5, 0$
5. User valuation (\$): $v_n = 0.50$

As we will see, by exploiting of the CES capacity profile $i_{nse}(t)|_{t=1,\ldots,T}$ and the charging power profile $i_{nsc}(t)|_{t=1,\ldots,T}$, our algorithm allows the CES manager to optimize the usage of the CES to avoid overutilization at popular times, underutilization at unpopular times, and to incentivize diverse charge/discharge patterns such that users' requests cancel one another. Additionally, we note that there is no restriction on how many requests per day that a user can submit. If a user submits a request at 8:00am, they can submit multiple other different requests later in the day that would be independent of their earlier requests.

Before we move on, let us discuss how the valuations v_{ns} can potentially be assigned by the users. As stated previously, there are various potential strategies for energy consumers to make use of CES. For example, users can employ the CES to shift their

electricity purchases to take advantage of inexpensive electricity rates during off-peak hours, or users can employ the CES to store locally generated renewable energy and use it at a later time. In all cases, in order for a user to choose to make use of the CES instead of defaulting to purchasing electricity from the grid, there must be an incentive to do so. In this work, we assume that the users are incentivized via cost savings; specifically, a user will only request a charge/discharge profile from the CES if the total cost that the user must pay to the CES manager is less than the cost of purchasing the same energy from the grid. The value v_{ns} is equal to the the magnitude of such cost savings as discussed next.

The proposed CES scheduling heuristic requires the users' submissions of their valuations of each potential CES schedule (potential charge/discharge schedule parameter 5 listed previously). For the purposes of this work, we assume the users' motivation to use the CES is to store excess solar energy (that was generated on-site at no cost) to use during later time periods or to charge the CES using inexpensive grid energy and discharge from the CES later in the day to avoid expensive electricity rates. For on-site solar usage, a user's valuation of potential schedule s is equivalent to the cost of electricity from the grid that is replaced by the stored solar:

$$v_{ns} = -\sum_{t} p_{\text{grid}}(t) i_{nsc}(t)|_{i_{nsc}(t)<0} \tag{3.1}$$

where $p_{\text{grid}}(t)$ is the price of electricity from the grid at time t and the negative values of $i_{nsc}(t)$ are the discharging power from the CES. If the user wanted to charge the CES during cheap electricity rates and discharge during expensive electricity rates, the

valuation (e.g., cost savings) of such a CES schedule would be calculated as:

$$v_{ns} = -\sum_t p_{\text{grid}}(t) i_{nsc}(t)|_{i_{nsc}(t)<0} - \sum_t p_{\text{grid}}(t) i_{nsc}(t)|_{i_{nsc}(t)>0}. \tag{3.2}$$

Consider the following example where a user would like to store 5 kWh of locally generated solar energy in the CES from 3:00pm to 5:00pm and then discharge 5 kW from 5:00pm to 6:00pm. Furthermore, assume the local grid's electricity rate from 5:00pm to 6:00pm is 0.11 ($/kWh). As such, in order for the user to prefer the CES instead of purchasing the electricity from the grid, the total cost for utilizing the CES must be less than 5 kW $\times$ 1 hour $\times$ 0.11 $/kWh = $ 0.55. In other terms, we can say that the user values that specific CES charge/discharge profile at $0.55.

CES Constraints

The community energy storage system has three parameters that constrain its operation[1]: 1) the CES can store up to $\hat{E}$ kWh at any given time, 2) the CES's maximum charging power $\hat{P}_c$ kW, and 3) the CES's maximum discharging power $\hat{P}_d$ kW. At any given time, the total stored energy, total charging power, and total discharging power of all the users combined must be less than the aforementioned parameters.

Charging and Discharging Cancellation

One important characteristic of the energy storage scheduling problem is that different users' requests to charge and discharge the CES can occur during the same time period, thus resulting in charge/discharge cancellations. As mentioned in the introduction, a

[1]We note that we do not include a battery model nor degradation in this work; however, this could easily be added to the framework. Any battery model limitations would reduce the number of feasible charging schedules $s \in \mathcal{S}_n$ and degradation costs could be included in the users' payment calculation (i.e., an extra term could be added to the payment $\tilde{p}_{ns^*}$ calculated in line 9 of Algorithm 1 to account for degradation costs).

charge/discharge cancellation occurs when user A commits to charging the CES at a given time and user B commits to discharging the CES at the same time, thus effectively cancelling each other's power usage of the CES at that time and allowing other users access to charge/discharge at that time slot.

The importance of charging and discharging cancellations is twofold. First, the occurrence of a charging and discharging cancellation decreases the total power being charged/discharged from the CES (recall the CES has a maximum power constraint); therefore, allowing other users access to that time slot. Second, a charge/discharge cancellation eliminates the usage of the CES altogether and instead users within the community are providing power to one another directly. That is, locally generated renewable power that would have been injected into the grid or stored in the CES is instead being used immediately by another user within the community.

Offline and Online Problem

In the body of this work, we first formulate the energy storage scheduling problem as an *offline* optimization and then use the offline problem to aid the design of a heuristic to solve the *online* problem. In the offline case, we assume the CES manager is clairvoyant and knows the entire sequence of N energy storage requests over the time span $t = 1, \ldots, T$. As such, the offline CES manager can create the optimal schedules for the energy storage requests and can achieve maximal value. However, the reality is that the CES manager does not know the users' desired charging and discharging times and storage capacity needs in advance. Instead, the energy storage requests are revealed one-by-one throughout the time span meaning that an online solution method is required for real world implementation. Additionally, the energy storage scheduling problem has obstacles that are not easily overcome in many online heuristics; namely, the lack of accurate statistics for the users' energy storage requests as there are many exogenous

factors that directly affect the time and capacity of such requests (e.g., stochastic renewable generation and weather affect the time and capacity of energy storage requests and random human behavior affects desired discharging times). As such, in the following we present an online solution that can account for *adversarially* chosen sequences of energy storage requests and still yield utility that is within a constant factor of the clairvoyant offline solution. Let us first state the offline problem.

Offline Problem Formulation

The state of the CES at any time t can be fully described by the following two variables: $y_e(t)$ the total energy capacity that is reserved at time t summed across all requests and $y_c(t)$ the total charging power that is scheduled for time t. In order to calculate $y_e(t)$ and $y_c(t)$, we introduce the decision variable x_{ns}. Specifically, when request n to use the CES is received, the CES manager must select one of the potential schedules $s \in \mathcal{S}_n$ or deny the request altogether. As such, the CES manager sets the variable x_{ns} equal to 1 if schedule s is selected for request n and 0 otherwise. If no CES schedule is selected, the request is denied and $x_{ns} = 0,\ \forall s$.

The total demands for energy capacity and charging power, $y_e(t)$ and $y_c(t)$ respectively, are calculated as follows:

$$y_e(t) = \sum_{\mathcal{N},\mathcal{S}_n} i_{nse}(t) x_{ns}, \tag{3.3}$$

$$y_c(t) = \sum_{\mathcal{N},\mathcal{S}_n} i_{nsc}(t) x_{ns}. \tag{3.4}$$

As stated in Section 3.2, if the CES manager has full knowledge of the sequence of CES requests, the optimal schedules can be found by solving the following *offline*

optimization:

$$\max_{x} \sum_{\mathcal{N},\mathcal{S}_n} v_{ns} x_{ns} \tag{3.5a}$$

subject to:

$$x_{ns} \in \{0,1\}, \qquad \forall n \in \mathcal{N}, s \in \mathcal{S}_n \tag{3.5b}$$

$$\sum_{\mathcal{S}_n} x_{ns} \leq 1, \qquad \forall n \in \mathcal{N} \tag{3.5c}$$

$$y_e(t) \leq \hat{E}, \qquad \forall t \in \mathcal{T} \tag{3.5d}$$

$$y_c(t) \leq \hat{P}_c, \qquad \forall t \in \mathcal{T} \tag{3.5e}$$

$$y_c(t) \geq -\hat{P}_d, \qquad \forall t \in \mathcal{T}. \tag{3.5f}$$

In (3.5a), the objective is to maximize the total value of CES schedules across all requests. Constraint (3.5b) is an integer constraint on the decision variable. Constraint (3.5c) ensures that only one CES usage schedule can be selected per request. Constraints (3.5d), (3.5e), and (3.5f) enforce the energy capacity limit, charging power limit, and discharging power limit of the CES, respectively.

Furthermore, to gain insight into how to formulate an online pricing heuristic for the CES problem, the offline optimization can be examined in the dual domain[2]. Specifically, we make use of Fenchel Duality and use the dual variables u_n, $p_e(t)$, $p_c(t)$, and $p_d(t)$ [34]. The dual variable u_n corresponds to the utility gained by the user who submitted request n. That is, their valuation of their assigned energy storage schedule minus the price of that schedule that they pay to the CES manager. We note that each user's utility should be positive if they are using the CES and 0 if their request is denied. Additionally, the

[2]We note that the integer constraint (3.5b) must be temporarily relaxed in order to formulate the offline dual. However, we also note that our competitive ratio results for our online pricing mechanism are for integer allocations.

dual variables $p_e(t)$, $p_c(t)$, $p_d(t)$ are associated with the total energy capacity constraint, total charging power constraint, and total discharging power constraint, respectively. Moreover, they can be viewed as the marginal prices that the users must pay for utilizing the limited storage, charging power, and discharging power of the CES. Additionally, in the remainder of the section, the Fenchel conjugate of a function $f(y(t))$ is defined as:

$$f^*(p(t)) = \sup_{y(t)\geq 0} \big\{p(t)y(t) - f(y(t))\big\}. \tag{3.6}$$

In this work and many other online combinatorial problems, making use of Fenchel conjugate functions yields a generalized dual problem that can be used to design online solution algorithms (e.g., online packing/covering [94], online paging/caching [95], online matching [96], etc.). Namely, the conjugate functions $f^*(p(t))$ that appear in the Fenchel dual problem's objective function could account for various convex cost functions due to increasing usage of limited resources[3] or scaling penalties. Furthermore, we note that the Lagrange dual is a special case of the more general Fenchel dual problem; moreover, the Fenchel dual can be derived from the Lagrange dual problem and the conjugate definition (shown in [97]). We note that Lagrangian duality is used in similar primal-dual works [98, 99, 100]; however, in general, the Fenchel dual typically presents a better structure for the design and analysis of online primal-dual algorithms that attempt to approximate solutions for NP-hard combinatorial problems such as the one we study in this work. We refer the reader to [81, 34, 101, 39, 97] for further reading on Fenchel duality in this setting and primal-dual methods.

The dual reformulation is needed for three major reasons:

1. The primal problem is NP-Hard and cannot be solved in reasonable time even for

[3]In this work, we do not explicitly make use of cost functions for utilizing limited resources (capacity and power); however, the capacity and power constraints' costs could be viewed as zero-infinite step functions, which would yield the same Fenchel conjugates as (3.8)-(3.10).

moderately sized systems.

2. The primal problem cannot be solved without knowledge of the future energy storage requests. Thus, an online optimization that works without knowledge of the future is required.

3. To take advantage of the dual variables for the energy capacity, charging power, and discharging power constraints of the problem which can be viewed as 'marginal prices' for these limited resources. Our proposed heuristic functions then estimate these dual variables in an online fashion (without knowledge of the future) and in the Theorem and the Appendix, are shown to ensure a level of welfare that is within a ratio of the optimal offline solution. This is in the form of the "competitive ratio" discussed in the Theorem and the Appendix.

With the aforementioned dual variables and Fenchel conjugate definition, the offline Fenchel dual of (3.5a)-(3.5f) is as follows:

$$\min_{u,p} \sum_{\mathcal{N}} u_n + \sum_{\mathcal{T}} \Big[f_e^*(p_e(t)) + f_c^*(p_c(t)) + f_d^*(p_d(t)) \Big] \tag{3.7a}$$

subject to:

$$u_n \geq v_{ns} - \sum_{\mathcal{T}} \Big[i_{nse}(t) p_e(t) + i_{nsc}(t) p_c(t) - i_{nsc}(t) p_d(t) \Big], \qquad \forall s \in \mathcal{S}_n, n \in \mathcal{N} \tag{3.7b}$$

$$u_n \geq 0, \qquad \forall n \in \mathcal{N} \tag{3.7c}$$

$$p_e(t), p_c(t), p_d(t) \geq 0, \qquad \forall t \in \mathcal{T}. \tag{3.7d}$$

We note that $f_e^*(p_e(t))$, $f_c^*(p_c(t))$, and $f_d^*(p_d(t))$ are the Fenchel conjugates for the energy capacity limit, charging power limit, and discharging power limit, respectively. Recall from Section 3.2, 1) the CES can store up to $\hat{E}$ kWh at any given time, 2) the CES's maximum charging power $\hat{P}_c$ kW, and 3) the CES's maximum discharging power $\hat{P}_d$ kW.

With these variables and the dual variables $p_e(t)$, $p_c(t)$, and $p_d(t)$, the Fenchel conjugates can be written as follows:

$$f_e^*(p_e(t)) = \hat{E}\, p_e(t), \tag{3.8}$$

$$f_c^*(p_c(t)) = \hat{P}_c\, p_c(t), \tag{3.9}$$

$$f_d^*(p_d(t)) = \hat{P}_d\, p_d(t). \tag{3.10}$$

Insight on Scheduling Decisions

In order to learn how to make scheduling decisions in the online case, let us first examine the offline Fenchel dual (3.7a)-(3.7d). The constraint (3.7b) gives insight into the optimal scheduling decisions for each request n. Specifically, if the utility gained u_n from request n is negative across all potential schedules, then the request to utilize the CES is denied and u_n is set equal to 0. However, when $u_n > 0$ then the request is accepted and the charging/discharging/storage schedule $s \in \mathcal{S}_n$ to be selected is the one that returns the maximal u_n. With this in mind, we can instead use the following equation to calculate the utility of request n:

$$u_n = \max\left\{0, \max_{\mathcal{S}_n}\left\{v_{ns} - \sum_{\mathcal{T}}\left[i_{nse}(t)p_e(t) + i_{nsc}(t)p_c(t) - i_{nsc}(t)p_d(t)\right]\right\}\right\}. \tag{3.11}$$

Equation (3.11) is derived from examining the KKT conditions for the primal (described in detail in Appendix A.3). Specifically, for any request n to use the CES, there is a

dual variable $u_n \geq 0$ from constraint (3.7b) which corresponds to the utility of request n. Moreover, we know that in the offline primal and dual solutions, no schedule can be selected unless constraint (3.7b) is active for a specific schedule. As such, we can set the utility equal to the maximum of 0 (corresponding to no schedule being selected due to negative utility gain) and the RHS of (3.7b) (corresponding to utility maximizing schedule being selected). In summary, if the dual variables (CES resource prices) $p_e(t)$, $p_c(t)$, and $p_d(t)$ are known or estimated, then equation (3.11) can be used to determine which schedule gets allocated for request n or if request n is denied altogether and u_n is set to 0 (we note that in Section 3.3 we present our methodology to estimate the dual variables/CES resource prices in real-time so that (3.11) can be solved in an online fashion).

We note that in order to solve for the utility gained u_n from the offline dual (3.7a)-(3.7d) (and the offline primal (3.5a)-(3.5f)), this requires full knowledge of the requests to use the CES beforehand. However, as discussed previously in Section 3.2, the manager of the CES does not know the sequence of requests beforehand and must make scheduling decisions as they arrive without knowledge of future requests. Moreover, as we show in the remainder of the section, we never have to solve the offline dual problem as presented in (3.7a)-(3.7d), as this would require knowledge of the entire sequence of usage requests, which the CES manager does not have. Instead, we make use of dual variable update functions (3.12), (3.15), (3.16) to estimate the dual variables in real-time. Then, these dual variables are used as 'prices' for the limited resources and our algorithm selects schedules w.r.t. these prices. We can show that our estimated dual variables will always yield feasible solutions to the primal problem. This is because the dual variable update functions are carefully selected to yield values that increase as the usage of the CES increases. Then, when a constraint is about to be violated, the dual variable update functions will output values high enough such that no energy storage schedule yields pos-

itive utility, meaning that requests will be denied if constraints are going to be violated. The gap in the objective value from the original (unrelaxed) primal problem (3.5a)-(3.5f) and our online heuristic is bounded in Theorem 3.3.1.

3.3 Online CES Scheduling Heuristic

Online Scheduling via Dual Variable Updates

In the following, we present a scheduling heuristic for optimizing usage of the CES that updates the dual variables $p_e(t)$, $p_c(t)$, and $p_d(t)$ in an online fashion as requests are revealed. Then, with the estimated dual variables, the algorithm solves equation (3.11) for each request to select the utility maximizing charging/discharging/storage schedule. Moreover, the online scheduling heuristic updates the dual variables for charging, discharging, and storage based only on $y_e(t)$ and $y_c(t)$ (the total energy capacity reserved at time t and the total charging power scheduled at time t, respectively).

The online scheduling procedure for the usage of the CES is outlined in Algorithm CommunityEnergyScheduling. When a CES usage request is received, the CES manager generates a set of feasible schedules S_n and then the best schedule, s^*, is chosen in line 8. We note that our algorithm is equivalent to a posted price mechanism where all the options are enumerated with corresponding prices for each. Users simply examine their valuations for each feasible schedule, subtract the current cost of each schedule ($\tilde{p}_{ns}$), and choose the utility maximizing schedule (i.e., no complex optimization needed, they simply choose the highest value option). We note that the total price that the customer pays for their allocated schedule is calculated in line 9 of Algorithm 2 and is denoted as $\tilde{p}_{ns^*}$ After each request is scheduled or denied, the CES manager updates the dual variables with the new values for charging and discharging power as well as energy capacity (lines 11-12).

Algorithm 3 CommunityEnergyScheduling

Input: $\hat{E}, \hat{P}_c, \hat{P}_d, L_{e,c,d}, U_{e,c,d}$

Output: x, p

Define the update functions $p(y(t))$ according to (3.12) - (3.16) for energy capacity, charging, and discharging.
Initialize $x_{ns} = 0$, $y_{e,c}(t) = 0$, $u_n = 0$.
Initialize prices $p(0)$ according to (3.12) - (3.16).
Repeat for all N CES requests:
Request n is received, generate feasible charging/discharging schedules $\mathcal{S}_n$
Update dual variable u_n according to (3.11).
if $u_n > 0$ **then**
    $(s^\star) = \arg\max_{\mathcal{S}_n} \left\{ v_{ns} - \sum_{t\in[t_{ns}^-, t_{ns}^+]} \left(i_{nse}(t)p_e(t) + i_{nsc}(t)p_c(t) - i_{nsc}(t)p_d(t) \right) \right\}$
    $\widetilde{p}_{ns^\star} = \sum_{\mathcal{T}} \left[i_{nse}(t)p_e(t) + i_{nsc}(t)p_c(t) - i_{nsc}(t)p_d(t) \right]$
    $x_{ns^\star} = 1$ and $x_{ns} = 0$ for all $s \neq s^\star$
    Update total demand $y(t)$ for energy capacity and charging power according to (3.3)-(3.4).
    Update dual variables $p(y(t))$ for energy capacity, charging, and discharging according to (3.12) - (3.16).
else
    $x_{ns} = 0, \quad \forall\ s \in \mathcal{S}_n$.
end if
if $\exists s^\star$ and $x_{ns^\star} = 1$ **then**
    Allocate request n the energy capacity, charging power, and discharging power from schedule s^*.
    Request n is fulfilled by schedule s^* for the price of $\widetilde{p}_{ns^\star}$ to the requester.
else
    Deny request n from using the CES.
end if

The three main benefits are as follows: 1) the online scheduling heuristic ensures that the utility gained from each scheduled request is positive for the user, 2) the online scheduling heuristic filters out low value charging and discharging requests in order to prevent the CES from being overused, and 3) the online scheduling heuristic promotes diverse charging and discharging schedules to take advantage of charge/discharge cancellations as mentioned in Section 3.2. The underlying framework of the dual variable update heuristic is similar to that of [39], in which the authors present an auction mechanism for optimizing the usage of computer hardware at data centers for cloud computing.

In our online CES scheduling heuristic, we expand upon the specialized functions proposed in [39] that approximate the optimal dual variables in an online fashion. These dual variable functions depend on the amount of energy capacity, charging power, and discharging power that is reserved at a future time t. The update functions increase slowly at first then increase rapidly as the CES power and capacity limits are approached. Additionally, when the power and capacity limits are met, the dual variable update functions ensure that no more schedules will be allocated by outputting dual variables high enough to ensure no schedule yields positive utility, thus enforcing the hard capacity and power limits. The specialized function to update the dual variable associated with the energy capacity of the CES is as follows:

$$p_e(t) = \Big(\frac{L_e}{6}\Big)\Big(\frac{6U_e}{L_e}\Big)^{\frac{y_e(t)}{\hat{E}}},\ y_e(t) \in [0, \hat{E}], \tag{3.12}$$

where U_e and L_e correspond to the maximum and minimum value per kWh of energy capacity per time unit, respectively, across all requests. We note that the CES manager does require knowledge of U_e and L_e beforehand to calculate initial values for the dual variables and to ensure limits are not breached. The maximum and minimum valuations

are calculated as follows:

$$L_e = \min_{n \in \mathcal{N}, s \in \mathcal{S}_n} \frac{v_{ns}}{3 \sum_{t \in [t_s^-, t_s^+]} i_{nse}(t)}, \tag{3.13}$$

$$U_e = \max_{n \in \mathcal{N}, s \in \mathcal{S}_n, t \in \mathcal{T}} \frac{v_{ns}}{i_{nse}(t)}, \; i_{nse}(t) > 0. \tag{3.14}$$

In addition to the energy capacity's dual variable update function in (3.12), the dual variables for the charging and discharging power of the CES also require update functions:

$$p_c(t) = \Big(\frac{L_c}{6}\Big)\Big(\frac{6U_c}{L_c}\Big)^{\frac{y_c(t)}{\hat{P}_c}}, \; y_c(t) \in [-\hat{P}_d, \hat{P}_c], \tag{3.15}$$

$$p_d(t) = \Big(\frac{L_d}{6}\Big)\Big(\frac{6U_d}{L_d}\Big)^{\frac{-y_c(t)}{\hat{P}_d}}, \; y_c(t) \in [-\hat{P}_d, \hat{P}_c]. \tag{3.16}$$

We note that the dual variable update functions for charging and discharging power, (3.15) and (3.16), are similar to the energy capacity dual variable function (3.12) except for the domain. The energy capacity function's input values, $y_e(t)$, are nonnegative and less than $\hat{E}$. The charging and discharging functions' input values can be negative and are within the range $y_c(t) \in [-\hat{P}_d, \hat{P}_c]$. With the 3 dual variable update functions (3.12), (3.15), and (3.16), we now have the means to calculate estimates for the optimal dual variables in order to solve (3.11) in an online fashion (i.e., at the reception of each request to use the CES). The full procedure can be seen in Algorithm 2 CommunityEnergyScheduling.

The heuristic presented in CommunityEnergyScheduling attempts to solve an online scheduling problem without full knowledge of the sequence of requests. As stated

before, we are able to compare the total welfare generated from our online heuristic to the total welfare generated by an omniscient offline CES manager. The comparison that we make is in the form of a competitive ratio. An online heuristic is said to be α-competitive when the ratio of welfare generated by the omniscient offline solution to the welfare generated by the online heuristic is bounded by $\alpha \geq 1$. The competitive ratio, α, is defined as $OPT/ALG_{worstcase} \geq 1$, where OPT is the welfare generated by the offline optimal solution and $ALG_{worstcase}$ is the worst-case welfare generated by the online algorithm. A value of 1 means the algorithm performs optimally and higher values of α indicate worse performance. In this work, we build upon results from [39] (and previous work [9, 8, 4]) and present a competitive ratio that accounts for the cancellation of complimentary resources (e.g., charging and discharging power, which previous works could not account for). For the following results, we assume that each CES request utilizes a small amount of the charging/discharging power and energy capacity of the CES to ensure that one schedule cannot prohibit numerous future schedules and that the ratios of users' maximum valuation to minimum valuation for charging and discharging power are equal, i.e., $\frac{U_c}{L_c} = \frac{U_d}{L_d} = \frac{U_{c,d}}{L_{c,d}}$ (to yield a singular $\alpha_{c,d}$ for both the charging and discharging of the CES).

Theorem 3.3.1 *The community energy storage system's schedules generated by* CommunityEnergyScheduling *in Algorithm 1 are α-competitive in welfare over N usage requests where $\alpha = \max\{\alpha_e, \alpha_{c,d}\}$ and α_e and $\alpha_{c,d}$ are defined as follows:*

$$\alpha_e = 2\ln\left(\frac{6U_e}{L_e}\right),$$
$$\alpha_{c,d} = 2\ln\left(\frac{6U_{c,d}}{L_{c,d}}\right).$$

Proof. The full proof can be found in Appendix A.4.

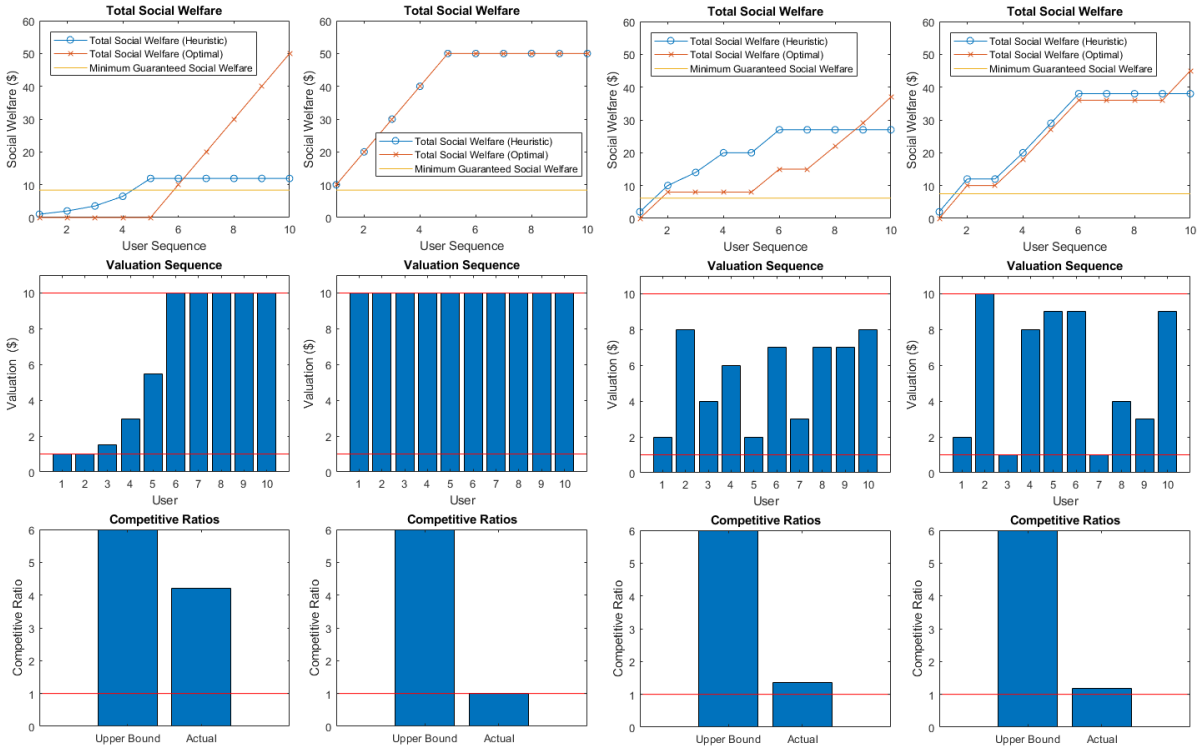

Figure 3.2: Simple example system with 10 user requests to use a 5kWh shared battery system. Each column corresponds to a different arrival sequence of user valuations.

3.4 Numerical Results

In the following, we present two different numerical results to showcase our heuristic. First, we describe an example system in Section 3.4 which explicitly details the CES requests' valuations and compares the social welfare generated from our proposed heuristic to the optimal offline case as the users submit requests to use the shared battery. We then present a larger case study for a shared battery system serving commercial customers in California in Section 3.4.

Intuitive Example

In order to showcase the details of our heuristic, we make use of a specific example setup. Namely, we consider a shared battery system that has a maximum charging rate of 5kW, maximum storage of 5kWh, and a maximum discharging rate of 5kW. Furthermore, we consider 10 unique users who want to purchase the exact same charging, storage, and discharging schedule and are interested in no other schedules (the specific schedule of interest is to charge 1kW from 8-9am, store 1kWh from 9-10am, and discharge 1kW from 10-11am). These 10 unique users arrive sequentially one after another and submit their bids to purchase the charging, storage, and discharging schedule (the specific times that each user submits their request are irrelevant as long as they are submitted sequentially and all before 8am). For this example, we assume the users' valuations are within \$1 and \$10. Due to the constraints of the shared battery system, it is clear that only 5 of the 10 users will be able to use the battery for that specific charging, storage, and discharging schedule. In the offline case, the optimal solution yielding maximal social welfare will select the 5 users with the highest valuations to use the shared battery. However, since the users submit their bids sequentially and their valuations are unknown a priori, our heuristic attempts to emulate the offline solution via dynamic prices that increase as the battery usage increases, thus filtering out users with low valuations.

In Figure 5.7, we present the results of 4 different user valuation sequences (each column corresponds to a different sequence of user valuations). Row 1 presents the social welfare results of our heuristic and the optimal offline solution. Row 2 presents the users' valuations (in order). Row 3 presents the competitive ratio upper bound from our theoretical results in addition to the actual competitive ratio for that column's request sequence. From left to right: Column 1 portrays the worst case user valuation sequence. This is because each user's valuation was carefully selected to equal the current price

of the schedule generated by our heuristic (i.e., each of the first 5 users have the minimum valuations that our pricing heuristic will accept while the last 5 users have the maximum valuation, thus leading to the worst possible competitive ratio). Note that the actual competitive ratio in this case is still below the theoretical upper bound. Column 2 portrays one of the many valuation sequences where the heuristic matches the offline optimal solution (i.e., competitive ratio = 1). Columns 3 and 4 present randomly generated valuation sequences (i.e., user valuations were drawn from a uniform distribution between \$1 and \$10) to showcase that our heuristic often yields competitive ratios close to 1.

Additionally, we compare each of the 4 arrival sequences in Fig. 5.7 to a First-Come-First-Serve (FCFS) heuristic that is the status quo scheduling method for any new CES implementation. Table 3.1 presents the percentage of the offline optimal welfare that is generated by both our Algorithm 1 and a FCFS heuristic.

	Sequence 1	Sequence 2	Sequence 3	Sequence 4
ALG1	24%	100%	73%	87%
FCFS	24%	100%	57%	66%

Table 3.1: Percentage of offline optimal welfare generated by Algorithm 1 and First–Come-First-Serve.

California Case Study

In this section we present results from a community energy storage system in California. Specifically, there are 10 loads (presented in Fig. 3.3.A) sourced from a commercial building load dataset [102]. The publicly accessible dataset [102] contains hourly load profile data for commercial building types and residential buildings in all TMY3 locations in the United States. The Typical Meteorological Year 3 (TMY3) provides one year of hourly data that best represents median weather conditions over a multiyear period

for a particular location. Across the 10 day time span, January 1st - January 10th, we also assume that each load is equipped with behind-the-meter solar generation that they would like to charge and discharge the CES with. The normalized solar generation for the California location [40] is presented in Fig. 3.3.B. We assume that each building is equipped with solar generation capacity to fulfill 80% of their peak load at maximum rating. Fig. 3.3.D presents the 10 loads once the solar generation is subtracted. Note that negative power means that the location is producing more power than is being consumed. We assume that all 10 buildings are able to use a 2500 kWh community energy storage system with maximum charge and discharge rates of 500 kW. Furthermore, we assume that the 10 buildings are connected to the local grid and pay time-of-use electricity rates [103] for energy that is not provided by their solar generation. The electricity rate used is the PG&E E-19 structure for buildings <1000 kW max demand and is shown in Fig. 3.3.C. For the purposes of this work, we do not consider net energy metering for the locations injecting excess solar generation back into the local grid as sending excess energy to the CES is preferred.

As noted in Section 3.2, the incentive for the buildings to use the CES comes from storing excess solar generation and using it at a later time. As such, whenever a building detects that it is producing more power than it needs, it submits a request to store that excess power in the CES. Specifically, on an hour-by-hour basis, each location submits requests to store their excess energy in the CES. In order to accomodate this, the CES manager limits the number of feasible charging/storage/discharging schedules to 96 for each request. Namely, all the excess generation that the building wants to inject into the CES during hour t must be discharged at the same rate during a future hour in the range $[t, t + 96]$ (i.e., in the next 4 days). The valuation for each of the 96 schedules is calculated via equation (3.1) (i.e., the predicted cost savings from using stored energy versus purchasing energy from the grid). As portrayed in Fig. 3.3.E, the total load

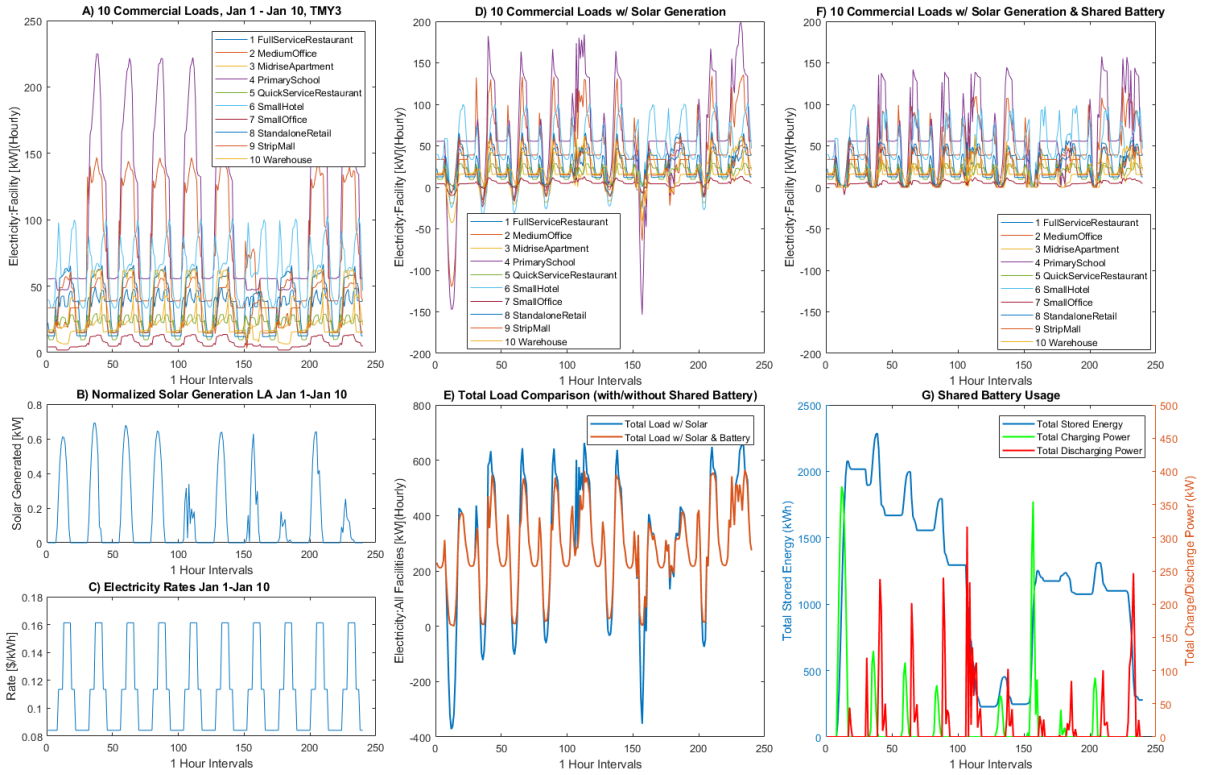

Figure 3.3: Simulation Results for California Test Case.

of all 10 buildings is greatly affected by the CES usage. Specifically, the cumulative load no longer goes negative (the red curve in 2.E), meaning that the buildings are not injecting solar back into the local grid. Instead, they are storing that power and using it to reduce peak demands at later times. This helps reduce electricity costs for the buildings in addition to reducing the stress on the load grid from injecting the excess solar generation. Last, in Fig. 3.3.G, we present the charging, discharging, and total energy stored in the CES throughout the 10 days.

Additional Case Study

In Figure 3.4, we present results for the same energy community as in Section 3.4; however, we include a large hospital as one of the loads (hospital also from dataset

[102]). As seen in plots A) and B) of Figure 3.4, the load (A) and net load (B) of the hospital are significantly larger than the other 10 loads. While it is possible that the large hospital might dominate usage of the CES, due to the fact that valuations are bounded per time slot for all users and users only purchase CES schedules if they are cheaper than the current grid electricity prices, all users end up with a fair chance at CES usage. Additionally, the smaller users have slightly different load patterns than the large hospital yielding many charge/discharge cancellations. Moreover, if the smaller users submit their CES requests before the hospital, they could exclude the hospital's large charge/discharge requests due to the capacity and power constraints. Last, we note that our results can begin to degrade if the size of the CES is not large enough to adequately supply all of the users. Specifically, in the case of the 10 small users and the large hospital, we assume that any CES manager would install a large enough battery to give all users a fair chance at usage. In plot C) of Figure 3.4, we show the total net load of all 11 users with no CES, with a 2500kWh battery (same as in Section 3.4) and a 5000kWh battery. We note that the 2500kWh battery was not large enough to store all of the excess solar generation; however, the 5000kWh battery nearly stored all excess solar generation.

3.5 Conclusion

We presented a solution to the problem facing the manager of a community energy storage system attempting to schedule the charging/discharging/energy storage of the system. We presented an online heuristic that updates dual variables in real-time as a solution to the problem. The heuristic acts as a pricing mechanism to ensure the CES yields positive utility and promotes charge and discharge cancellations to reduce the CES's usage at popular times. The heuristic is able to handle the inherently stochastic

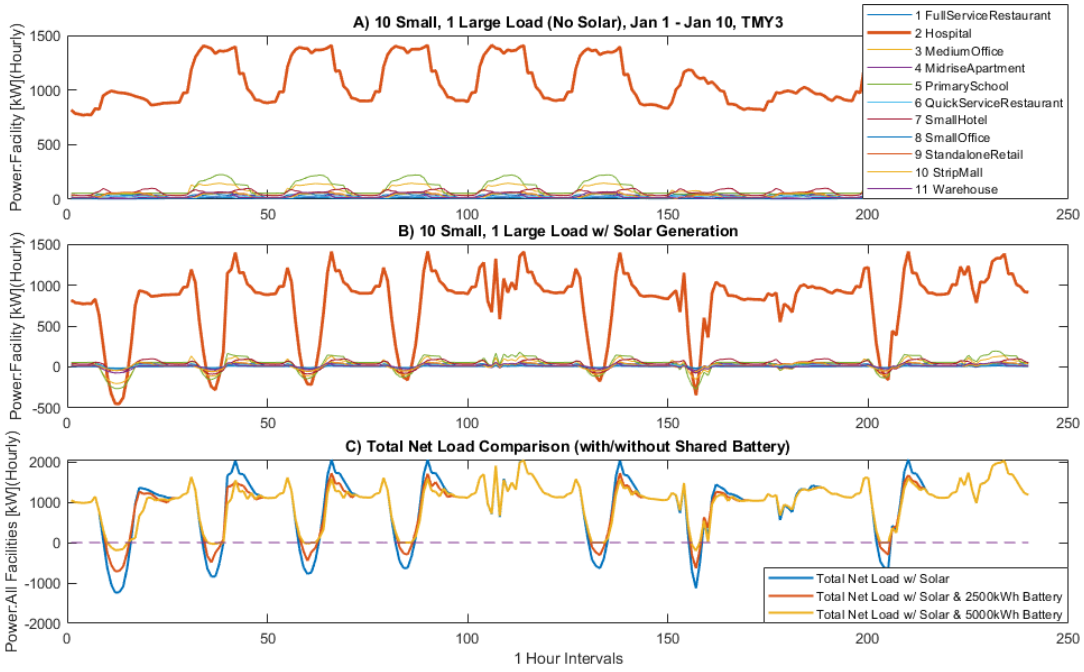

Figure 3.4: Load profiles for energy community and a large hospital.

nature of the requests to charge and discharge from the CES (stemming from weather uncertanties and randomness in users' electricity usage patterns). The heuristic can handle adversarially chosen request sequences and will always yield total welfare within a factor of $\frac{1}{\alpha}$ of the offline optimal welfare. An intuitive example was presented to showcase the heuristics performance for various request sequences and a larger case study was presented for 10 commercial buildings sharing a CES. The proposed algorithm performs well if the CES is sized large enough to adequately serve the users of the system. The results in Fig. 3.4 demonstrate the issues that arise when the CES's capacity is sized incorrectly and the effectiveness of the CES is limited.

Chapter 4

Safe Learning for Real-Time Electricity Pricing with Grid Constraints

4.1 Introduction

In order to integrate the increasing volume of intermittent renewable generation in modern power grids, aggregators are exploring various methods to manipulate both residential and commercial loads in real-time. As a result, various demand response (DR) frameworks are gaining popularity because of their ability to shape electricity demand by broadcasting time-varying signals to customers; however, most aggregators have not implemented complex DR programs beyond peak shaving and emergency load reduction initiatives. One reason for this is the customers' unknown and time-varying responses to dispatch signals, which can lead to economic uncertainty for the aggregator and reliability concerns for the grid [104]. The aggregator could explicitly request price sensitivity infor-

This work was supported by NSF grants #1847096 and #1737565 and UCOP Grant LFR-18-548175.

mation from its customers; however, this two-way negotiation has a large communication overhead and most customers cannot readily characterize their price sensitivities, and even if they could, they might not be willing to share this private information. As such, aggregators prefer the 1-way *passive* approach because it does not require any real-time feedback from the customer and it does not require new communication infrastructure for reporting preferences (e.g., a web portal, phone application, etc.). With this in mind, future load shaping initiatives for renewable integration (i.e., more complex objectives than peak shaving) need to be able to *passively* learn customers' response to dispatch signals only from historical data of past interactions [105].

Recently, much work has been done for aggregators attempting to learn customers' price responses whilst implementing peak shaving DR programs. The authors of [106] present a data-driven strategy to estimate customers' demands and develop prices for DR. In [107], the authors use linear regression models to derive estimations of customers' responses to DR signals. Similarly, [108] develops a joint online learning and pricing algorithm based on linear regression. In [109], the authors present a contract-based DR strategy to learn customer behavior while broadcasting DR signals. The authors of [110] present an online learning approach based on piecewise linear stochastic approximation for an aggregator to sequentially adjust its DR prices based on the behavior of the customers in the past. In [111], the authors develop a risk-averse learning approach for aggregators operating DR programs. In [112], a learning algorithm for customers' utility functions is developed and it is assumed that the aggregator acts within a two-stage (day-ahead and real-time) electricity market. Additionally, the authors of [113] present a learning framework for forecasting individual loads and DR capabilities and find that users with more variable consumption patterns are more effective DR participants. Using a similar framework as in this work, a multi-armed bandit (MAB) formulation is used in [114, 115] to determine which customers to target in DR programs.

In addition to learning how customers respond to DR signals, an aggregator must also consider power system constraints to ensure reliable operation (e.g., nodal voltage, transformer capacities, and line flow limits). In real distribution systems, it is critical that these constraints are satisfied at every time step to ensure customers receive adequate service and to avoid potential grid failures even without sufficient knowledge about how customers respond to price signals (i.e., in early learning stages) [116, 117]. One paper similar to ours that considers these realistic constraints, [118], presents a least-square estimator approach to learn customer sensitivities and implements DR in a distribution network. The authors of [118] show that their least-square algorithm's parameter estimation error converges to zero over time, thus the algorithm's regret is sublinear while also accounting for the distribution network's constraints.

Similar to the aforementioned papers, the work presented in this chapter considers the problem of an aggregator *passively* learning the customers' price sensitivities while running a load shaping program. However, our approach permits more complex load shaping objectives (e.g., tracking a daily target load profile) and varies in terms of both load modeling and learning approach from all the above papers. Specifically, we present a multi-armed bandit (MAB) heuristic akin to Thompson sampling (TS) to tackle the trade-off between exploration of untested price signals and exploitation of well-performing price signals while ensuring grid reliability. It is important to note that the standard TS heuristic cannot guarantee that grid reliability constraints are upheld during the learning process. As such, we present two *modified* versions of TS while retaining the fundamental principles TS is based on. Furthermore, we provide discussion on how the constraints are upheld (i.e., operational reliability) for the modified heuristics, discussion on the performance of the heuristics compared to a clairvoyant solution, and simulation results highlighting the strengths of the method.

In our work, we make use of a load clustering technique in order to *exploit the known*

physical structure of the problem and make use of our prior knowledge of how flexible electric appliances behave to lower the problem dimensionality. We note that grouping (clustering) loads for dimensional reduction is common in DR literature [119]. Some pertinent examples include [120] where the authors aggregate heterogeneous thermostatically controlled loads (TCLs) using an LTI "bin" model, [121] where the authors group EVs into "classes" depending on their charging availability, [122] where the authors present a load profile clustering method for load data classification based on information entropy, piecewise aggregate approximation, and spectral clustering, [123] where the authors present aggregate models for classes of TCLs that include statistical information of the population, systematically deal with heterogeneity, and account for a second-order effects, [124] where the authors propose a clustering technique for determining natural segmentation of customers and identification of temporal consumption patterns in the smart grid domain, and [125] where the authors develop cohorts, or groups of consumers with similar consumption patterns, from correlations between daily loads.

The **main contributions** of this work are as follows:

- We use the multi-armed bandit (MAB) framework to model the stochastic and unknown nature of customers' daily aggregate response to electricity prices.
- We make use of an appliance clustering methodology to provide a mesoscopic model of the price responsive demand of a large population of flexible appliances and reduce the dimensionality of the learning problem.
- Our learning framework can account for daily variabilities and realistic grid reliability constraints that are critical for daily operation in spite of uncertainty about customers' price response.
- We present two modified heuristics based on Thompson sampling (TS) as solutions

to the constrained learning and pricing problem.

- We provide a performance guarantee in the form of a regret bound and discussion on the reliability guarantees of the approach as well as a distribution system case study demonstrating the efficacy of the approach.

The remainder of Chapter 4 is organized as follows: Section 4.2 presents the aggregator's daily objective as well as the customers' load model. Section 4.3 describes the multi-armed bandit formulation for the electricity pricing problem, presents the modified TS heuristic, and discusses its performance and reliability. Section 4.4 presents simulation results that showcase the efficacy of the approach. The online appendix [3] contains a table of notation and proofs.

4.2 Problem Formulation

The Aggregator's Objective

The aggregator's main goal is to select dispatch signals to manipulate customer demand according to a given optimization objective that varies daily. Specifically, we consider the case where the aggregator broadcasts a dispatch signal $\mathbf{p}_\tau = [p(t)]_{t=1,\dots,T}$ to the population of customers each day (we use $t = 1, \dots, T$ to index time of day and $\tau = 1, \dots, \mathcal{T}$ to index days). The set of dispatch signals available for use by the aggregator is denoted as $\mathcal{P}$. In this chapter, without a loss of generality, we will assume that the dispatch signal sent to customers for load shaping purposes is a real-time pricing (RTP) signal[1].

[1]The reader should note that this choice is not fundamental to the development of the modified learning heuristics we present in this chapter. It only allows us to provide a concrete characterization of the response to dispatch signals by mathematically modeling the customers as cost-minimizing agents equipped with home energy management systems in Section 4.2.

The aggregator's cost function could cover a broad range of goals including (but not limited to) manipulating the population's load to match a target profile, minimizing the distribution grid's electricity cost from the regional retailer, or solving for the dispatch of multiple generators, if a market is operated at the distribution system level.

In this work, on each day τ, we assume the aggregator's cost function is a fixed and known nonlinear function $f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)$ that depends on the load profile $\mathbf{D}_\tau(\mathbf{p}_\tau)$ of the population in response to the daily broadcasted price $\mathbf{p}_\tau$ and a random exogenous parameter vector $\mathbf{V}_\tau$[2]. The population's load profile on day τ, $\mathbf{D}_\tau(\mathbf{p}_\tau)$, is a $T \times 1$ vector with the t^{th} element corresponding to the population's power demand during time period t. The exogenous and given $T \times 1$ vector $\mathbf{V}_\tau$ varies daily and can correspond to a daily target profile reflecting renewable generation forecasts, weather predictions, and grid conditions. We consider the exogenous vectors to be i.i.d. drawn from a distribution defined on a finite sample space $\mathcal{V}$, with each outcome drawn with a nonzero probability. We would like the reader to note that this assumption is only made for convenience for our theoretical regret performance guarantee in Theorem 4.3.1. In a real-world implementation, the daily exogenous parameters could be correlated across days (e.g., due to weather, seasons, weekday/weekend, etc.). However, this correlation does not affect the safety guarantees of our algorithm or its applicability (i.e., it only affects our formal regret results).

The aggregator must ensure that the broadcasted price signals do not result in load profiles that violate distribution system reliability constraints (e.g., nodal voltage, transformer capacities, or line flow limits). As such, if the aggregator had full information about how the population responds to price signals (i.e., full knowledge of $\mathbf{D}_\tau(\mathbf{p}_\tau)$), the

[2]We note that the function f need not have a closed form representation and thus can represent the solution of an economic dispatch problem with multiple generators, which can still be handled through our framework. However, without loss of generality and purely for brevity of notation, here we focus on common distribution systems which usually lack two-sided markets, and thus we focus on load profile manipulation for renewable integration and distribution system protection.

aggregator can solve the following optimization problem on day τ to select the optimal price $\mathbf{p}_\tau^\star$:

$$\mathbf{p}_\tau^\star = \underset{\mathbf{p}_\tau \in \mathcal{P}}{\arg\min}\, f\big(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau\big) \tag{4.1}$$

$$\text{s.t.} \qquad g_j\big(\mathbf{D}_\tau(\mathbf{p}_\tau)\big) \leq 0, \;\; \forall j = 1, \ldots, J \tag{4.2}$$

where $g_j(\cdot)_{j=1,\ldots,J}$ is used to represent the reliability constraints for the distribution system. We note that these general constraints need not be linear for the proposed Thompson sampling approach.

However, as explained in the introduction, knowledge of customers' price response is unavailable to the aggregator. Recall, 1) the aggregator does not want to directly query customers for their price sensitivities, 2) most customers cannot readily characterize their price sensitivities, and 3) customers might not be willing to share this private information. Accordingly, the aggregator needs a method to sequentially choose daily price signals to simultaneously 1) control their daily incurred cost; 2) learn the customers' price response models; and 3) ensure the distribution system constraints are not violated at any time.

Distribution System Operational Constraints

As stated previously, there are various operational constraints within a distribution system that should be met in order to ensure adequate service for customers and to prevent grid failures. In the aggregator's daily optimization in Section 4.2, the constraints are formulated as general functions $g_j(\cdot)_{j=1,\ldots,J}$. Specifically, these general functions represent distribution system parameters (i.e., the nodal voltage $u_\tau(t)$ and power flow through

distribution lines $f_\tau(t)$) that should obey the following constraints:

$$u_\tau(t) \geq u^{min}, \ \ \forall t, \tau, \tag{4.3}$$

$$u_\tau(t) \leq u^{max}, \ \ \forall t, \tau, \tag{4.4}$$

$$f_\tau(t) \leq S^{max}, \ \ \forall t, \tau, \tag{4.5}$$

where u^{min}, u^{max}, and S^{max} correspond to the lower voltage limit, upper voltage limit, and power flow limit, respectively, for the population's connection to the distribution grid. We note that $u_\tau(t)$ and $f_\tau(t)$ can be easily derived from the population's load profile $\mathbf{D}_\tau(\mathbf{p}_\tau)$ (See Section 4.4). Now that we have described the aggregator's objective and the distribution system's constraints, we next describe the customers' load model as well as their price response model.

Load Flexibility Model

It is hard to approach the problem of learning the response of a population of customers to complex dispatch signals such as RTP as a complete "black box problem", i.e., by just observing the broadcasted price and the load response. There are many reasons for this, including 1) the existence of random or exogenous parameters which lead to variability in the temporal and geographical behavior of electricity demand; 2) the variability of the control objective on a daily basis (e.g., due to randomness in renewable generation outputs, market conditions, or baseload); and 3) the small size of the set of observations that one can gather compared to the high dimensional structure of the load (there are only 365 days in a year, so only 365 sets of prices can be posted). Hence, in this chapter, we will be *exploiting the known physical structure* of the problem and making use of our statistical prior knowledge of how the load behaves to lower the problem dimensionality.

Specifically, to lower the dimensionality for the learning problem, we explore the fact

that flexible loads only show limited number of "load signatures" (justified due to the automated nature of load response through home energy management systems, the limited types of flexible appliances, and the common electricity usage patterns that emerge from electricity customers as shown in [113, 126]). Let us assume that electric appliances can belong to a finite number of clusters $c \in \mathcal{C}$. For each cluster c, we denote $\mathcal{D}_c$ as the set of feasible daily power consumption schedules that satisfy the energy requirements of the corresponding appliances. Any power consumption schedule, $[d_c(t)]_{t=1,\dots,T} = \mathbf{D}_c \in \mathcal{D}_c$, would satisfy the daily power needs of an appliance in cluster c. For example, consider a cluster that represents plug-in electric vehicles (EVs) that require E_c kWh in the time interval $[t_1, t_2]$ with a maximum charging rate of ρ_c kW. Accordingly, the set $\mathcal{D}_c$ of daily feasible power consumption schedules is given by:

$$\mathcal{D}_c = \Big\{ \mathbf{D}_c | \sum_{t=t_1}^{t_2} d_c(t) = E_c;\ 0 \leq d_c(t) \leq \rho_c \Big\}. \tag{4.6}$$

Another specific cluster example is that of electric appliances that are uninterruptible but can perform load shifting (e.g., a dishwasher cannot be interrupted once it is turned on but the start time of the cycle can be shifted). Let $\Pi_c(\cdot)$ denote the load profile of uninterruptible cluster c appliances once they are turned on. For example, $\Pi_c(\cdot)$ could be a rectangular pulse function that outputs the rated power of the appliance, ρ_c for the duration of the appliance's cycle and 0 otherwise. To relay their load flexibility, cluster c users can specify a time interval $[t_{c,1}, t_{c,2}]$ within which the appliance cycle must start (e.g., a user wants the dishwasher to be finished before dinner). Thus the home energy management system can calculate the best values for the time shift, denoted by t_c, as long as it lies within the interval $[t_{c,1}, t_{c,2}]$. The set $\mathcal{D}_c$ of daily feasible power consumption

schedules for appliances in this cluster is given by:

$$\mathcal{D}_c = \Big\{ \mathbf{D}_c | d_c(t) = \Pi_c(t - t_c);\ t_c \in [t_{c,1}, t_{c,2}] \Big\}. \tag{4.7}$$

For discussion on characterizing the sets for other flexible appliances, including interruptible (Section III.B in [127], non-interruptible (Section III.D in [127]), and thermostatically controlled loads (Section III.C in [127]), we refer the reader to reference [127].

By adopting this model, the total power consumption flexibility of a population of customers can be characterized as a function of how many appliances belong to each cluster within the population. Let us denote a_c as the number of appliances in cluster c (note that this will vary on a day by day basis as described in the next section). With this notation, we can write the set of feasible daily power consumption profiles for the population, $\mathcal{D}$:

$$\mathcal{D} = \sum_{c \in \mathcal{C}} a_c \mathcal{D}_c, \tag{4.8}$$

where the summation and scalar multiplication operations are defined in the sense of Minkowski addition[3].

We would like to note that choosing the number of clusters in the model is a control knob that can be tuned by the aggregator as shown in [128]. Using a higher number of potential appliance clusters will increase the accuracy of the load model (i.e., reduce the quantization error in the reproduction of the individual load profiles) and yield better performance in the daily optimization once the true parameters have been sufficiently learned by the aggregator. However, increasing the number of load clusters increases the size of the problem space and increases the randomness in the customers' daily loads

[3]For two sets A and B defined on a finite dimensional Euclidean space, the Minkowski sum is defined as $A + B = \{\mathbf{a} + \mathbf{b} \mid \mathbf{a} \in A, \mathbf{b} \in B\}$.

thus slowing down the learning rate of the algorithm. The number of clusters will vary depending on the system being analyzed as well as the aggregator's preferences.

Price Response Model

In this section, we discuss how the total population responds to dynamic electricity prices given the load flexibility model in (4.8) and how clustering is used to reduce the dimensionality of the problem. There are two main ways dynamic pricing affects the power consumption: 1) *Automated per cluster response*: Within each load cluster c (i.e., given pre-specified preferences such as EV charging deadlines or AC temperature set points), we assume that the customer chooses the power consumption profile $\mathbf{D}_c \in \mathcal{D}_c$ that minimizes their electricity cost dependent on the daily broadcasted price $\mathbf{p}_\tau$. For appliances in cluster c on day τ, we assume all will choose the same minimum cost power consumption profile:

$$\widetilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau) = \underset{\mathbf{D}_c \in \mathcal{D}_c}{\arg\min} \sum_{t=1}^{T} p(t) d_c(t). \tag{4.9}$$

We assume that each appliance will always choose the cost minimizing power consumption profile out of the available profile set to combat the fact that the available profile sets $\mathcal{D}_c$ for each cluster can be infinitely large. Thus, we have effectively reduced the dimensionality of the problem as we know *a priori* how each cluster will respond to each price signal (i.e., each cluster will always select its cost minimizing profile). Due to the automated nature of home energy management systems, each cluster selecting its cost minimizing profile is a reasonable assumption once the customers have defined their flexibility preferences, e.g., the desired charge amounts and deadlines for EVs [129, 130]. 2) *Preference Adjustment*: We also consider the fact that customers may respond to price signals by adjusting their preferences. Consider the following example: two customers

(Customer-A and Customer-B) live in the same neighborhood but have different sensitivities to electricity prices. If electricity prices are high on a hot summer day, Customer-A might shutdown their air conditioner to avoid a large electricity bill; however, Customer-B prioritizes comfort over cost-savings, and leaves their air conditioner on, no matter the cost. As shown in the previous example, the number of appliances in each cluster, i.e., a_c in (4.8), also depends on the daily posted price vector $\mathbf{p}_\tau$, and are now denoted as $a_c(\mathbf{p}_\tau)$.

Combining the *automated per cluster response* and *preference adjustment*, we can define the population's load on day τ in response to the posted price $\mathbf{p}_\tau$ as follows:

$$\mathbf{D}^\star_\tau(\mathbf{p}_\tau) = \sum_{c\in\mathcal{C}} a_c(\mathbf{p}_\tau)\widetilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau). \tag{4.10}$$

As stated before, if the aggregator has full knowledge of the customers' price responses, which reduces to having full knowledge of the preference adjustments $a_c(\mathbf{p}_\tau)$, then the aggregator can pick the daily price vector $\mathbf{p}^\star_\tau$ in order to shape the population's power consumption according to (4.1). However, as we cannot assume this, we model the $a_c(\mathbf{p}_\tau)$'s as random variables with parameterized distributions, ϕ_c, based on the posted price signal $\mathbf{p}_\tau$ and an unknown but constant parameter vector $\boldsymbol{\theta}^\star$. Here, $\boldsymbol{\theta}^\star$ represents the *true model* for the customers' sensitivity to the price signals. This allows for the complex response of the customer population to be represented as a single vector, thus reducing the dimensionality of the problem. We note that while $a_c(\mathbf{p}_\tau)$ may only take integer values in reality, we believe it is justified to relax this integrality constraint and allow it to take continuous values with large enough appliance population size. With this in mind, we would like to highlight three properties of the price response model:

1. The preference adjustment models $a_c(\mathbf{p}_\tau)$ are stochastic and their distributions ϕ_c are parameterized by $\mathbf{p}_\tau$ and $\boldsymbol{\theta}^\star$. This is due to exogenous factors outside of the

aggregator's scope that influence customers' power consumption profiles resulting in a level of stochasticity in the responses to prices (i.e., customers will not respond to prices in the same fashion each day).

2. The probability distributions of $a_c(\mathbf{p}_\tau)$ (i.e., ϕ_c) are unknown to the aggregator, i.e., the aggregator does not know the true parameter $\boldsymbol{\theta}^\star$ of the stochastic model.

3. The realizations of $a_c(\mathbf{p}_\tau)$ are not directly observable by the aggregator. The aggregator can only monitor the population's total consumption profile $\mathbf{D}_\tau$ and cannot observe the decomposed response of each cluster $a_c(\mathbf{p}_\tau)\widetilde{\mathbf{D}}_{c,\tau}(\mathbf{p}_\tau)$ independently.

Because we have introduced stochasticity to customers' price response models, we appropriately alter the aggregator's optimization problem for selecting the price signal on day τ to account for the distributions ϕ_c:

$$\mathbf{p}_\tau^\star = \underset{\mathbf{p}_\tau \in \mathcal{P}}{\arg\min}\, \mathbb{E}_{\{\phi_c\}_{c\in\mathcal{C}}}\big[f\big(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau\big)\big] \tag{4.11}$$

$$\text{s.t.} \quad \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}\big[g_j\big(\mathbf{D}_\tau(\mathbf{p}_\tau)\big) \leq 0\big] \geq 1-\mu, \;\; \forall j \tag{4.12}$$

where μ is the aggregator's desired reliability metric for the distribution system constraints. In (4.11), the aggregator now considers minimizing an expected cost and is subject to probabilistic reliability constraints in (4.12) that depend on the distributions ϕ_c of the preference adjustment models $a_c(\mathbf{p}_\tau)$.

We note that the formulated chance constraints are enforced with respect to uncertainty in the clusters' price sensitivity parameters, not to the exogenous context vector $\mathbf{V}_\tau$. In this work, we assume the daily exogenous vector is fully known each day and does not add uncertainty to the problem. However, uncertainties in the exogenous vector are important to real-world systems such as the power grid and can be accommodated by our approach by adding external noise to these vectors in the same fashion as noise being

added to the population's load. This, of course, would further slow down the learning rate of the algorithm due to the added noise reducing the effectiveness of each posterior update.

Clearly, the aggregator needs to learn the underlying parameters of the stochastic models ϕ_c of how customers respond to price signals in order to select price signals for load shaping initiatives (i.e., the aggregator needs to learn $\boldsymbol{\theta}^\star$). Our proposed learning approach and pricing strategy for an electricity aggregator is detailed in the next section.

4.3 Real-Time Pricing via Multi-Armed Bandit

Multi-Armed Bandit Overview

We utilize the multi-armed bandit (MAB) framework to model the iterative decision making procedure of an aggregator implementing a daily load shaping program [131, 132, 133]. The MAB problem can be described as a decision making problem where an agent has a set of available actions but can only take one action per round. After an action is taken, the agent experiences a cost that is dependent on the action taken. The agent can only learn about the distribution of costs from each action by experimenting. Throughout this iterative procedure, the agent faces the core dilemma: should the agent *exploit* actions that have yielded small costs, or *explore* actions that have not been tested thoroughly? The goal in a MAB problem is to develop a strategy for selecting actions that balance this trade-off and minimize the cumulative cost over a given time span. More thorough explanation and background of the MAB problem can be found in [134].

For the electricity pricing problem, the MAB framework exemplifies the *exploration-exploitation* trade-off dilemma faced by an aggregator each day. Namely, should the aggregator choose to broadcast untested prices (i.e., *explore*) to learn more information about the customers? Or should the aggregator choose to broadcast well-performing

prices (i.e., *exploit*) to manipulate the daily electricity demand?

To evaluate the performance of an algorithm that aims to tackle the exploration-exploitation trade-off, one commonly examines the algorithm's *regret*. Formally, regret is defined as the cumulative difference in cost incurred over $\mathcal{T}$ days between a clairvoyant algorithm (i.e., the optimal strategy that is aware of the customers' price responses) and any proposed algorithm that does not know the customers' price responses:

$$R_{\mathcal{T}} = \sum_{\tau=1}^{\mathcal{T}} f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau) - f(\mathbf{D}_\tau(\mathbf{p}^\star), \mathbf{V}_\tau). \tag{4.13}$$

Instead of considering the cumulative difference in objective function value, an alternative metric for regret is to count the number of times that suboptimal price signals are selected over the $\mathcal{T}$ days. For this, we introduce the following notation: let $\mathbf{p}^{\mathbf{V}_\tau,\star}$ denote the optimal price signal for the true model of the population's price response $\boldsymbol{\theta}^\star$ when the daily exogenous parameter $\mathbf{V}_\tau$ is observed on day τ. Any price signal $\mathbf{p}_\tau \neq \mathbf{p}^{\mathbf{V}_\tau,\star}$ is considered a suboptimal price. Moreover, we denote $N_\tau(\mathbf{p}, \mathbf{V})$ as the number of times up to day τ that the algorithm simultaneously observes the exogenous parameter $\mathbf{V}$ and selects the price signal $\mathbf{p}$. As such, the total number of times that suboptimal price signals are selected over $\mathcal{T}$ days is:

$$\sum_{\mathbf{V}\in\mathcal{V}} \sum_{\mathbf{p}\in\{\mathcal{P}\setminus\mathbf{p}^{\mathbf{V},\star}\}} N_{\mathcal{T}}(\mathbf{p}, \mathbf{V}) = \sum_{\tau=1}^{\mathcal{T}} \mathbb{1}^{\{\mathbf{p}_\tau \neq \mathbf{p}^{\mathbf{V}_\tau,\star}\}}, \tag{4.14}$$

where $\mathbb{1}^{\{\cdot\}}$ is the indicator function that is set equal to one if the criteria is met and zero otherwise. Subsequently, in an iterative decision making problem such as this, the question arises: *how can an aggregator learn to price electricity with bounded regret, and what are the regret bounds we can provide for a proposed algorithm given dynamically changing grid conditions and reliability constraints?* In the following sections, we present

a modified Thompson sampling heuristic for the electricity pricing problem to simultaneously learn the true model $\boldsymbol{\theta}^\star$ for the population, select the daily price signals, ensure grid reliability, and provide a regret guarantee.

Thompson Sampling

Thompson sampling (TS) is a well-known MAB heuristic for choosing actions in an iterative decision making problem with the *exploration-exploitation* dilemma [134, 135, 136]. Two other well-studied frameworks, greedy algorithms and upper-confidence bound (UCB) algorithms, have shown promise in this problem area. However, greedy algorithms are inferior to Thompson sampling in regret performance and UCB algorithms are restricted to simpler linear optimizations [137, 138, 139], whereas Thompson sampling can readily handle more general objective functions such as those adopted in our work [140]. Additionally, a novel aspect of our work is that we have shown how to modify the Thompson sampling heuristic to account for reliability constraints with a theoretical guarantee (Proposition 1). There are no other bandit optimization approaches known to be able to handle general objective functions with safety constraints. Relevant works here include the analysis of the performance of the UCB algorithm in the linear MAB setting with linear safety-constraints [141], and well as linear TS with linear constraints [142]. In the latter work, it is shown that in the linear case, the presence of linear constraints do not negatively affect the regret performance of TS, which is remarkable and could be a preliminary justification as to why TS performs well in our work in the presence of general (non-linear) cost and constraint functions.

Simply put, the integral characteristic of Thompson sampling is that the algorithm's knowledge on day τ of the unknown parameter $\boldsymbol{\theta}^\star$ is represented by the prior distribution $\pi_{\tau-1}$. Each day the algorithm samples $\widetilde{\boldsymbol{\theta}}_\tau$ from the prior distribution, and selects an action assuming that the sampled parameter is the true parameter. The algorithm then makes

an observation dependent on the chosen action and the hidden parameter and performs a Bayesian update on the parameter's distribution π_τ based on the new observation. Because TS samples parameters from the prior distribution, the algorithm has a chance to explore (i.e., draw new parameters) and can exploit (i.e., draw parameters that are likely to be the true parameter) through out the run of the algorithm.

Constrained Thompson Sampling

In this section, we present the MAB heuristic titled Con-TS-RTP adopted to the electricity pricing problem. Con-TS-RTP is a modified Thompson sampling algorithm where the daily optimization problem is subject to constraints (standard TS algorithms do not have constraints)[143].

When initializing π_0, the initial distribution on the customers' unknown parameter can be selected by the aggregator. If the aggregator has access to prior information regarding the true parameter, then they could initialize the prior as a distribution of their choice. However, if the aggregator has no prior knowledge, a uniform distribution among all available parameters may be used to model the lack of knowledge of the aggregator.

Each day, the algorithm observes the daily target profile $\mathbf{V}_\tau$, draws a parameter $\widetilde{\boldsymbol{\theta}}_\tau$ from the prior distribution, broadcasts a price signal to the customers, observes the load profile of the population in response to the broadcasted price, and then performs a Bayesian update on the parameter's distribution π_τ based on the new observation. We note that there are no restrictions on the class of optimization problem to be solved each day; however, in order for our regret guarantee to hold, the aggregator *must* be able to find the globally optimal solution and can use any desired solution method to do so. In our experimental examples, we assume that θ's and $\mathbf{p}_\tau$'s are chosen from discrete sets in order to be able to guarantee that an enumeration method could solve for the globally

optimal price signals each day in spite of non-convexities that arise.

The observation on day τ is denoted as $\mathbf{Y}_\tau = \mathbf{D}^\star_\tau(\mathbf{p}_\tau)$ and we assume that each $\mathbf{Y}_\tau$ comes from the observation space $\mathcal{Y}$ that is known a priori. When performing the Bayesian update, the algorithm makes use of the following *likelihood function*: $\ell(\mathbf{Y}_\tau; \mathbf{p}, \boldsymbol{\theta}) = \mathbb{P}_\theta(\mathbf{D}^\star_\tau(\mathbf{p}_\tau) = \mathbf{Y}_\tau | \mathbf{p}_\tau = \mathbf{p})$. This function calculates the likelihood of observing a specific load profile when broadcasting price $\mathbf{p}$ and the true parameter is $\boldsymbol{\theta}$. The pseudocode for Con-TS-RTP applied to the constrained electricity pricing problem is presented in Algorithm 4.

Discussion on Regret Performance of Con-TS-RTP

The regret analysis of Con-TS-RTP is inspired by the results in [140] for TS with nonlinear cost functions. The authors in [144] extended the regret results from [140] by analyzing the effects of an objective function that is dependent on exogenous parameters such as $\mathbf{V}_\tau$. The analysis in the aforementioned papers provides bounds on the total number of times that suboptimal price signals selected by the algorithm over $\mathcal{T}$ days as specified in equation (4.14). The regret guarantee we provide in this work extends the result further, allowing for constraints in the daily optimization that are dependent on the sampled $\widetilde{\boldsymbol{\theta}}_\tau$. As such, our regret guarantee applies to the Con-TS-RTP algorithm with constraints as formulated in *Constraint Set A* in Algorithm 4. We refer the reader to the online appendix [3] as well as [144] and [140] for further discussion on the derivation of Theorem 4.3.1.

Assumption 4.3.1 *(Finitely many price signals, observations).* $|\mathcal{P}|, |\mathcal{Y}| < \infty$.

Assumption 4.3.2 *(Finite Prior, "Grain of truth") The prior distribution π is supported over finitely many particles: $|\Theta| < \infty$. The true parameter exists within the parameter space: $\boldsymbol{\theta}^\star \in \Theta$. The initial distribution π_0 has non-zero mass on the true parameter $\boldsymbol{\theta}^\star$*

Algorithm 4 CON-TS-RTP

Input: Parameter set Θ; Price set $\mathcal{P}$; Observation set $\mathcal{Y}$; Voltage constraints u^{min}, u^{max}; Power flow constraint S^{max}, Reliability metrics μ, ν

Initialize π_0 based on aggregator's available prior knowledge of customer sensitivity.

1: **for** Day index $\tau = 1...\mathcal{T}$ **do**

2: Sample the daily hidden parameter $\widetilde{\boldsymbol{\theta}}_\tau \in \Theta$ from the aggregator's prior distribution $\pi_{\tau-1}$.

3: Observe the daily exogenous parameter $\mathbf{V}_\tau$.

4: Broadcast the daily price signal:

$$\hat{\mathbf{p}}_\tau = \arg\min_{\mathcal{P}} \mathbb{E}_{\{\phi_c\}_{c\in\mathcal{C}}}\left[f(\mathbf{D}_\tau(\mathbf{p}_\tau), \mathbf{V}_\tau)|\boldsymbol{\theta} = \widetilde{\boldsymbol{\theta}}_\tau\right]$$

Subject to:

Constraint Set A:

$$\begin{cases} \textbf{\textit{A.1:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \geq u^{min}|\boldsymbol{\theta} = \widetilde{\boldsymbol{\theta}}_\tau] \geq 1-\mu, \quad \forall t \\ \textbf{\textit{A.2:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \leq u^{max}|\boldsymbol{\theta} = \widetilde{\boldsymbol{\theta}}_\tau] \geq 1-\mu, \quad \forall t \\ \textbf{\textit{A.3:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[f_\tau(t) \leq S^{max}|\boldsymbol{\theta} = \widetilde{\boldsymbol{\theta}}_\tau] \geq 1-\mu, \quad \forall t \end{cases}$$

Constraint Set B:

$$\begin{cases} \textbf{\textit{B.1:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \geq u^{min}|\boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \\ \textbf{\textit{B.2:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \leq u^{max}|\boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \\ \textbf{\textit{B.3:}}\ \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[f_\tau(t) \leq S^{max}|\boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \end{cases}$$

5: Observe the population's load response to price $\mathbf{p}_\tau$: $\mathbf{Y}_\tau = \mathbf{D}^\star_\tau(\mathbf{p}_\tau)$.

6: Update the aggregator's knowledge of the true parameter in the posterior:

$$\forall S \subseteq \Theta : \pi_\tau(S) = \frac{\int_S \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \boldsymbol{\theta})\pi_{\tau-1}(d\boldsymbol{\theta})}{\int_\Theta \ell(\mathbf{Y}_\tau; \hat{\mathbf{p}}_\tau, \boldsymbol{\theta})\pi_{\tau-1}(d\boldsymbol{\theta})}$$

7: **end for**

(i.e., $\mathbb{P}_{\pi_0}[\boldsymbol{\theta}^\star] > 0$*).*

Assumption 4.3.3 *(Unique optimal price signal). There is a unique optimal price signal* $\mathbf{p}^{\mathbf{V},\star}$ *for each exogenous parameter* $\mathbf{V} \in \mathcal{V}$*.*

Theorem 4.3.1 *Under assumptions 4.3.1-4.3.3 and Constraint Set A in Algorithm 4, for* $\delta, \epsilon \in (0,1)$*, there exists* $\mathcal{T}^\star \geq 0$ *s.t. for all* $\mathcal{T} \geq \mathcal{T}^\star$*, with probability* $1-\delta$*:*

$$\sum_{\mathbf{V}\in\mathcal{V}} \sum_{\mathbf{p}\in\{\mathcal{P}\setminus\mathbf{p}^{\mathbf{V},\star}\}} N_{\mathcal{T}}(\mathbf{p},\mathbf{V}) \leq B + C(\log \mathcal{T}), \tag{4.15}$$

where $B \equiv B(\delta,\epsilon,\mathcal{P},\mathcal{Y},\Theta)$ *is a problem-dependent constant that does not depend on* $\mathcal{T}$*, and* $C(\log\mathcal{T})$ *depends on* $\mathcal{T}$*, the sequence of selected price signals, and the Kullback-Leibler divergence properties of the bandit problem (i.e., the marginal Kullback-Leibler divergences of the observation distributions* $KL\big[\ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta}^\star),\ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta})\big]$ *(The complete description of the* $C(\log\mathcal{T})$ *term is left to the online appendix).*

Proof. The proof is in the online appendix [3].

Con-TS-RTP with Improved Reliability Constraints

In order for the aggregator to ensure safe operation of the distribution grid while running the Con-TS-RTP algorithm, the reliability constraints need to hold for the true price response model $\boldsymbol{\theta}^\star$ each day. However, with the constraints formulated as in Algorithm 4's *Constraint Set A*, the distribution system constraints are only enforced for the sampled $\widetilde{\boldsymbol{\theta}}_\tau$ and not necessarily the true parameter $\boldsymbol{\theta}^\star$. This entails that the distributions $\{\phi_c\}_{c\in\mathcal{C}}$ are parameterized by the sampled $\widetilde{\boldsymbol{\theta}}_\tau$; therefore, they are inaccurate if any parameter $\widetilde{\boldsymbol{\theta}}_\tau \neq \boldsymbol{\theta}^\star$ is sampled. This could potentially lead to many constraint violations

throughout the run of the algorithm resulting in inadequate service for the customers and grid failures.

Due to the importance of reliable operation of the distribution system, we present a modification to the Con-TS-RTP algorithm (i.e., replacing *Constraint Set A* with *Constraint Set B* in Algorithm 4) to increase the reliability of the selected prices and resulting load profiles with respect to the grid constraints. Specifically, we propose alternate constraints that depend on the algorithm's current knowledge of the true parameter, instead of the sampled parameter. In other words, instead of depending on $\widetilde{\boldsymbol{\theta}}_\tau$, the proposed alternate constraints depend on the prior distributions $\pi_{\tau-1}$ as follows:

$$\mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \geq u^{min} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \tag{4.16}$$

$$\mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[u_\tau(t) \leq u^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \tag{4.17}$$

$$\mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}[f_\tau(t) \leq S^{max} | \boldsymbol{\theta} \sim \pi_{\tau-1}] \geq 1-\nu, \quad \forall t \tag{4.18}$$

where ν is a small constant (detailed in Proposition 4.3.1). When considering constraints (4.16)-(4.18) in the Con-TS-RTP algorithm, the algorithm will select more conservative price signals each day that can guarantee the distribution system's constraints are met with high probability by using the information in the updated prior distributions. Before analyzing the modified algorithm's reliability, we make the following assumption:

Assumption 4.3.4 *There exists $\xi^\star > 0$ and $\lambda \geq 0$, such that for all $\boldsymbol{\theta} \neq \boldsymbol{\theta}^\star$, $KL\big[\ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta}^\star), \ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta})\big] \geq \xi^\star$, where*

$$\xi^\star_{\boldsymbol{\theta},\mathbf{p}} = \max_{x\in\mathbb{Z}^{>0}} \left\{ \frac{-\lambda}{x} - \frac{4}{\sqrt{x}} \sqrt{\frac{\log|\mathcal{Y}||\mathcal{P}|}{\delta} + \frac{\log x}{2}} \sum_{\mathbf{Y}\in\mathcal{Y}} \left| \log \frac{\ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta}^\star)}{\ell(\mathbf{Y};\mathbf{p},\boldsymbol{\theta})} \right| \right\} \tag{4.19}$$

and

$$\xi^\star = \max_{\boldsymbol{\theta}\in\Theta,\mathbf{p}\in\mathcal{P}} \xi^\star_{\boldsymbol{\theta},\mathbf{p}}.$$

Assumption 4.3.4 ensures that as the aggregator performs the steps in Algorithm 3, the algorithm's Bayesian updates of the prior distribution π_τ will likely never decrease the mass of the true parameter $\boldsymbol{\theta}^\star$ below a certain threshold. Specifically, with Assumption 4.3.4, it can be shown (as in [140]) that with probability $1-\delta\sqrt{2}$ the following holds for all $\tau \geq 1$:

$$\pi_\tau(\boldsymbol{\theta}^\star) \geq \pi_0(\boldsymbol{\theta}^\star)e^{-\lambda|\mathcal{P}|}, \tag{4.20}$$

where $\lambda \geq 0$ is a chosen parameter (from Assumption 4.3.4) that dictates the minimum reachable mass of the true parameter via Bayesian updating. With the modified constraints (4.16)-(4.18) and the minimum mass of the true parameter specified in (4.20), the reliability of Con-TS-RTP can be characterized as follows:

Proposition 4.3.1 *Under assumptions 4.3.1-4.3.4, with ν in equations* (4.16)-(4.18) *chosen such that $\nu \leq \mu\pi_0(\boldsymbol{\theta}^\star)e^{-\lambda|\mathcal{P}|}$, with probability $1-\delta\sqrt{2}$, the Con-TS-RTP algorithm with Constraint Set B will uphold the probabilistic distribution system constraints as formulated in* (4.12) *for each day $\tau = 1, \ldots, \mathcal{T}$.*

Proof. The proof is in Appendix A.5.

Remark: The novelty of Con-TS-RTP is that we can ensure with high probability an unsafe price signal is never selected. We can tune the safety parameter to determine what level of risk is acceptable to the aggregator. We note that the selection of an unsafe price signal has no effect on the learning capability of the algorithm. The Con-TS-RTP algorithm will learn regardless of safe/unsafe price signals. The algorithm will

never crash/stop prematurely due to the selection of an unsafe price signal; however, the local distribution grid might surpass safety limits on transformers or line flow limits due to an unsafe price selection on select very limited days, at which points protective measures (e.g., relays) should be used to ensure physical grid safety. We note that this is natural for any learning algorithm dealing with stochastic conditions and unknown system parameters. Contingencies can never be avoided 100%, similar to other grid operation paradigms that deal with uncertain conditions (e.g., wholesale market dispatch with renewables or possible transmission system contingencies). They could only be avoided with a certain high probability when making dispatch decisions. However, it is understood that other protective measures should always be put in place to avoid physical system damage in case of contingencies.

4.4 Experimental Evaluation

Test Setup: Radial Distribution System

In this section we describe the power distribution system and the corresponding network parameters for the test case. We consider an actual radial distribution system from the ComEd service territory in Illinois, USA (adopted from [145] and shown in Fig. 4.1) represented by the undirected graph $\mathcal{G}$ which includes a set of nodes (vertices) $\mathcal{N}$ and a set of power lines (edges) $\mathcal{L}$. In this work, we consider each node as one population with its own daily load profile; however, each node could be an aggregation of smaller entities downstream of the local distribution connection point. The undirected graph is organized as a tree, with the root node representing the distribution system's substation where it is connected to the regional transmission system. We denote N as the total number of nodes in the network excluding the root node. The nodes are indexed as $i = 0, \ldots, N$, and the node corresponding to $i = 0$ (i.e., the root node) is the substation. The power

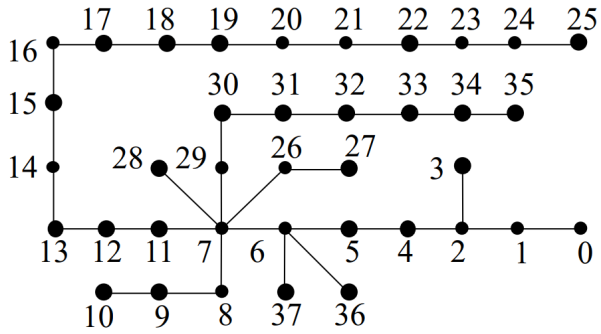

Figure 4.1: Radial distribution system.

lines are indexed by $i = 1, \ldots, N$ where the i-th line is directly upstream of node i (i.e., line i feeds directly to node i). In the following, we denote the parent vertex of node i as $\mathcal{A}_i$ and the set of children vertices of node i as $\mathcal{K}_i$.

Furthermore, we assume the aggregator has access to measurement data at each node's local connection point. Specifically, the aggregator measures the active power demands at each node i at time t on day τ denoted as $d^P_{i,\tau}(t)$. In order to ensure the delivered power is suitable for the electricity customers, the aggregator also monitors node i's local voltage at time t on day τ denoted as $v_{i,\tau}(t)$. In the following, we denote the active power daily load profile of node i on day τ as $\mathbf{D}^P_{i,\tau} = [d^P_{i,\tau}(t)]_{t=1,\ldots,T}$. Additionally, the aggregator records the active power flows $f^P_{i,\tau}(t)$ on each line $i \in \mathcal{L}$. We note that reactive power should also be monitored in distribution systems, even though it is generally not priced and customers do not consider it in determining their optimal load response to prices. As such, we use the superscript Q for the reactive power at a node, $d^Q_{i,\tau}(t)$, and for reactive power flow on a line, $f^Q_{i,\tau}(t)$. Each line in the distribution system has its own internal resistance denoted as R_i, reactance denoted as X_i, and power limit denoted as S^{max}_i. The parameters for the distribution system are listed in Table 4.1.

Line	R ($10^{-3}\Omega$)	X ($10^{-3}\Omega$)	S^{max} (KVA)	Line	R ($10^{-3}\Omega$)	X ($10^{-3}\Omega$)	S^{max} (KVA)
1	24.2	48.2	54	20	129.5	30.9	10.8
2	227.3	743.5	84	21	15.1	5.4	14.4
3	76.3	18.2	10.8	22	50.8	12.1	10.8
4	43.6	142.7	84	23	69.1	16.5	10.8
5	25.8	84.4	84	24	31.6	11.2	14.4
6	10.5	10.7	40.2	25	96.3	23	10.8
7	23.2	23.6	40.2	26	110.7	112.6	40.2
8	75.1	26.7	14.4	27	2.1	0.7	14.4
9	114.4	27.3	10.8	28	242.1	86.2	14.4
10	110.8.3	67.7	14.4	29	27.3	27.8	40.2
11	63.7	22.7	14.4	30	174.6	62.1	16.2
12	278.7	99.2	14.4	31	43	15.3	10.8
13	254.2	10.8.5	14.4	32	207.8	74	10.8
14	21.8	5.2	10.8	33	109.4	38.9	14.4
15	57.3	20.4	14.4	34	50.5	18	14.4
16	126.7	45.1	14.4	35	165.2	58.8	14.4
17	48.6	11.6	10.8	36	49.5	17.6	14.4
18	95.1	22.7	10.8	37	5.8	2.1	14.4
19	137.3	32.8	10.8				

Table 4.1: Distribution system parameters.

Power Flow Model

In order to solve for the power flow and nodal voltages of the power distribution system, we make use of the *LinDistFlow* model[146], which is a linear approximation for the AC power flow model[4]. The *LinDistFlow* model has been extensively studied and verified to be competitive to the nonlinear AC flow model on many realistic feeder topologies including radial [148, 149, 150, 151]. The *LinDistFlow* model reduces computational

[4]The reader should note that the proposed learning approach is not limited to the *LinDistFlow* model. There are other power flow models that can be utilized such as [147].

complexity by making use of the following linear power flow and voltage equations:

$$d^P_{i,\tau}(t) + \sum_{j \in \mathcal{K}_i} f^P_{j,\tau}(t) = f^P_{\mathcal{A}_i,\tau}(t);\ \forall t, \tau, i, \tag{4.21}$$

$$d^Q_{i,\tau}(t) + \sum_{j \in \mathcal{K}_i} f^Q_{j,\tau}(t) = f^Q_{\mathcal{A}_i,\tau}(t);\ \forall t, \tau, i, \tag{4.22}$$

$$u_{\mathcal{A}_i,\tau}(t) - 2\big(f^P_{i,\tau}(t)R_i + f^Q_{i,\tau}(t)X_i\big) = u_{i,\tau}(t); \forall t, \tau, i \tag{4.23}$$

where (4.21) accounts for active power and (4.22) accounts for reactive power. In (4.23) we make use of the operator $u_{i,\tau}(t) = \big(v_{i,\tau}(t)\big)^2$ to provide a linear voltage drop relationship across the distribution system. For the scope of this work, we assume that the substation connection to the regional transmission system (node $i = 0$) is regulated and has a fixed voltage $v_{0,\tau}(t) = 12.5\text{kV}, \forall t, \tau$.

Distribution System Operational Constraints

The nodal voltages and line flows calculated in (4.21)-(4.23) should obey the following constraints for reliable operation:

$$u_{i,\tau}(t) \geq u_i^{min}, \qquad \forall t, \tau, i \in \mathcal{N}, \tag{4.24}$$

$$u_{i,\tau}(t) \leq u_i^{max}, \qquad \forall t, \tau, i \in \mathcal{N}, \tag{4.25}$$

$$f^P_{i,\tau}(t)^2 + f^Q_{i,\tau}(t)^2 \leq (S_i^{max})^2, \qquad \forall t, \tau, i \in \mathcal{L}, \tag{4.26}$$

where (4.24)-(4.25) are the nodal voltage constraints and (4.26) corresponds to the power constraints for each distribution line.

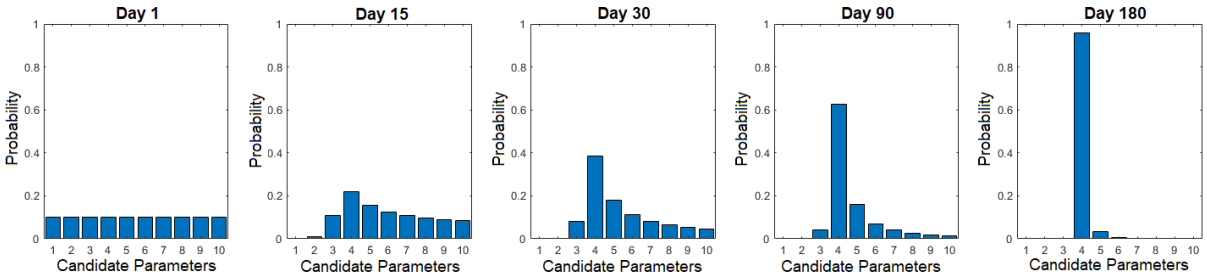

Figure 4.2: The 5 plots above portray the evolution of the aggregator's knowledge of the population's hidden parameter at node 10 throughout the learning procedure. The true parameter is parameter 4. From left to right: Day 1 (initialized to uniform distribution, i.e., no knowledge of the true parameter), Day 15 prior, Day 30 prior, Day 90 prior, and Day 180 prior. At day 180, the aggregator is about 95% certain that parameter 4 is the true parameter.

Load Model and Multi-armed Bandit Formulation

In this test case, we consider 6 time slots each day, each 4 hours long. We consider 10 unique target load profile vectors, with the daily target profile $\mathbf{V}_\tau$ for day τ drawn from a uniform distribution each morning. Each of the 10 target load profile vectors corresponds to a desired load curve to accommodate different levels of forecasted renewable generation. Furthermore, the aggregator transmits daily price signals $\mathbf{p}_{i,\tau}$ to each node within the system. The aggregator has a high and low price for each of the 6 time slots resulting in 2^6 possible daily price signals. Since the aggregator is shaping the electricity demand at each node within the distribution system, each node has its own cost $f\big(\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}), \mathbf{V}_\tau\big)$ that is dependent on the node's daily demand and the target profile. In this test case, we assume the cost function for each node is the squared deviation of the node's electricity demand from the target profile: $f\big(\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}), \mathbf{V}_\tau\big) = |\mathbf{D}_{i,\tau}(\mathbf{p}_{i,\tau}) - \mathbf{V}_\tau|^2$, thus equally penalizing over-usage and under-usage of electricity. We note that the units are KW^2 and if the aggregator had a converting function for the squared deviation (KW^2) to \$U.S.D., then we could calculate the monetary losses of the system. In our experimental examples, we make use of discrete sets for the available θ's and $\mathbf{p}_\tau$'s to guarantee that

an enumeration-based method could solve for the globally optimal price signals each day in spite of problem non-convexities.

We consider 20 unique load flexibility clusters in this test case. Each cluster's parameters represent the varying start/stop times, total energy demands, and power limitations common to EV loads in residential areas and are of the form presented in equation (4.6). We note that we generated the population's load price response directly using the same clustering model (i.e., the actual load response in the simulation is at the level of 20 clusters and can be well represented by the 20 clusters *plus additive noise*. For a discussion on the effects of poor clustering, we refer the reader to Section 4.4). Each node in the distribution system is comprised of these 20 load clusters with its own unique sensitivities $a_{i,c}(\mathbf{p}_\tau)$ for each cluster. Each sensitivity parameter is selected as $a_{i,c}(\mathbf{p}_{i,\tau}) \sim \mathcal{N}(\frac{\beta_c}{\boldsymbol{\theta}_i^\star \mathbf{p}_{i,\tau}}, \sigma^2)$ each day where β_c is a cluster specific constant known by the aggregator (we note that β_c represents *a priori* knowledge of customers' preferences and could come from behavioral studies; however, our framework does not require this and β_c can be completely omitted in cases where prior information is unavailable). Each node's price sensitivity, i.e., parameter to be learned, $\boldsymbol{\theta}_i^\star$, is a vector of length 6 and the set of possible parameters, Θ, contains 10 unique vectors. Unless noted, the reliability parameter chosen for the Con-TS-RTP algorithm is $\nu = 0.1$.

Note on reactive power: We note that reactive power is generally price insensitive; however, reactive power is present in a distribution system and affects the constraints of the system. Reactive power flows alter how the price sensitive loads are limited by the operational constraints of the system (i.e., active and reactive flows on lines affect the capacity available for the price responsive loads). Due to the lack of data as to how much reactive power is present in the distribution system due to our appliance clusters and otherwise, for our numerical examples, we omit the inclusion of reactive power to only view the appliance clusters' active load profiles within the distribution system. For

further discussion on this, we refer the reader to papers that fully capture the effects of reactive power in such problems such as [152] and [153] in which the authors showcase techniques to handle distribution systems with chance constraints.

In the following sections, without loss of generality, we assume that reactive power is not responsive to the pricing signals. We note that our proposed learning approach can accommodate reactive power flows (*LinDistFlow* can as well); however, our goal was to show proof of concept of our learning/pricing approach with active customer loads, thus reactive power flow will be examined in future work.

Results

We simulated the Con-TS-RTP algorithm for 365 days for an aggregator attempting to learn the sensitivities of the nodes in the system and shape their demands. In the following, we highlight the results of the simulation at node 10 of the radial distribution system. Figure 4.2 presents the evolution of the prior distribution for node 10's hidden parameter. On day 1, the prior was initialized to a uniform distribution among the candidate parameters, and by day 180 the weight on the true parameter exceeded 0.95.

Figure 4.3 presents the regret performance of Con-TS-RTP at node 10. As seen in Figure 4.3, the regret curve flattens after day 130 as the algorithm never chooses a suboptimal price signal after this day.

Figure 4.4 presents node 10's deviation from a specific daily target profile. On days 2, 3, 4, 53, and 365 the same target profile (i.e., $\mathbf{V}_2 = \mathbf{V}_3 = \mathbf{V}_4 = \mathbf{V}_{53} = \mathbf{V}_{365}$) was drawn and the aggregator selected different price signals to shape the node's demand. As seen in Fig. 4.4, the deviation from the target profile on day 365 is less than the deviation on the other days as the algorithm has learned the true parameter and selects the optimal price signal to shape the load.

In Figure 4.5, we present the distribution system constraint violations that were

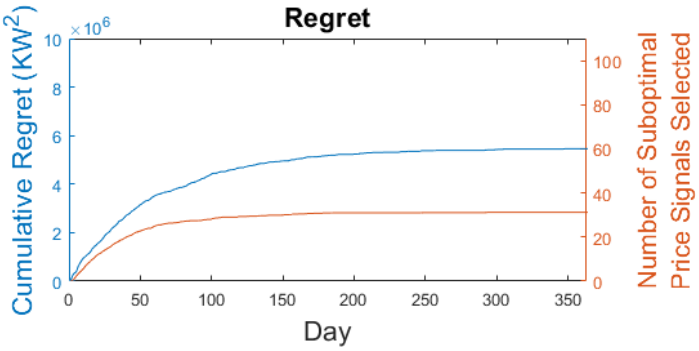

Figure 4.3: Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that the y-axis (left) units are KW^2 for the squared load deviation from the target profile.

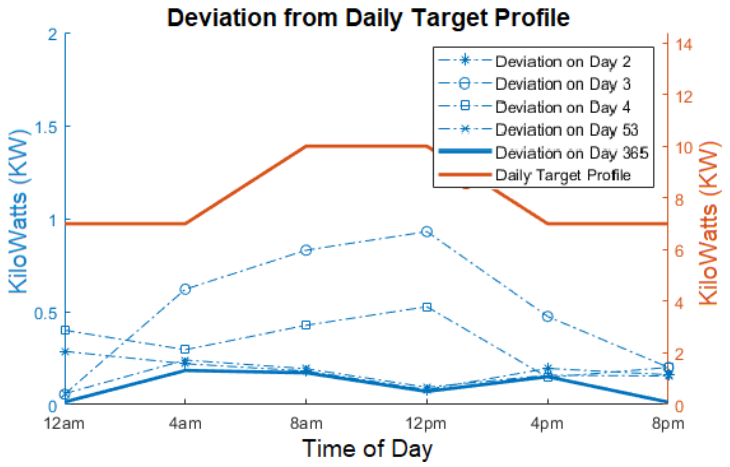

Figure 4.4: Deviation of node 10's demand from a specific daily target profile.

avoided by using Con-TS-RTP instead of an unconstrained TS algorithm. Clearly, in the early learning stages, the unconstrained TS algorithm does not have accurate knowledge of the hidden parameters and violates the distribution system constraints often. Con-TS-RTP is more conservative with its exploration of untested price signals and avoids the constraint violations made by the unconstrained TS algorithm. Last, we note that the simulation was implemented in Matlab and CVX on an i7 processor with 16gb of RAM. The 365 day simulations were run in less than 5 minutes.

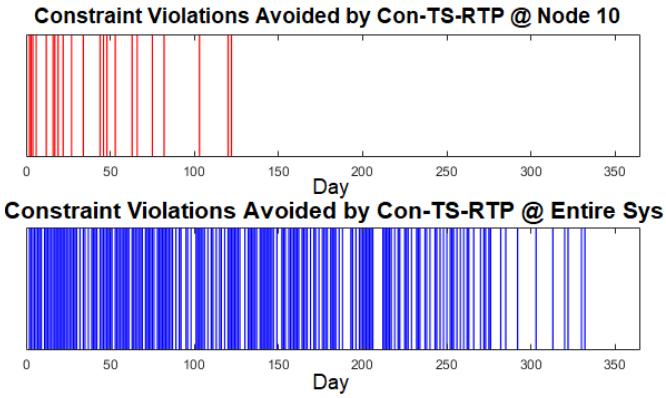

Figure 4.5: Top: Distribution system constraint violations at node 10 avoided by using Con-TS-RTP instead of an unconstrained TS. Bottom: Distribution system constraint violations across the entire system avoided by using Con-TS-RTP instead of an unconstrained TS.

Effects of Clustering

In this section, we portray the effects of selecting different numbers of clusters to represent a true load as well as the effects of selecting too few clusters on the performance of our Con-TS-RTP algorithm. First, in Figure 4.6, we perform a simple demonstration. We considered a population of 100 EVs with random charging requests and then constructed clusters to view the accuracy of the clustered load profiles versus the actual load profile. As shown in Figure 4.6, using 1, 5, or 10 clusters to represent the EV population results in load profiles quite different from the actual; however, with 20 clusters, the load profile begins to match the actual profile.

Furthermore, in Figure 4.7 we show the effects of reducing the number of clusters in the load model on the regret performance of our Con-TS-RTP algorithm. Specifically, we focus on the same setup as Section 4.4 with the exception that we have the Con-TS-RTP algorithm use a 10 cluster model instead of the 20 cluster model for the population to see the effects of an inaccurate cluster model. As shown in Figure 4.7, the regret curve for this case never flattens and the algorithm is never able to select the optimal price signal. This is because the algorithm's model of the load (i.e., the 10 clusters) is unable to

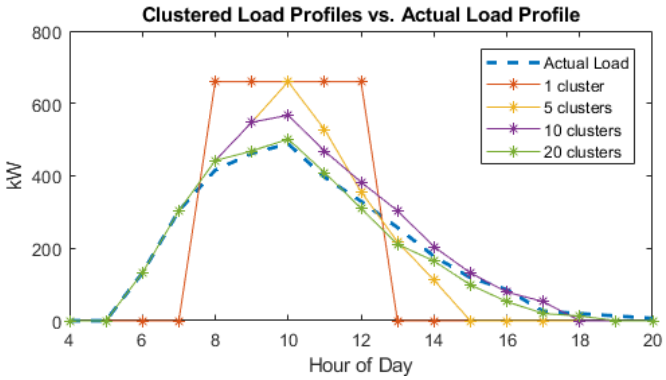

Figure 4.6: Effects of changing the number of clusters to model an actual load. Specifically, load profiles for 4 cluster models compared to the actual load profile for a population of 100 charging EVs.

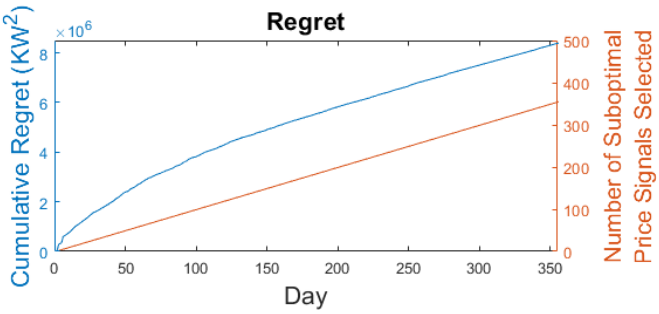

Figure 4.7: Effects of using too few clusters for the population's load model. We show the regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$ for a 10 cluster model instead of the 20 cluster model as previously shown in Fig. 4.3. Due to the inaccuracies of the 10 cluster model, the algorithm is never able to select the optimal price signals.

accurately model the population's response and causes the algorithm to select incorrect prices every day.

Evolving Price Sensitivity

In this section, we show an example of what happens when customers' sensitivities change over time and how a Bayesian learning approach can naturally adapt and account for these dynamic changes. Specifically, we simulated the same system as in Figure 4.3, but on day 250, we altered the true $\boldsymbol{\theta}_i^\star$ parameter. As seen in Figure 4.8, the regret curves

first flatten around day 125, then increase at day 250, and then flatten again near day 325. This shows that Con-TS-RTP was able to successfully learn the first and second true parameter without any modifications to the algorithm. The algorithm naturally shifts its belief about the true parameter as it observes outputs that do not match its current belief.

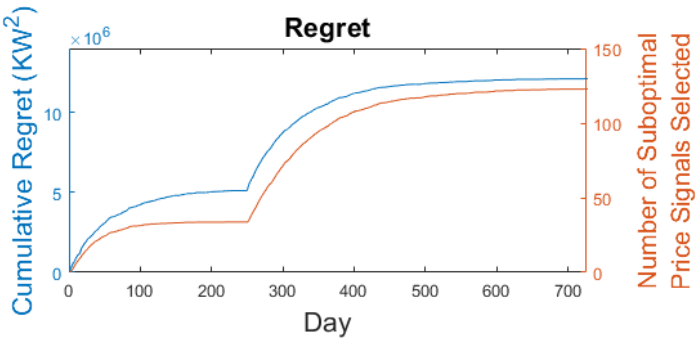

Figure 4.8: Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that on day 250, the hidden parameter was altered.

Non-repeating Target Profiles

In the previous case study, we assumed a low number of target profiles (10 profiles) to satisfy the assumptions we have made for our theoretical results. In this section, we demonstrate how extending the number of target profiles to 365 does not negatively affect the performance of the algorithm in practice. Furthermore, we ensure that once a target profile has been viewed by the aggregator, it is never drawn again. Thus, each day the aggregator is posting a price to shape the population's load to match a target profile that it has never seen before. As shown in Figure 4.9, enlarging the set of target profiles does not slow down the learning process. Note that in Figure 4.9 the regret flattens near trial 100 which matches the duration of the learning period seen in Figure 4.3 (i.e., in simulation, the aggregator is still able to learn the true parameter when the number of target profiles is increased from 10 to 365, resulting in similar regret curves).

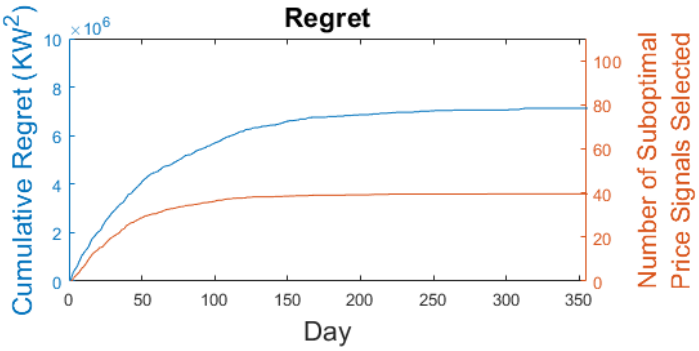

Figure 4.9: Regret performance of Con-TS-RTP at node 10 with $\nu = 0.1$. Note that on each day, the sampled $\mathbf{V}_\tau$ has never been seen by the aggregator.

Effects of Varying the System Reliability Metric

In this section, we discuss the effects of varying the system reliability parameters in the daily optimization's constraints (i.e., altering the value of ν for the system constraints formulated as in (4.16)-(4.18)). As described in Sections 4.2 and 4.3, the reliability metric dictates the aggregator's allowable probability of a constraint violation under its current belief distribution about the unknown parameter. Decreasing ν is restricting the algorithm to avoid violations and setting $\nu = 1$ is equivalent to solving the daily optimization without the constraints altogether. In Figure 4.10, we simulated the system with varying reliability parameters. Specifically, each curve shown is the average regret at node 10 over 20 independent simulations. As shown in Figure 4.10, the regret increases as the desired reliability increases (smaller ν). This is because the aggregator is forced to select more conservative prices during the learning procedure to ensure that the constraints are met with higher probability.

Comparison with Two-Stage Learning

In this section, we present a comparison of the Con-TS-RTP approach versus a 2-stage "learn" and then "optimize" algorithm, where the first stage consists of pure exploration and the second stage purely exploits the knowledge gained in the first stage. The sim-

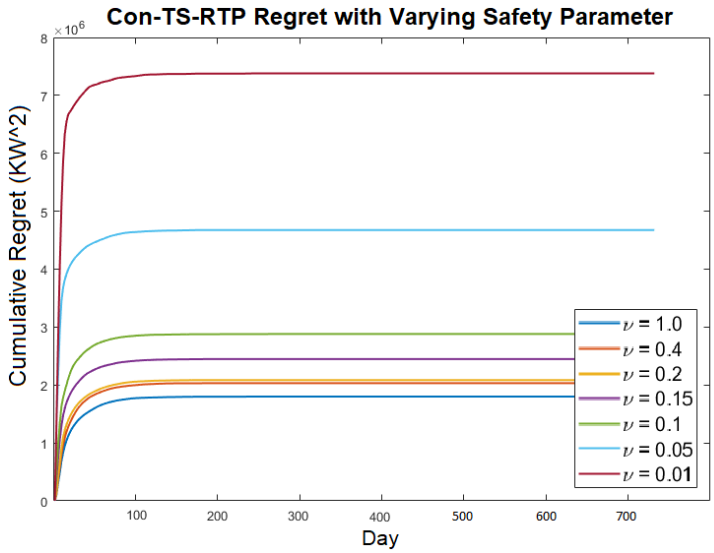

Figure 4.10: Regret curves for various system reliability metrics. Each curve is an average of 20 independent simulations.

ulation setup is the same as the setup used in Section 4.4. A description of the 2-stage algorithm used is as follows: The aggregator decides the duration of the learning stage *a priori*, (in Figure 4.11, we present regret curves for learning stages with durations of 5, 15, and 25 days) and during this learning stage, the aggregator randomly selects price signals from a predetermined safe set of prices (i.e., prices high enough such that constraints cannot be violated), observes the populations' responses, and performs posterior updates. Then, after the learning stage is complete, for the remainder of time the aggregator broadcasts the best price signals with respect to the knowledge of the unknown parameter at the end of the learning stage (the selected price signal will ensure safety but might be potentially suboptimal depending on the duration of the learning stage). Clearly, the two most significant shortcomings of the 2-stage approach are: 1) arbitrarily bad performance during the learning stage due to random price selection; and 2) difficulty selecting a sufficient duration of the learning stage. As seen in Figure 4.11, this 2-stage myopic algorithm results in linear regret in the 5 day and 15 day learning stage curves. Due to an insufficient number of posterior updates, the aggregator is forced to

post suboptimal price signals to ensure safety given its noisy knowledge of the unknown parameter after the learning stage is over. On the other hand, the 25 day learning stage is able to converge to the optimal price signals, but the performance during the learning stage causes fast growth of regret whereas Con-TS-RTP is able to avoid all of the aforementioned shortcomings.

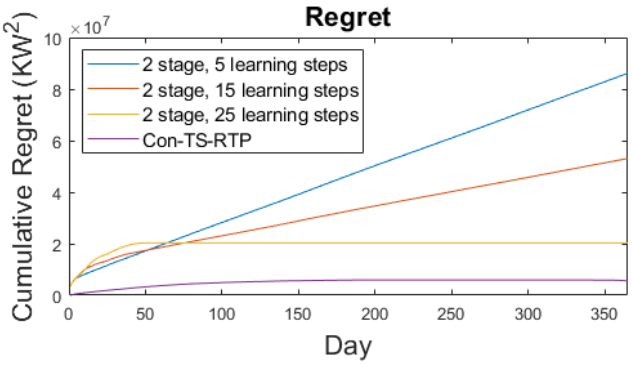

Figure 4.11: Regret performance of Con-TS-RTP and a 2-Stage algorithm at node 10 with $\nu = 0.1$. Note that the 5 day (blue) and 15 day (red) learning algorithms were unable to converge to the optimal price signals. The blue and red curves never flatten because their learning stages were too brief to adequately learn the customers' preferences and are unable to select the optimal price signals, resulting in a linearly growing regret. However, the 25 day (yellow) learning stage algorithm is able to adequately learn the population's parameters and select optimal prices after that.

4.5 Conclusion

In this work, we presented a multi-armed bandit problem formulation for an electricity aggregator attempting to run a real-time pricing program for load shaping (e.g., to reduce demand at peak hours, integrate more intermittent renewables, track a desired daily load profile, etc). We made use of a constrained Thompson sampling heuristic, Con-TS-RTP, as a solution to the *exploration/exploitation* problem of an aggregator *passively* learning customers' price sensitivities while broadcasting price signals that influence customers to alter their demand to match a desired load profile. The proposed

Con-TS-RTP algorithm permits day-varying target load profiles (i.e., multiple target load profiles reflecting renewable forecasts and desired demand patterns) and takes into account the actual operational constraints of a distribution system to ensure that the customers receive adequate service and to avoid potential grid failures. Additionally, our setup accounts for complex electricity usage patterns of the customers by classifying different load clusters based on electricity demand and load flexibility. We discussed a regret guarantee for the proposed Con-TS-RTP algorithm which bounds the total number of suboptimal price signals broadcasted by the aggregator. Furthermore, we discussed an operational reliability guarantee that ensures the power distribution system constraints are upheld with high probability throughout the run of the Con-TS-RTP algorithm.

Chapter 5

Real-World Implementations

5.1 Mobility-Aware Smart Charging of Electric Bus Fleets

5.1.1 Introduction

Due to the potential reduction in operational costs [154], elimination of tailpipe emissions [155], and encouragement from government agencies [156], transit systems have started to purchase electric buses over the traditional diesel or compressed natural gas (CNG) buses. At surface level, replacing traditional buses with electric buses might seem like a simple task; however, there are many obstacles preventing a transit system from simply assigning electric buses to existing routes that were previously served by diesel buses.

The two most fundamental obstacles are the restricted travel distance and lengthy recharge time of electric buses. Even with recent advances in electric transportation and battery technology, modern electric buses are commonly restricted to operate within 20%-95% state of charge (SOC) to prevent stressing the batteries and reducing lifespan

[157]. Combining this SOC limitation with the high cost of large battery packs, most electric buses are currently inferior to diesel/CNG buses in operational range. Second, the recharging process of an electric bus takes significantly more time than the refueling process of a diesel/CNG bus[157]. Additionally, due to the lengthy recharge time and limited charging infrastructure, the transit system dispatcher must be mindful of how the fleet's recharging infrastructure is managed in order to provide adequate energy to serve routes.

Despite the aforementioned challenges, the promise of eliminating large amounts of greenhouse gas emissions from transit buses has enticed early adopters to operate fleets of electric buses since the early 21st century [154]; however, it is likely that these electric bus fleets are operating suboptimally in their recharging strategies and route assignments [158]. Accordingly, there has been increasing interest in the optimal operation and infrastructure planning of electric bus fleets.

The first category of work that studies optimized charging for electric bus fleets considers the assignment of buses to routes as given, i.e., the times at which each bus is parked and is available to recharge is predetermined. Specifically, the authors of [159] present an optimization model for installing charging infrastructure and sizing batteries for a cost-effective electric bus fleet. Similarly, the authors of [158] consider infrastructure planning as well as fleet composition and the recharging process, with the goal of minimizing total cost of ownership (TOC) of the fleet. Moving away from infrastructure planning, the authors of [160] present a method to minimize battery aging costs of an electric bus fleet recharging at nighttime. The authors of [161] present the cost savings from controlling the charging thresholds for a fleet of electric buses serving one route continuously in Tallahassee, Florida. Similarly, the authors of [162] present a MILP framework for scheduling bus charging and show the potential cost savings from an electric bus fleet in Davis, California. Furthermore, [163] presents a charging strategy for

electric buses with fast charging infrastructure.

Considering both route assignment and charge scheduling (i.e., the mobility-aware setting) the authors of [164] present a k-greedy solution method to maximize travel distance of each electric bus within the fleet. A work similar to ours, [165], presents a linear formulation for route assignment and charge scheduling; however, the aim is to minimize the number of electric buses needed to replace an existing diesel fleet. Hence, the variability of electricity costs are not considered.

Similar to the aforementioned papers, the work presented in this section considers both the route assignment and charge scheduling problem of an electric bus fleet. However, the presented approach is able to improve upon previous mobility-aware work by accounting for time-varying electricity prices, utilizing on-site solar energy generation, and providing a minimal cost schedule for the fleet's daily operation.

Organization: Section 5.1.2 describes the problem of a fleet dispatcher operating a fleet of electric buses and proposes a mixed integer linear program (MILP) formulation that solves for the minimal cost route assignments and recharging schedule. Section 5.1.3 presents the results of the MILP for the real-world example of Stanford's Marguerite Shuttle Transit System.

5.1.2 Problem Description

We consider a fleet dispatcher attempting to optimize an electric bus transit system. Specifically, the fleet dispatcher aims to assign electric buses to serve the daily trips and schedule the recharging of the buses to minimize electricity cost (e.g., recharging during the inexpensive electricity rates of nighttime or when solar generation is abundant while still fulfilling all required bus routes). In the following, we consider the case where the physical infrastructure (e.g., buses, chargers, parking spots, etc.) and time-tables

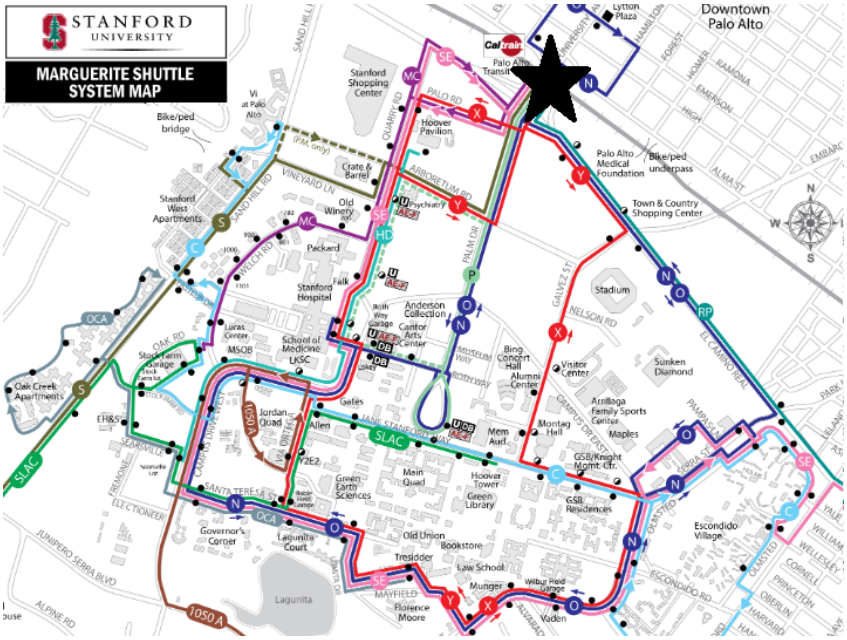

Figure 5.1: Primary service area for Stanford University's Marguerite Shuttle. Trip origins at Caltrain Palo Alto Transit Center (star). Full system map available at: https://transportation.stanford.edu/marguerite

(e.g., routes, stops, start/end times, etc.) are already established within the transit system, but not yet optimized for the aforementioned objective (as is the case for the Stanford University Marguerite Shuttle, discussed in Section 5.1.3). Given the transit system's fixed time-table and electric bus infrastructure, the fleet dispatcher seeks to answer questions such as the following:

1. Which electric bus should be assigned to each route at each time?
2. When should each electric bus be recharged?
3. Does the system need to utilize spare diesel buses to supplement the electric buses?
4. Would more infrastructure benefit the daily operation of the electric bus fleet?
5. What size of on-site solar generation system is needed to fully supply the fleet with renewable energy?

Let us consider the Stanford Marguerite Shuttle Transit System (Figure 5.1.1) which consists of 38 electric buses, 23 diesel buses, 23 electric bus chargers, and total of 20 daily routes. Currently, the assignment of buses to routes and their recharging strategy follows rules adopted by operators that work well in practice by ensuring sufficient charge is available for service. However, as we demonstrate in our numerical case study, the current assignment results in significant losses for the transit system in terms of daily operational costs and can be improved upon through a joint charge and route assignment policy. As such, in order to optimize the decision making problem of the fleet dispatcher, we formulate a mixed-integer-linear-program (MILP) to solve for both the optimal recharging schedules and route assignments for an electric bus transit system.

MILP Formulation

In the electric bus transit system, we consider one central transit center (i.e., bus depot) from which all the buses start and finish their routes as well as recharge. The buses are required to serve numerous routes throughout the service area, and each route must be served multiple times each day (i.e., the electric bus fleet is required to fulfill multiple *trips* for each route). We denote $\mathcal{S}$ as the set of scheduled trips across all routes that need to be fulfilled. For each trip $i \in \mathcal{S}$, let a_i and b_i denote the start and end time of trip i. More specifically, these are the times that a bus leaves the depot and later returns if serving trip i. If trip i is a one-way route that does not loop back to the depot, we account for the extra duration for the bus to return to the depot in b_i accordingly (i.e., the trip end time b_i accounts for "deadhead" travel). Similarly, if a route does not start at the depot, we account for the deadhead travel time to the starting location in a_i.

In order to capture the state of charge of each bus at any time t, we discretize the day into T time steps (e.g., five minute intervals) and $\mathcal{T}$ is the set of time steps for an entire day. Furthermore, let d_i be the energy consumption per time step for a bus serving trip i

(while we assume that varying traffic conditions across different routes can affect energy consumption rates, we assume that the buses are identical in their energy consumption when they serve the same route). Let $\mathcal{K}$ be the set of electric buses and $\mathcal{N}$ be the set of electric bus chargers installed at the central depot. For each charger $n \in \mathcal{N}$, u_n is the charging rate. Additionally, let $\mathbf{p} = [p(t)]_{t \in \mathcal{T}}$ be the vector of electricity prices for an entire day. We denote as E^k_{min} and E^k_{max} the minimum and maximum energy levels for bus k, respectively. The fleet dispatcher usually sets $E^k_{min} > 0, \forall k \in \mathcal{K}$ for safety precautions. Let $g(t)$ be the available on-site solar generation at time t, which we assume is known at the time of dispatch. Moreover, we assume that the electricity used from the on-site solar generation is free for the operator. Last, we denote the initial energy level of bus k as e^k_0.

Next, we describe the decision variables used in the MILP formulation. We set the binary variable $X^k_i(t)$ to 1 if bus k is serving trip i at time t and 0 otherwise. We set the binary variable $Z_k(t)$ to 1 if bus k is charging at time t and 0 otherwise. We set the binary variable $Y^k_n(t)$ to 1 if bus k is occupying charger n at time t and 0 otherwise. We use the variable $E^k(t)$ to track the energy level of bus k at time t. Lastly, let $V(t)$ be the total amount of electricity that the dispatcher purchases from the grid at time t, and $S(t)$ be the amount of electricity that buses obtain from the available on-site solar generation at time t. With the necessary notation and decision variables, the joint charging and routing MILP for the electric bus fleet can be formulated as follows:

$$\text{Minimize} \quad \sum_{t\in\mathcal{T}} \; p(t)V(t) \tag{5.1a}$$

Subject to:

$$Z^k(t) + \sum_{i\in\mathcal{S}} X_i^k(t) \leq 1, \qquad \forall k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1b}$$

$$\sum_{k\in\mathcal{K}} X_i^k(t) = 1, \qquad \forall i \in \mathcal{S}, t \in [a_i, b_i] \tag{5.1c}$$

$$X_i^k(t+1) = X_i^k(t), \qquad \forall i \in \mathcal{S}, k \in \mathcal{K}, t \in [a_i, b_i - 1] \tag{5.1d}$$

$$\sum_{k\in\mathcal{K}} Y_n^k(t) \leq 1, \qquad \forall n \in \mathcal{N}, t \in \mathcal{T} \tag{5.1e}$$

$$\sum_{n\in\mathcal{N}} Y_n^k(t) = Z^k(t), \qquad \forall k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1f}$$

$$E^k(t) = E^k(t-1) + \sum_{n\in\mathcal{N}} u_n Y_n^k(t) - \sum_{i\in\mathcal{S}} d_i X_i^k(t), \qquad \forall k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1g}$$

$$\sum_{n\in\mathcal{N}} \sum_{k\in\mathcal{K}} Y_n^k(t) u_n = V(t) + S(t), \qquad \forall t \in \mathcal{T} \tag{5.1h}$$

$$E_{min}^k \leq E^k(t) \leq E_{max}^k, \qquad \forall k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1i}$$

$$X_i^k(t) \in \{0, 1\}, \qquad \forall i \in \mathcal{S}, k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1j}$$

$$Y_n^k(t) \in \{0, 1\}, \qquad \forall n \in \mathcal{N}, k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1k}$$

$$Z^k(t) \in \{0, 1\}, \qquad \forall k \in \mathcal{K}, t \in \mathcal{T} \tag{5.1l}$$

$$0 \leq S(t) \leq g(t), \qquad \forall t \in \mathcal{T} \tag{5.1m}$$

$$E^k(0) = e_0^k, \qquad \forall k \in \mathcal{K} \tag{5.1n}$$

$$E^k(T) = e_0^k, \qquad \forall k \in \mathcal{K}. \tag{5.1o}$$

The objective in equation (5.1a) aims to minimize the daily electricity cost of recharging the bus fleet. Constraint (5.1b) ensures that a bus is either charging, serving a trip, or parked in the depot (without charging). Constraint (5.1c) ensures that all the required

daily trips will be served by a bus. Constraint (5.1d) ensures that one unique bus will serve each trip (i.e., a trip cannot be interrupted to switch buses). Constraint (5.1e) ensures that a bus can only occupy one charger per time slot. Constraint (5.1f) guarantees that if a bus is occupying a charger, then it is charging. Constraint (5.1g) calculates the energy level of each bus in each time epoch. Specifically, the energy level at time t is equal to the energy level at time $t-1$ plus the charged energy if the bus was charging or minus the spent energy if the bus was serving a trip. Constraint (5.1h) ensures that buses obtain electricity from either the grid or on-site solar. Constraint (5.1i) ensures that the buses operate above a desired minimum energy threshold. Constraints (5.1j)-(5.1l) are binary constraints on the decision variables. Constraint (5.1m) ensures that the solar energy used by the bus fleet is less than or equal to available solar generation at time t. Lastly, constraint (5.1n) sets the initial energy of each bus and constraint (5.1o) ensures that the energy level of the fleet returns to the initial value so the same route assignments and charge schedule can be used for the next day.

Behind-the-Meter Solar Integration

To exploit free on-site solar energy and to avoid injecting excess power back into the distribution grid, the fleet dispatcher prioritizes recharging the buses during periods when solar generation is available. Only if there is not enough solar energy, then the fleet dispatcher should purchase electricity from the grid. As stated in Section 5.1.2, to accommodate behind-the-meter solar integration, the dispatcher's MILP formulation makes use of a daily solar forecast, $g(t)|_{t=1,\dots,T}$. This can be estimated from forecast models, including those that use weather forecasts, and previous years' solar irradiance data. We note that if the solar generation is over-estimated, then the fleet will have to purchase more expensive grid energy potentially during peak times such as midday. As such, a conservative estimate is preferred as cheaper electricity can be procured in the

Table 5.1: PG&E E-20 Rate Structure

Time Interval	**Label**	**Price**
12:00am-8:30am	Off-Peak	$0.08422/kWh
8:30am-12:00pm	Partial-Peak	$0.11356/kWh
12:00pm-6:00pm	Peak	$0.16127/kWh
6:00pm-9:30pm	Partial-Peak	$0.11356/kWh
9:30pm-12:00am	Off-Peak	$0.08422/kWh

late night period. Future work could investigate moving-horizon solution methods to account for stochastic solar generation and update the route and charge assignments in real-time as solar energy data becomes available.

5.1.3 Case Study

As stated in the introduction, the motivation for the proposed MILP for electric bus fleets is the real-world Stanford Marguerite Shuttle Transit System (Figure 5.1.1). The Marguerite Shuttle System is free, open to the public, and operates seven days a week all year traversing the Stanford campus and surrounding areas. More specific information can be found at https://transportation.stanford.edu/marguerite.

Stanford Marguerite Shuttle System Information

Currently, the Marguerite fleet consists of 23 diesel buses and 38 electric buses from BYD split into 10 K7 models with battery capacity of 197kWh, 10 K9 models and 18 K9M models, both with 324kWh battery capacity. Additionally, the central depot is equipped with 23 double port electric bus chargers where each port can deliver up to 40kW. Each bus can be charged from one or two ports for a total power of 80kW. For the electricity rates, we consider PG&E's E-20 electricity rate structure for off-peak, partial-peak, and peak hours. The electricity rates are given in Table 5.1. Furthermore, the Marguerite Shuttle system serves up to 20 unique routes on any given day. Across

Table 5.2: Stanford Marguerite Shuttle Route Information

Route Name	**Daily Trips**	**Trip Miles**
C Line	33	7.00
C Limited	11	4.60
MC Line (AM/PM)	46	3.00
MC Line (Mid Day)	11	5.10
P Line (AM/PM)	56	2.50
P Line (Mid Day)	11	4.00
Research Park (AM/PM)	24	10.40
X Express (AM)	12	1.20
X Line	44	4.60
X Limited (AM)	10	2.00
X Limited (PM)	10	1.50
Y Express (PM)	20	1.20
Y Line	44	4.60
Y Limited (AM)	10	2.40
Y Limited (PM)	10	2.00
Totals	**352 trips/day**	**1431.50 miles/day**

all 20 routes, 15 of them are mainly fulfilled by electric buses, meaning that the electric bus fleet is required to make 352 trips per day, during weekdays. The specific routes and mileages are listed in Table 5.2. For the purposes of this numerical example, the solar forecast used was an average daily solar generation calculated from October 2019 with a maximum generation of 1 MW. The solar forecast is displayed in Figure 5.2.

Simulation Results

The proposed MILP was implemented in Matlab making use of CVX and Mosek. All numerical experiments were run on a laptop with 16 GB of RAM and 3.5 GHz Intel i7 processor. This section reports on the charging schedule, route assignments, and cost savings when comparing the proposed MILP solution with on-site solar generation, without on-site solar generation, and the status quo (i.e., the status quo is the actual operations of the Stanford Marguerite Fleet from 7-October-2019) which does not yet exploit free on-site solar generation.

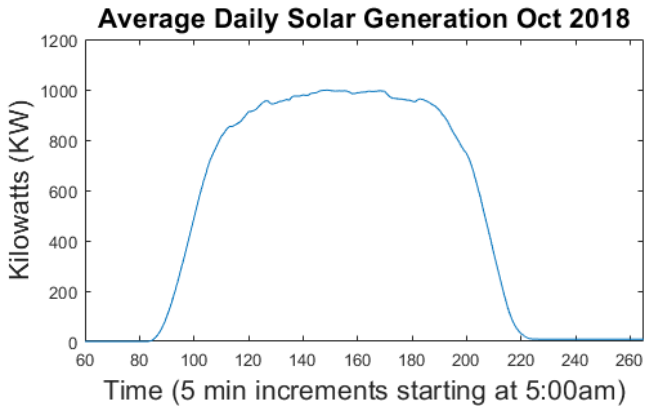

Figure 5.2: Average daily solar generation for a 1 MW on-site installation. Data averaged from CAISO renewable database in October 2019.

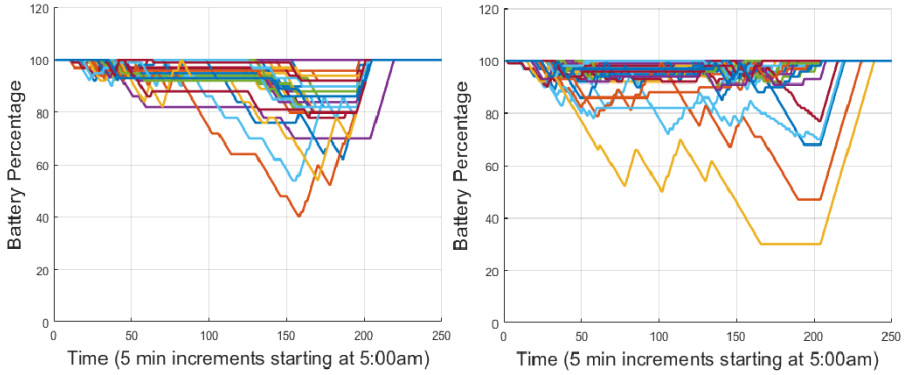

Figure 5.3: Left: Battery levels for each electric bus when considering a fleet without available on-site solar. Right: Battery levels for each electric bus when optimizing with available on-site solar generation.

Figure 5.3 presents the energy levels of each bus in the fleet during the day when the dispatch is generated through our proposed MILP. Time on the x-axis begins at 5:00am, as this is the start of the earliest route that must be fulfilled. The left plot shows the energy levels of the buses when the MILP is not utilizing on-site solar generation. The right plot shows the battery levels of the buses when the MILP accounts for on-site solar generation. It will become more clear when examining Figure 5.4 that the buses charge more during midday in the right plot than the left, to make use of the free on-site solar.

Figure 5.4 presents the total charging power of the fleet across the entire day. The red curve presents the total charging power for the MILP solution that does not exploit on-site

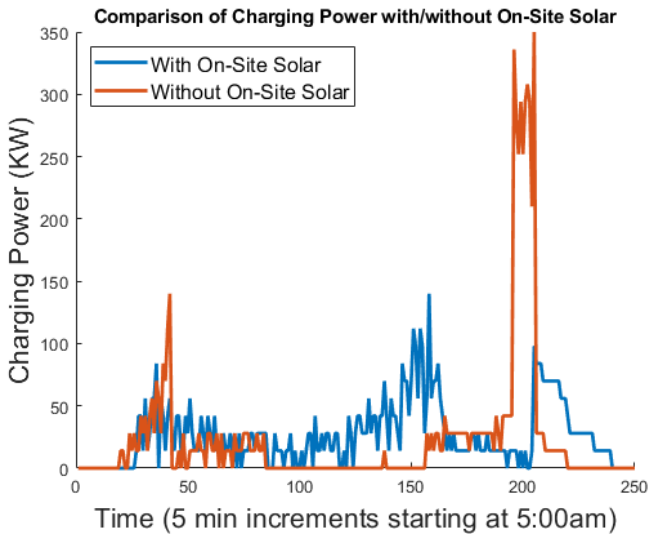

Figure 5.4: Total charging power of the fleet throughout the day. Blue: Solution accounting for on-site solar generation. Red: Solution does not include on-site solar generation.

solar generation. Conversely, the blue plot shows the fleet's total charging power from the MILP solution that does account for on-site solar generation. It is clear from this plot that the solution that accounts for on-site solar (blue) is able to charge in the middle of the day when solar is abundant; however, the solution that does not exploit solar (red) does not charge during the midday as the electricity prices are highest at this time. Instead, the fleet has a spike in charging power in the evening when electricity rates are decreased. This large transient in the evening could be detrimental to grid stability, increase in harmonics, accelerate aging of grid assets (i.e. transformers) and could potentially lead to demand charges for the fleet dispatcher due to high power consumption. As such, the solution making use of on-site solar generation with a forecasting method is preferable.

Last, Figure 5.5 presents the daily electricity costs for the three different test cases. Case A: Status Quo. We had access to the data from the operations of the Stanford Marguerite fleet on 7-October-2019 and calculated the cost of charging the fleet under the E-20 rate structure. As such, under normal operation, the daily operational cost was

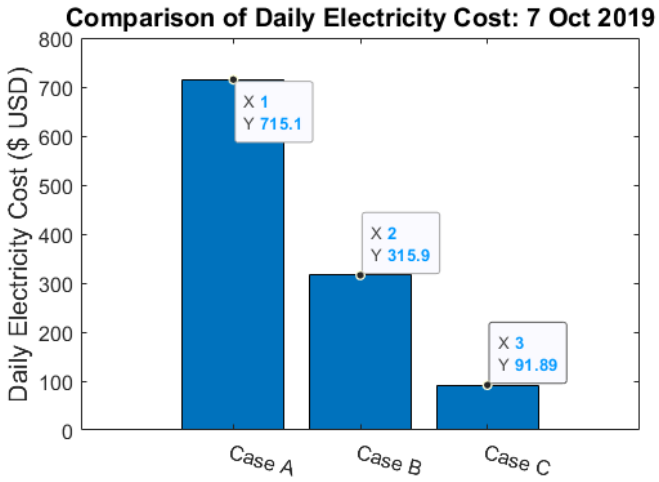

Figure 5.5: Price Comparison for 3 difference regimes: Case 1: Status Quo, electric bus charging data obtained from real-implementation (Stanford Marguerite Shuttle) on 7-Oct-2019. Case 2: Mobility-Aware MILP solution for same routes and buses as Case A, *without* on-site solar generation. Case 3: Mobility-Aware MILP solution for same routes and buses as Case A, *with* on-site solar generation.

$715.10 USD. Case B corresponds to the solution of the proposed MILP with the same routes, buses, and chargers as Case A; however, the mobility-aware solution reassigned buses to new trips and rescheduled the charging of each bus. In Case B, the MILP solution did not account for on-site solar and the daily cost was $267.90 USD. Last, Case C was identical to Case B; however, the MILP accounted for the on-site solar generation and had access to the daily solar forecast. As such, the daily cost was reduced to $61.89 USD. From these results, it is evident that the fleet dispatcher benefits from the MILP formulation for routing and charging (55% decrease in cost in Case B).

5.1.4 Conclusion

In this work, we investigated the joint route assignment and charge scheduling problem of a transit system dispatcher operating a fleet of electric buses in order to maximize solar energy integration and reduce energy costs. We considered a complex bus transit system with preexisting routes, limited charging infrastructure, limited number of electric

buses, and time-varying electricity rates. We presented a mixed integer linear program (MILP) that yields route assignments and charging schedules using daily solar forecasts. We presented numerical results from a real-world case study with Stanford University's Marguerite Shuttle to demonstrate the cost-saving benefits of our solution and highlight the significant cost savings compared to the status quo.

Future work includes investigating a moving-horizon solution approach to account for stochastic solar generation. Additionally, we would like to add traditional diesel routes to the optimization to further minimize emissions and to expand the clean operation of the electric bus fleet. Further future work can include performing field test experiments with real buses during operational hours, determining the optimal solar capacity to fully charge the electric bus fleet, and quantify the value and size of onsite solar and battery combination for resiliency.

Acknowledgement: The authors would like to thank the Stanford Transportation team for the support, discussions, and information about operations. This work was funded by the California Energy Commission under grant EPC-17-020. SLAC National Accelerator Laboratory is operated for the US Department of Energy by Stanford University under Contract DE-AC02-76SF00515.

5.2 Real-Time Electric Vehicle Smart Charging at Workplaces

5.2.1 Introduction

Many large companies across the United States are installing Electric Vehicle Supply Equipment (EVSEs) within their parking lots to allow for employees and visitors to charge their Electric Vehicles (EVs) during their stay [166]. There are numerous benefits that can come from workplace charging stemming from the highly flexible nature of EV charging at workplaces [10, 167] (i.e., EVs are plugged in for long durations and their charging can be shifted depending on other conditions including behind-the-meter solar generation or electricity rates). For example, Google (who has installed EVSEs at several of its Bay Area locations for employee EV charging) allows employees to leave their EVs connected to the EVSE for long durations and does not force them to unplug the EV after reaching full charge, whereas other workplaces might require employees to move EVs after charging is complete. The approach of not enforcing EVs to be removed once fully charged may require more investment in EVSEs to satisfy the needs of the employees, but the investment would allow for greater flexibility that can be harnessed by a smart charging algorithm to benefit the workplace (by providing financial savings) and the utility (by enabling its customers to respond to price signals) [168, 4].

The Grid Integration, Systems, and Mobility (GISMo) group within the SLAC National Accelerator Laboratory and the Smart Infrastrucuture Systems Lab (SISL) at UCSB are working on a project to study the impacts of electric vehicles on the distribution system and designing tools and algorithms to be used by workplace charging station operators to efficiently manage asset operations and minimize electricity charges. In collaboration with Google, the project has access to historical charging data from previous

years at multiple Bay Area workplace parking locations for Google (within PG&E's service territory). The locations exhibit typical workplace charging behavior in that most session start times occur around 8 to 9 A.M. and most session end times occur between 4 and 6 P.M. The dataset contains 15-minute interval data for over 10,000 charging sessions and between 700 and 1000 sessions per month. The datasets include pertinent information such as charging session start times, end times, 15 minute average power delivered, total energy delivered, and many others.

Challenges: There are several key challenges to designing an algorithm to schedule the charging profiles of numerous EVs. First and foremost, the algorithm must run in real-time without knowledge of the future EV arrivals. The algorithm must adapt its planned power schedules as more information is revealed (i.e., as more EVs arrive to the parking lot). Second, contrary to most smart charging algorithms presented in the literature, the algorithm must be able to function with limited information from each EV [169], [170]. Specifically, when an EV plugs in, the algorithm does not get access to the EV's State of Charge (SoC) nor does it know the EV's future departure time (most level 2 chargers do not sense EV SoC nor do they request user input for future departure times). As such, our smart charging algorithm must predict how much maximum charge an EV may consume as well as the EV's future departure time. Such challenges have been acknowledged in past papers including [171, 172, 173, 174]. Third, all of the EV charging schedules within a parking lot are coupled due to the local transformer capacity constraint [175]. As such, the algorithm cannot over-allocate power at any given time; therefore, the algorithm should make use of a model of the future EV arrivals to avoid over allocating power due to unexpected arrivals.

Contribution: In this work, we present an EV smart charging algorithm for workplace parking lots equipped with EVSEs that operates in real-time to minimize electricity cost from time-of-use electricity rates and demand charges while ensuring that the owners of

the EVs receive adequate levels of charge and the entire system safely operates within the local transformer capacity constraint. Notably, our algorithm is SoC and departure time agnostic and uses both scenario generation to account for each EV's unknown future departure time as well as certainty equivalent control to account for the unknown EV arrivals in the future. We build models from the Google dataset for each day of the week and our algorithm uses these models as the expected future when optimizing the EV charging schedules.

5.2.2 Problem Description

Base Optimization

Before we discuss the smart charging algorithm that operates in real-time with a predicted model of the future, let us consider the simpler problem of scheduling the charging power to a single EV if we know the arrival time, departure time, and maximum energy. Specifically, let us break one day into 15 minute segments (96 segments total) and let us denote $t = 1, \ldots, T$ as the time steps each day and $T = 96$. Furthermore, let x be a Tx1 indicator vector where the t-th element is 1 if the EV is plugged in at time t and 0 if not. Additionally, let p be a Tx1 vector where the t-th element corresponds to the electricity rate at time t from the local grid (\$/kWh). Last, let us denote our decision variable e as a Tx1 vector where the t-th element indicates how much energy will be delivered to the EV in time epoch t. With this notation, we can formulate the

charge scheduling optimization problem for the singular EV as follows:

$$\max_{e} \quad w_1 \log(e^T x + 1) - w_2 p^T e \tag{5.2a}$$

subject to:

$$0 \leq e \leq e_{max}, \tag{5.2b}$$

$$e_{del\_min} \leq e^T x \leq e_{del\_max}. \tag{5.2c}$$

The objective function (5.2a) has two weight coefficients, w_1 and w_2, for each of the two terms. The first term corresponds to the utility that the owner of the EV receives for the energy that their EV receives. The logarithmic utility term was chosen to model the diminishing returns in user utility for EVs receiving excessive amounts of energy (e.g., the first 20kWh charged is more valuable to the EV owner than the second 20kWh). The second term corresponds to the cost of the energy that is purchased from the local grid (we will add demand charges in the next section). The weights, w_1 and w_2, let us adjust the relative importance of user utility from receiving energy and energy cost. Constraint (5.2b) ensures that the energy delivered at each time step is non-negative and not greater than the energy e_{max} that the EVSE can deliver in a single timestep. Constraint (5.2c) ensures that the amount of energy the EV receives is greater than a minimum amount e_{del_min} and less than a maximum amount e_{del_max}.

5.2.3 Real-Time Smart Charging Algorithm

In this section, we modify the base optimization presented in Section 5.1.2 that will be solved at each 15 minute time step $t = 1, \ldots, 96$ each day. Specifically, assume that at an arbitrary time t, there are currenly I EVs plugged in at the location.

Departure time scenario generation: When an EV plugs in to an EVSE, we generate N

potential departure times for that EV and create a scenario in our optimization problem for each potential departure time and solve the optimization across all scenarios. As time progresses, if a potential departure time is no longer feasible (i.e., the potential departure time is the current time step and the EV has not yet departed), then that scenario is removed from the optimization via dynamic scenario weights (let C_n be a weight coefficient for each scenario that is set to 0 if the scenario is no longer feasible.). Furthermore let $x_{i,n}$ be the Tx1 binary vector that indicates when EV i is available to charge in scenario n.

Certainty equivalent control for future EV arrivals: We make use of our dataset to generate a model for an average day that consists of estimated arrival times, departure times, and energy requests for each day of the week. We then use these daily models in the real-time optimization to account for the unknown future EV arrivals. Specifically, at time t, let us assume that there are J EVs in the certainty equivalent daily model that are expected to arrive in the future. Let x_j be the Tx1 binary vector indicating when EV j is available to charge. The decision variables that determine how much energy is delivered at a given time t are Tx1 vectors labeled as $e_1, \ldots, e_I$ for the actual EVs plugged in and $e_{I+1}, \ldots, e_{I+J}$ for the future EV arrivals from the model. Additionally, let p_d be the demand charge (\$/kW) to be assessed on the monthly peak load. The optimization

at time t can be written as:

$$\max_{\substack{e_1,\ldots,e_I \\ e_{I+1},\ldots,e_{I+J}}} \sum_{i=1}^{I}\sum_{n=1}^{N} \frac{1}{C_n}\Big[w_1 \log({e_i}^T x_{i,n} + 1) - w_2 p^T e_i\Big] \tag{5.3a}$$
$$+\sum_{j=I+1}^{I+J} \Big[w_1 \log({e_j}^T x_j + 1) - w_2 p^T e_j\Big] - w_2 p_d \hat{e}_{inc}$$

subject to:

$$0 \le e_k \le e_{max}, \qquad \forall k = 1, \ldots, I+J \tag{5.3b}$$

$${e_i}^T x_{i,n} \ge e_{del\_min}, \qquad \forall i = 1, \ldots, I, \tag{5.3c}$$
$$\forall n = 1, \ldots, N,$$

$${e_j}^T x_j \ge e_{del\_min}, \qquad \forall j = I+1, \ldots, I+J, \tag{5.3d}$$

$$\sum_{k=1}^{I+J} e_k(t) \le e_{trans}, \qquad \forall t = 1, \ldots, T, \tag{5.3e}$$

$$\hat{e}_{inc} \ge \sum_{k=1}^{I+J} e_k(t) - \hat{e}_{old}, \quad \forall t = 1, \ldots, T, \tag{5.3f}$$

$$\hat{e}_{inc} \ge 0. \tag{5.3g}$$

The first term of the objective function (5.3a) accounts for all I EVs currently plugged in and their N potential departure times each while the second term of the objective function accounts for all J EVs in the future model. The third term ($w_2 p_d \hat{e}_{inc}$) accounts for the demand charge from increasing the current month's peak demand ($\hat{e}_{old}$) by $\hat{e}_{inc}$. Constraint (5.3b) ensures that the energy delivered is nonnegative and less than the EVSE max e_{max}. Constraint (5.3c) ensures a minimum amount of energy is delivered to each EV currently plugged in, constraint (5.3d) ensures a minimum amount of energy is delivered to each EV in the future model, and constraint (5.3e) ensures that the sum of all energy delivered by the EVSEs at each time t does not exceed the transformer constraint e_{trans}.

Constraints (5.3f)-(5.3g) keep track of any increase to the current month's peak load for the demand charge. We note that $\hat{e}_{old}$ corresponds to the previous peak energy demand that has been observed during the month. The pseudocode for the daily algorithm can be viewed in Algorithm 5 below.

Algorithm 5 Real-Time Smart Charging

```
for each day do
    Update current parking lot state
    for each 15 minute interval t do
        if new departure from parking lot then
            Update parking lot state
        end if
        if new arrival to parking lot then
            Generate N potential departure times for new arrival
            Update Parking lot state
        end if
        Formulate optimization for time t:
        for each EV i plugged in at time t do
            Add EV i to total objective function (5.3a)
            Add EV i to active constraints (5.3b)-(5.3g)
        end for
        for each future EV j in daily model t_model > t do
            Add EV j to total objective function (5.3a)
            Add EV j to active constraints (5.3b)-(5.3g)
        end for
        Solve optimization (5.3a)-(5.3g) for time t
        Store planned energy schedule for each EV i
        Set each EVSE's output power for the current 15 minute interval
        Update peak load ê_old for demand charge calculation (if a new peak load is observed)
    end for
end for
```

5.2.4 Case Study Results

We examine a two week period from June 17 - June 29 in 2019 at a Bay Area workplace from our Google EV dataset. The location has 57 level 2 EVSEs with 50-100 EVs arriving each weekday and is under PG&E's E-19 rate structure.

First, the EV charging session data was filtered by weekday and then filtered again by arrival time. Namely, each charging session was put into one of 12 possible groups

corresponding to 2-hour windows for the arrival times (e.g., an EV charging session that started at 9:48am would be stored in the 8:00am-10:00am group). Once this was done, daily arrival time histograms were generated and the average stay duration and average energy consumption were calculated for each of the 12 groups. The average arrivals per weekday, the average arrivals per 2 hour window, the groups' average stay durations, and the groups' average energy consumption were then used to create the algorithm's future model each day and to generate potential departure times for each EV arrival. The daily average arrival time histogram can been viewed in Figure 5.6. The 2 hour groups' average energy consumption and average stay duration can be viewed in Table 5.3. Last, we note that all of these simulations were done in Python with CVX and Mosek on a Laptop with an i7 processor and 16gb of RAM. Solving the daily smart charging optimization (5.3a)-(5.3g) at each 15 minute time step took less than a second, so the algorithm is fit to run in real-time. Moreover, the optimization problem's complexity is not affected by the number of arriving EVs each day; rather, the problem size grows only as the number of chargers increases. Additionally, for implementation, (2a)-(2g) has to be solved every 15 minutes but for the 57 chargers in our case study, (2a)-(2g) was solved in less than a second. Thus, the algorithm is scalable and there is significant extra time for computation for a larger dataset (i.e., more chargers at the parking lot).

Example Day's Charging Schedule Evolution

Figure 5.7 shows an example as to how the real-time optimization changes the planned power output as time progresses and more EVs arrive. These plots correspond to a Monday with 78 EVs arriving to 57 EVSEs with $(w_1, w_2) = (2, 1)$ and a transformer capacity constraint of 150kW. Starting with the Top Left: This plot shows the planned charging schedule that is calculated at 12:00am on the given Monday for the entire location. There are no actual EV arrivals at this early time in the morning so the entire

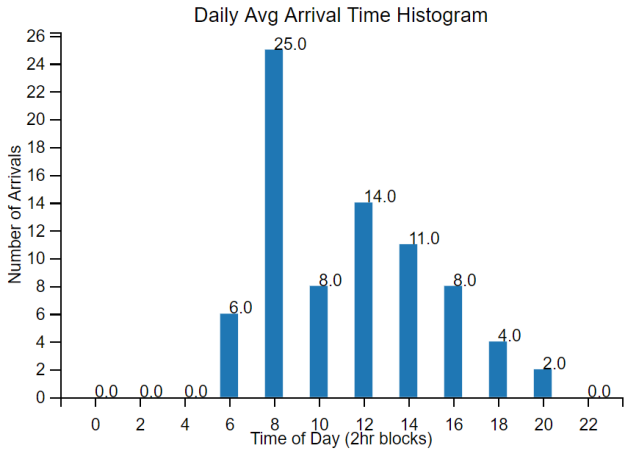

Figure 5.6: Average arrival time histogram for weekdays in 2019, split into 2 hour blocks, to be used for the daily model.

Arrival Time Block	Avg Charge Amount (kWh)	Avg Stay Duration
12:00am-2:00am	5.56	4 hrs 38 mins
2:00am-4:00am	4.00	2 hrs 02 mins
4:00am-6:00am	12.91	3 hrs 52 mins
6:00am-8:00am	14.63	5 hrs 29 mins
8:00am-10:00am	15.79	6 hrs 02 mins
10:00am-12:00pm	9.27	6 hrs 02 mins
12:00pm-2:00pm	7.41	11 hrs 15 mins
2:00pm-4:00pm	6.80	16 hrs 06 mins
4:00pm-6:00pm	7.14	16 hrs 27 mins
6:00pm-8:00pm	6.61	23 hrs 19 mins
8:00pm-10:00pm	6.78	25 hrs 54 mins
10:00pm-12:00am	7.74	10 hrs 01 mins

Table 5.3: Average charge amount (kwh) and average stay duration for different arrival time blocks.

planned charging schedule is created by looking at the future model for the day (shaded blue region). Top Right: This shows the planned charging power as of 10:45am in the morning. Everything that is not in the shaded blue region corresponds to time-steps in the past, meaning that the red plot in the white region corresponds to actual charging power. However, everything in the blue region is still estimated via the future model. Bottom Left: This plot shows the planned charging schedule as of 1:30pm. Bottom Right Plot: This plot shows the charging schedule as of 11:30pm, when there are no

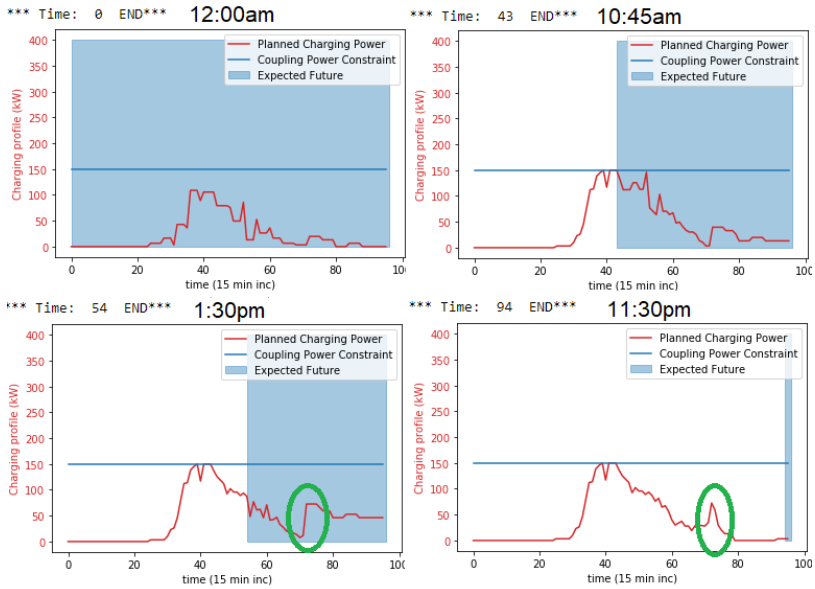

Figure 5.7: Evolution of planned charging power. Top Left: At 12:00am the algorithm is fully using the expected future as a model for what could happen that day. Top Right: At 10:45am, many EVs have arrived and the power output has started to deviate from the expected future. Bottom Left: Planned power output at 1:30pm. Bottom Right: At 11:30pm, there are no more arrivals, Red curve shows the daily power output of the 57 EVSEs. Additionally, the green circles indicate spikes in total charging power in the evening due to the algorithm waiting for cheap electricity rates to charge multiple EVs, discussed in Section 5.2.4.

more arrivals for that day. Additionally, the smart charging algorithm is able to ensure that the actual charging schedule stays below the transformer capacity constraint.

Figure 5.8 presents a comparison of a weekday's predicted daily charging schedule and the actual daily charging power that ocurred that day. As seen in Figure 5.8 the daily model does a good job predicting the future load for this location. Furthermore, note that most of the power is scheduled to be delivered during the partial-peak electricity rates during the mid-morning (8:30am-12:00pm) and the charging power decreases rapidly during the peak electricity rates mid-day (12:00pm-6:00pm).

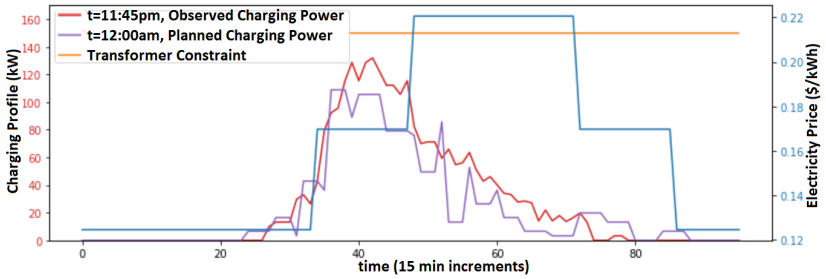

Figure 5.8: A comparison of the model's predicted daily charging schedule vs. the actual daily charging power.

Test Case Comparisons

Furthermore, in Table 5.4, we present simulation results for several test cases. In these tests, we vary (w_1, w_2) in the objective function between (2,1) and (10,1) to show how the algorithm swaps priority from minimizing electricity cost to maximizing user utility for receiving energy. Additionally, we varied the transformer constraint that couples all the EVs together between 250kW and 100kW. In Table 5.4, we also show the total energy delivered over the 2 week period and the total cost of purchasing that energy from the grid from energy rates as well as the total cost due to demand charges (displayed as percentages compared to the status quo values). The last 2 columns in Table 5.4 show whether or not the test case included a constraint that forces the EVSEs to charge at a certain rate for the first hour that a new EV is plugged in. The idea behind this is to ensure that EVs will receive some minimum amount of energy, even if they are only plugged in for a short duration, or if they arrive during peak electricity rates. We experimented with forcing the EVs to charge for the first hour of a new session at their maximum rate, half their maximum rate, a quarter of their maximum rate, and without this constraint altogether. The second to last column indicates whether or not this constraint was included, and the last column indicates the charge rate that was used for this first hour (p_{max} = 6.6kW). The first row of this table shows the total energy delivered and the total energy cost for the status quo (i.e., no algorithm in place, just the

Test #	w_1	w_2	Coupling Constraint (kW)	Energy Delivered	Electricity Purchase Cost	Demand Charge Cost	Forced Initial Charge	Initial Charge Rate
Status Quo	n/a	n/a	n/a	100%	100%	100%	n/a	n/a
01	2	1	250	50.13%	45.78%	64.50%	Yes	$\frac{1}{2}p_{max}$
02	2	1	150	50.01%	45.83%	88.76%	Yes	$\frac{1}{2}p_{max}$
03	2	1	125	50.68%	46.33%	73.96%	Yes	$\frac{1}{2}p_{max}$
04	2	1	110	50.44%	45.94%	65.09%	Yes	$\frac{1}{2}p_{max}$
05	2	1	110	44.20%	36.40%	65.09%	No	n/a
06	2	1	100	45.05%	37.45%	59.17%	No	n/a
07	10	1	100	81.93%	81.36%	59.17%	No	n/a
08	10	1	150	83.85%	81.03%	88.76%	No	n/a
09	10	1	250	84.21%	80.98%	89.94%	No	n/a
10	10	1	100	Infeasible	Infeasible	Infeasible	Yes	$\frac{1}{2}p_{max}$
11	10	1	110	81.76%	81.09%	65.09%	Yes	$\frac{1}{2}p_{max}$
12	10	1	150	83.49%	81.80%	88.76%	Yes	$\frac{1}{2}p_{max}$
13	10	1	150	Infeasible	Infeasible	Infeasible	Yes	p_{max}
13	10	1	250	87.25%	84.81%	97.63%	Yes	p_{max}

Table 5.4: Results for 13 different test cases.

energy consumed and cost for June 17-June 29, 2019). The $(w_1, w_2) = (10, 1)$ results in Tests 07-13 indicate that the smart charging algorithm was able to reduce energy costs and demand charges while delivering adequate energy (greater than 80% of the status quo energy) to the EVs.

Transformer Capacity Constraints

As shown in Table 5.4, we vary the transformer capacity constraint that couples the charging power of all the EVs. In the status quo, there is no coupling constraint and the total load peaks at 169 kW. However, with a smart charging algorithm in place, we can constrain the total load. This would allow for the location to use a smaller transformer capacity or increase their other non-EVs loads (e.g., the nearby offices can safely use more power without worrying about exceeding the transformer capacity due to the EVSEs). Figure 5.9 shows a 5 day comparison of the Test 08 load (see Table 5.4) compared to the status quo load. As seen in this plot, the smart charging algorithm was able to enforce a transformer capacity limit at 150kW without sacrificing much energy

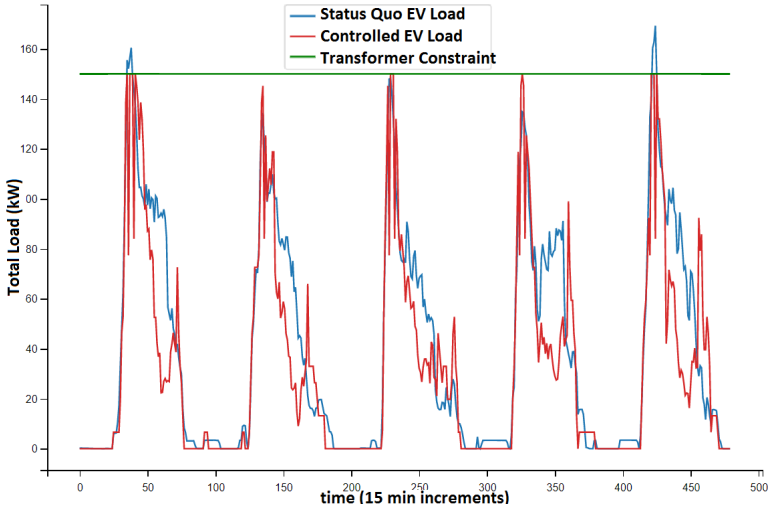

Figure 5.9: Comparison of the Monday-Friday weekday load for the status quo (i.e., no algorithm in place) vs. the smart charging algorithm with parameters as listed in Test 08. Additionally, the smart charging algorithm enforces the transformer constraint (green line).

delivered. Furthermore, the red curve (the smart charging profile) drops below the blue curve (the status quo) during the middle of each day to avoid the peak electricity rates from 12:00pm-6:00pm.

Infeasibilities Due to Transformer Capacities

Additionally, we would like to discuss the effects of removing the constraint forcing the EVSEs to charge at a certain rate for the first hour an EV is plugged in. Recall, this constraint was added to ensure that EVs are charged before they depart. If this constraint is removed, the smart charging algorithm becomes overly optimistic about each EV's departure time. Specifically, the algorithm optimistically predicts that the EVs will stay until the off-peak electricity rates in the evening and plans the daily charging schedule as seen in Plots 3,4 of Figure 5.7 (circled in green). Including the constraint that forces the EVSEs to charge a new arrival for the first hour removes the second peak that is seen in Plots 3,4 of Figure 5.7.

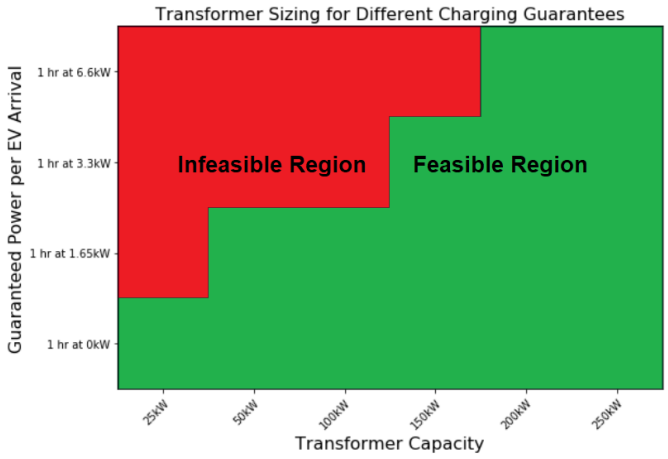

Figure 5.10: The feasible and infeasible regions of operation for various transformer capacities and forced initial charge rates for new EV arrivals.

The constraint that forces EVSEs to charge new EV arrivals for the first hour of being plugged in also affects whether or not the optimization is feasible each day. Namely, many EVs arrive around the start of the workday at 8:00am and plug in to an EVSE. If many EVSEs are forced to charge near their maximum rating at the same time, the total load of the parking lot can exceed the transformer capacity. As such, in Figure 5.10 we show the feasible and infeasible regions of operation for the smart charging algorithm with w_1=10 and w_2=1.

5.2.5 Conclusion

In this section, we presented a smart charging algorithm for a workplace parking lot equipped with EVSEs that operates in real-time to minimize electricity cost from time-of-use electricity rates and demand charges while ensuring that the owners of the EVs receive adequate levels of charge. Our algorithm uses both scenario generation to account for each EV's unknown future departure time as well as certainty equivalent model predictive control to account for the unknown EV arrivals in the future. We build models from our Google dataset for each day of the week and our algorithm uses

these models as the expected future when optimizing the EVs charging schedules. In the future, we have access to the meter data for loads other than the EVSEs for all locations in our dataset as well as the single-line diagram specifications and we hope to include more operational constraints and account for the non-EV loads in the smart charging algorithm. Furthermore, some locations have energy storage and solar generation which we hope to include into future algorithms. Finally, future work will deploy this framework in real-time field operation.

Acknowledgment: This work was funded by the California Energy Commission under grant EPC-17-020. SLAC National Accelerator Laboratory is operated for the US Department of Energy by Stanford University under Contract DE-AC02-76SF00515. Thank you to Rolf Schreiber from Google for providing data.

Chapter 6

Conclusions

6.1 Review

In this thesis, we have presented several problems within the cyber-human-physical-system space and proposed novel methodologies to improve the operation of said systems.

Namely, in Chapter 2, we leveraged recent advancements in online optimization for system scheduling with application to electric vehicle (EV) admission control and smart-charging at public parking facilities as well as scheduling for recharging and routing in EV fleets. Furthermore, in Chapter 3 we extended this framework to account for *coupled* resources (e.g., charging and discharging power) in order to schedule the operation of a community energy storage (CES) system. Additionally, in Chapter 4 we modified the machine learning framework called Thompson Sampling for the design of effective price signals for an electricity aggregator passively learning customers' price sensitivities while running a load shaping program and providing theoretical safety guarantees on critical infrastructure constraints. Last, in Chapter 5 we presented real-world implementation methods and results for optimizing workplace EV charging (at SLAC and Google campuses) in an online fashion and scheduling the charging/routing of a real electric bus fleet

(the Stanford Marguerite Shuttle) to minimize operational costs.

6.2 Future Directions

There is no doubt that large-scale cyber-human-physical-systems will continue to increase in usage, size/complexity, and impact within our societies. As such, optimizing their operation will continue to be of paramount importance. Stemming from the problems that I have worked on over the last 5 years, I foresee the following areas as both interesting problem settings and potentially fruitful areas to study.

Online Optimization in Modern Intermodal Transportation

To improve personal mobility in societal transportation, new mobility concepts and technologies are on the rise. Specifically, mobility-on-demand services such as Lyft and Uber have established themselves as a viable replacement to traditional vehicle ownership and public transit (for some users). Taking this one step further, Autonomous-Mobility-on-Demand (AMOD) system literature and test cases are on the rise as well. Here, we have fleets of autonomous (electric) vehicles rotating throughout urban areas providing ride-sharing services and charging at distributed charging stations (but with no driver). Additionally, map services such as Google Maps provide multiple modes of transit when directing a user from point A to point B (e.g., by foot, by bike, by car, by public transit, by ride-share). While these technologies do enlarge the options for personal mobility, they do not reduce total congestion or reduce emissions. Specifically, if these technologies operate independently of one another, they might actually increase congestion at various locations within an urban area and leave other modes of transportation underutilized (e.g., many people are influenced to use an AMOD service resulting in high road congestion while few are influenced to use existing public transportation). To this end, I believe a unified

approach to personal urban mobility is required. Such a unified approach would take advantage of all available modes of transportation to minimize congestion and emissions while providing adequate service (i.e., users' trips would consist of walking segments, ride-share segments, and public transit segments). Such an intermodal system would clearly require a scalable real-time (online) optimization framework in order to route and schedule the multiple modes of transportation for large population areas.

Learning to Characterize User Preferences in Modern Systems

Today, the internet of things (IoT) allows for unprecendented levels of interaction between people, their devices, and the cloud. That is, we are generating an increasing amount of data from repeated human interactions with large-scale systems. This is occurring in 'sharing economies' (think Uber and Lyft services), personal investing (countless smartphone apps available for instant trading of traditional equities and commodities and in upcoming spaces such as crypto and NFTs), social media, online advertising, personal urban mobility (sensitivities to public transportation costs, unhappiness due to inconvenient walking segments in routes, sensitivities to road tolls), electricity usage (appliance usage and thermostat settings), and online content recommendation systems, to name a few. In all of these settings, information about users' preferences is incredibly valuable to improve the operation and reliability of the systems involved; however, users often do not want to share their preferences (due to personal information privacy concerns) or they might be unable to characterize their preferences altogether (preferences could be complex and time-varying). As such, it is advantageous if these systems can passively learn user preferences from historical interactions and update them as needed. This eliminates the need for direct communication with users about their preferences and only takes advantage of the information that users are already willingly submitting to the system.

Open-Source Online (Real-Time) Optimization Software

During my Ph.D., I designed and tested numerous online optimization frameworks for various settings. Throughout this time, I read and reviewed a multitude of papers in the areas related to my work in online optimization and was exposed to various small simulation projects. However, throughout this time span, I came across no open-source online optimization software packages that were adaptable to various problem settings. As such, I believe there is a huge opportunity for the development of an online optimization software package. The package would be required to be modular to allow for users to modify, remove, and add to the functionality. The package would include a diverse set of online algorithms, various modifiable problem settings, and various models for future information. Moreover, tutorials for using, modifying, and expanding the framework would be critical for its success. I envision this as a multi-year project with several contributors designing and providing long-term support for the open-source software. The software package itself and the results and examples (in various upcoming CHPS systems or other online problems) from using the software would be interesting works on their own.

Appendix A

Proofs for Selected Results

A.1 Proof for Proposition 2.1.1

Proof: For α-competitiveness, the authors of [39] show that marginal pricing functions, operational cost functions, and Fenchel conjugates for the limited resources need to satisfy the *Differential Allocation-Payment Relationship* given by:

$$\big(p_g^l(t) - f_g^{l'}(y_g^l(t))\big)\mathrm{d}y_g^l(t) \geq \frac{1}{\alpha_g^l(t)} f_g^{l*'}(p_g^l(t))\mathrm{d}p_g^l(t) \tag{A.1}$$

for all $l \in \mathcal{L}, t = 1, \ldots, T$. For the energy-procurement operational cost in (2.28) and its Fenchel conjugate (2.35) respectively, the following derivatives are:

$$f_g^{l'}(y_g^l(t)) = \begin{cases} 0, & y_g^l(t) \in [0, s_l(t)) \\ \pi_l(t), & y_g^l(t) \in [s_l(t), s_l(t) + G_l(t)] \end{cases} \tag{A.2}$$

and

$$f_g^{l*'}(p_g^l(t)) = \begin{cases} s_l(t), & p_g^l(t) \in [0, \pi_l(t)) \\ s_l(t) + G_l(t), & p_g^l(t) \geq \pi_l(t). \end{cases} \tag{A.3}$$

The derivative of the proposed pricing function (2.43) is:

$$\mathrm{d}p_g^l(y_g^l(t)) = \begin{cases} \left(\frac{L_g}{2Rs_l(t)}\right)\left(\frac{2R\pi_l(t)}{L_g}\right)^{\frac{y_g^l(t)}{s_l(t)}} \ln\left(\frac{2R\pi_l(t)}{L_g}\right)\mathrm{d}y_g^l(t), & y_g^l(t) < s_l(t), \\ \left(\frac{L_g-\pi_l(t)}{2R(s_l(t)+G_l(t))}\right)\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)^{\frac{y_g^l(t)}{s_l(t)+G_l(t)}} \ln\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)\mathrm{d}y_g^l(t), & y_g^l(t) \geq s_l(t), \end{cases} \tag{A.4}$$

where $R = \sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l})$.

When $y_g^l(t) < s_l(t)$, $f_g^{l'}(y_g^l(t)) = 0$ and $f_g^{l*'}(p_g^l(t)) = s_l(t)$. As such, after inserting the derivative (A.4) in (A.7), we can show that the Differential Allocation-Payment Relationship holds when $\hat{\alpha}_g^l(t) \geq \ln\left(\frac{2R\pi_l(t)}{L_g}\right)$ as long as $R \geq \lceil \frac{eL_g}{2\max_{\mathcal{L},\mathcal{T}} \pi_l(t)} \rceil$. The constraint on R ensures $\hat{\alpha}_g^l(t) \geq 1$.

Similarly, when $y_g^l(t) \geq s_l(t)$, $f_g^{l'}(y_g^l(t)) = \pi_l(t)$ and $f_g^{l*'}(p_g^l(t)) = s_l(t) + G_l(t)$. As such, after inserting the derivative (A.4) in (A.7), we can show that the Differential Allocation-Payment Relationship holds when $\hat{\hat{\alpha}}_g^l(t) \geq \ln\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)$. Now, let $\alpha_g^l(t) = \max\{\hat{\alpha}_g^l(t), \hat{\hat{\alpha}}_g^l(t)\}$ and because (A.7) holds for the pricing function, cost function, and Fenchel conjugate, the remainder follows from Lemma 1 and Theorem 2 in [39]. ■

A.2 Proof for Proposition 2.1.2

Proof: Similar to Proposition 2.1.1, we show the pricing function, operational cost function, and Fenchel conjugate for the limited resource satisfy the Differential

Allocation-Payment Relationship in (A.7) with underestimated solar generation amounts $\underline{s}_l(t, t_{current})$. The derivatives of the energy-procurement operational cost in (2.28) and its Fenchel conjugate (2.35) remain the same as in (A.2) and (A.3), respectively. The derivative of the proposed pricing function (2.43) with underestimated solar $\underline{s}_l(t, t_{current})$ is given by:

$$\mathrm{d}p_g^l(y_g^l(t), t_{current}) = \begin{cases} \left(\frac{L_g}{2R\underline{s}_l(t,t_{current})}\right)\left(\frac{2R\pi_l(t)}{L_g}\right)^{\frac{y_g^l(t)}{\underline{s}_l(t,t_{current})}} \\ \times \ln\left(\frac{2R\pi_l(t)}{L_g}\right)\mathrm{d}y_g^l(t), & y_g^l(t) < \underline{s}_l(t, t_{current}), \\ \left(\frac{L_g-\pi_l(t)}{2R(\underline{s}_l(t,t_{current})+G_l(t))}\right)\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)^{\frac{y_g^l(t)}{\underline{s}_l(t,t_{current})+G_l(t)}} \\ \times \ln\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)\mathrm{d}y_g^l(t), & y_g^l(t) \geq \underline{s}_l(t, t_{current}), \end{cases} \tag{A.5}$$

where $R = \sum_{\mathcal{L}} M_l(C_l + E_l + \frac{1}{M_l})$.

When $y_g^l(t) < \underline{s}_l(t, t_{current})$, $f_g^{l'}(y_g^l(t)) = 0$ and $f_g^{l*'}(p_g^l(t)) = s_l(t)$. As such, after inserting the derivative (A.5) in (A.7), we can show that the Differential Allocation-Payment Relationship holds when $\hat{\alpha}_g^l(t) \geq \frac{\overline{s}_l(t,t_{current})}{\underline{s}_l(t,t_{current})} \ln\left(\frac{2R\pi_l(t)}{L_g}\right)$ as long as $R \geq \lceil \frac{eL_g}{2\max_{\mathcal{L},\mathcal{T}} \pi_l(t)} \rceil$. However, since the upper forecast is nonincreasing and the lower forecast is nondecreasing, the worst case forecast will occur at $t_{current} = 1$. As such we can write $\hat{\alpha}_g^l(t) \geq \frac{\overline{s}_l(t,1)}{\underline{s}_l(t,1)} \ln\left(\frac{2R\pi_l(t)}{L_g}\right)$.

Similarly, when $y_g^l(t) \geq \underline{s}_l(t, t_{current})$, Setting $f_g^{l'}(y_g^l(t)) = \pi_l(t)$ minimizes the LHS of (A.7) and $f_g^{l*'}(p_g^l(t)) = s_l(t) + G_l(t)$ maximizes the RHS. As such, after inserting the derivative (A.4) in (A.7), we can show that the Differential Allocation-Payment Relationship holds when $\hat{\hat{\alpha}}_g^l(t) = \left(\frac{\overline{s}_l(t,t_{current})+G_l(t)}{\underline{s}_l(t,t_{current})+G_l(t)}\right) \ln\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)$. Again, because of the properties of the forecast, we know the worst forecast occurs at $t_{current} = 1$ and write: $\hat{\hat{\alpha}}_g^l(t) \geq \frac{\overline{s}_l(t,1)+G_l(t)}{\underline{s}_l(t,1)+G_l(t)} \ln\left(\frac{2R(U_g-\pi_l(t))}{L_g-\pi_l(t)}\right)$.

Now, we define $\alpha_g^l(t)$ as in (2.19) because if the lower bound solar forecast is zero,

$\underline{s}_l(t,1) = 0$, we know we are in the case where $y_g^l(t) \geq \underline{s}_l(t, t_{current})$. When the solar forecast is nonzero, $\underline{s}_l(t,1) \neq 0$, we take the maximum of $\hat{\alpha}_g^l(t)$ and $\hat{\hat{\alpha}}_g^l(t)$ to account for the worst case ratio. With $\alpha_g^l(t)$ defined as in (2.19), equation (A.7) holds for the proposed pricing function, operational cost function, and Fenchel conjugate, the remainder of the proof follows from Lemma 1 and Theorem 2 in [39]. ■

A.3 Notes on the Origin of Equation (3.11)

Notes on origin of Equation (3.11):

One of the necessary KKT conditions for x_{ns}^* to be a critical point is the following:

$$x_{ns}^* \frac{\partial h}{\partial x_{ns}} = 0, \quad \forall s \in \mathcal{S}_n, n \in \mathcal{N}$$

where $h(x,u,p_z)|_{z=e,c,d}$ is the Lagrangian function:

$$h(x,u,p_z) = \sum_{\mathcal{N},\mathcal{S}_n} v_{ns}x_{ns} - \sum_{\mathcal{N}} u_n \Big[\sum_{\mathcal{S}_n} x_{ns} - 1\Big] - \sum_{\mathcal{T}} \sum_{z=e,c,d} p_z(t) \Big[\sum_{\mathcal{N},\mathcal{S}_n} x_{ns} i_{nsz}(t) - \hat{Z}\Big]$$

and $\frac{\partial h}{\partial x_{ns}}$ is its derivative:

$$\frac{\partial h}{\partial x_{ns}} = v_{ns} - u_n - \sum_{\mathcal{T}} \sum_{z=e,c,d} p_z(t) i_{nsz}(t), \ \forall s \in \mathcal{S}_n, n \in \mathcal{N}.$$

We know if x_{ns}^* is nonzero (i.e., a schedule is allocated to user n) then $\frac{\partial h}{\partial x_{ns}}$ must be 0. Setting the $\frac{\partial h}{\partial x_{ns}} = 0$ leaves us with:

$$u_n = v_{ns} - \sum_{\mathcal{T}} \Big[i_{nse}(t)p_e(t) + i_{nsc}(t)p_c(t) - i_{nsc}(t)p_d(t)\Big] \forall s \in \mathcal{S}_n, n \in \mathcal{N}. \tag{A.6}$$

With this in mind, we can instead use equation (3.11) to calculate the utility of request

n (as we know the maximal utility schedule will always be selected).

A.4 Proof for Theorem 3.3.1

The proof of Theorem 3.3.1 requires the following Definition and Lemmas.

Definition 1. (From [39]) *The Differential Allocation-Payment Relationship for a given parameter $\alpha \geq 1$ is:*

$$\big(p(t) - f'(y(t))\big)\mathrm{d}y(t) \geq \frac{1}{\alpha(t)} f^{*'}(p(t))\mathrm{d}p(t) \tag{A.7}$$

for all $t \in [0, T]$ and for all shared resources (energy capacity, discharging power, and charging power) where $f'(y(t))$ is the derivative of an operational cost function and $f^{*'}(p(t))$ is the corresponding Fenchel conjugate's derivative.

In the following, let $i_{nsd}(t)$ be the CES *discharging power profile* for request n in feasible schedule s at time t. We note that the charging power profile $i_{nsc}(t)|_{t=1,\ldots,T}$ and the discharging power profile $i_{nsd}(t)|_{t=1,\ldots,T}$ are negatives of one another. We add this variable to separate the charging power dual variable updates from the discharging power dual variable updates for ease of exposition. Specifically, let $i_{nsd}(t) = -i_{nsc}(t), \forall t$. Additonally, we add the variable $y_d(t)$ to denote the total discharging power at time t. Specifically, let $y_d(t) = -y_c(t), \forall t$. The variable $y_d(t)$ can be calculated similarly to $y_c(t)$ in (3.4) as: $y_d(t) = \sum_{\mathcal{N},\mathcal{S}_n} i_{nsd}(t) x_{ns}$.

Lemma A.4.1 *(From [39]) If the Differential Allocation-Payment Relationship holds for $\alpha \geq 1$, then each energy storage request n and the chosen charge/discharge schedule $s_n^\star$ satisfy the following:*

$$\widetilde{p}_{ns^\star} - \sum_{t\in[t_{ns}^-,t_{ns^\star}^+]}\Big(\Delta f_e(y_e(t))^{(n,n-1)} + \Delta f_c(y_c(t))^{(n,n-1)} + \Delta f_d(y_d(t))^{(n,n-1)}\Big) \geq \frac{1}{\alpha}(D^n - D^{n-1} - u_n)$$

where

$$\Delta f_e(y_e(t))^{(n,n-1)} = f_e(y_e(t))^{(n)} - f_e(y_e(t))^{(n-1)}$$

$$\Delta f_c(y_c(t))^{(n,n-1)} = f_c(y_c(t))^{(n)} - f_c(y_c(t))^{(n-1)}$$

$$\Delta f_d(y_d(t))^{(n,n-1)} = f_d(y_d(t))^{(n)} - f_d(y_d(t))^{(n-1)}$$

$$\widetilde{p}_{ns^\star} = \sum_{\mathcal{T}}\Big[i_{nse}(t)p_e(t) + i_{nsc}(t)p_c(t) + i_{nsd}(t)p_d(t)\Big].$$

Proof of Lemma A.4.1: We expand out $D^n - D^{n-1} =$

$$u_n + \sum_{t\in[t_n^-,t_{ns^\star}^+]}\Big(\Delta f_e^*(p_e(t))^{(n,n-1)} + \Delta f_c^*(p_c(t))^{(n,n-1)} + \Delta f_d^*(p_e(t))^{(n,n-1)}\Big)$$

where

$$\Delta f_e^*(y_e(t))^{(n,n-1)} = f_e^*(y_e(t))^{(n)} - f_e^*(y_e(t))^{(n-1)}$$

$$\Delta f_c^*(y_c(t))^{(n,n-1)} = f_c^*(y_c(t))^{(n)} - f_c^*(y_c(t))^{(n-1)}$$

$$\Delta f_d^*(y_d(t))^{(n,n-1)} = f_d^*(y_d(t))^{(n)} - f_d^*(y_d(t))^{(n-1)}.$$

The lemma follows by summing the *Differential Payment-Allocation Relationship* over all shared resources (energy capacity, discharge power, and charge power) and over the entire time period. □

Lemma A.4.2 *(From [39]) If the Differential Allocation-Payment Relationship holds for* $\alpha \geq 1$ *then* $P^n - P^{n-1} \geq \frac{1}{\alpha}(D^n - D^{n-1})$ *for all* n.

Proof of Lemma A.4.2: If energy storage request n is denied for all schedules $s \in \mathcal{S}_n$, then $P^n - P^{n-1} = D^n - D^{n-1} = 0$. Otherwise, the change of the primal objective is:

$$P^n - P^{n-1} = v_{ns^\star} - \sum_{t \in [t_n^-, t_{ns^\star}^+]} \left(\Delta f_e(y_e(t))^{(n,n-1)} + \Delta f_c(y_c(t))^{(n,n-1)} + \Delta f_d(y_d(t))^{(n,n-1)} \right)$$

where $v_{ns^\star} = u_n + \widetilde{p}_{ns^\star}$. By Lemma A.4.1, we get that

$$P^n - P^{n-1} \geq u_n + \frac{1}{\alpha}(D^n - D^{n-1} - u_n).$$

With $u_n \geq 0$ and $\alpha \geq 1$, then $P^n - P^n \geq \frac{1}{\alpha}(D^n - D^{n-1})\ \forall n \in \mathcal{N}$. □

Lemma A.4.3 *(From [39]) If there is a constant* $\alpha \geq 1$ *such that the incremental increase of the primal and dual objective values differ by at most an* α *factor, i.e.,* $P^n - P^{n-1} \geq \frac{1}{\alpha}(D^n - D^{n-1})$*, for every energy storage request* n*, then the heuristic is* 2α*-competitive.*

Proof of Lemma A.4.3: Summing up the inequality at each step n, we have

$$\begin{aligned} P^N &= \sum_n (P^n - P^{n-1}) \\ &\geq \frac{1}{\alpha} \sum_n (D^n - D^{n-1}) \\ &= \frac{1}{\alpha}(D^N - D^0). \end{aligned}$$

Now, we use the fact that the initial primal value is $P^0 = 0$ and by weak duality, $D^N \geq OPT$. Next, we assume $D^0 \leq \frac{1}{2}OPT$, we have that $P^N \geq \frac{1}{2\alpha}OPT$. Thus, the online heuristic is 2α-competitive. □

Lemma A.4.4 *The online pricing heuristic* (3.12) *is α_e-competitive in welfare generated from the scheduling of energy capacity in the shared battery where*

$$\alpha_e = 2 \ln \Big(\frac{6U_e}{L_e} \Big).$$

Proof of Lemma A.4.4: We will show that the pricing heuristic in (3.12) satisfies the *Differential Payment-Allocation Relationship* in equation (A.7) with parameter α_e. Then the rest of the Lemma follows from Lemmas A.4.1, A.4.2, and A.4.3.

The scheduled energy capacity of the shared battery has no cost to the battery manager but cannot exceed the total capacity limit of the battery $\hat{E}$ (in other terms, the cost function $f_e(y_e(t))$ for the energy capacity can be seen as a zero-infinite step function with the step occurring right after $\hat{E}$). Furthermore, the pricing function (3.12) never allows $y_e(t)$ to exceed $\hat{E}$ so the derivative $f'_e(y_e(t)) = 0$ while $y_e(t) \leq \hat{E}$ (and $y_e(t) \leq \hat{E} \ \ \forall t$ due to (3.12) outputting prices too high for any user once the used battery capacity is at $\hat{E}$).

Next, the derivative of the Fenchel conjugate (3.8) for the energy capacity is as follows:

$$f_e^{*'}(p_e(t)) = \hat{E}. \tag{A.8}$$

The derivative of the proposed pricing function (3.12) is

$$\mathrm{d}p_e(t) = \Big(\frac{L_e}{6\hat{E}}\Big)\Big(\frac{6U_e}{L_e}\Big)^{\frac{y_e(t)}{\hat{E}}} \ln\Big(\frac{6U_e}{L_e}\Big)\mathrm{d}y_e(t). \tag{A.9}$$

After inserting $f_e'(y_e(t))$, $f_e^{*'}(p_e(t))$, and $\mathrm{d}p_e(t)$ in (A.7), we can show that the Differential Allocation-Payment Relationship holds when choosing $\alpha = \alpha_e = \ln\left(\frac{6U_e}{L_e}\right)$. Because (A.7) holds for the dual variable update function, cost function, and Fenchel conjugate, the remainder of the proof follows from Lemmas A.4.1, A.4.2, and A.4.3. □

Definition 2. The *Generalized Differential Allocation-Payment Relationship* for the payment and remuneration of two coupled resources (resources a and b) for a given parameter $\alpha \geq 1$ is:

$$\begin{aligned}&\big[p_a(t) - f_a'(y_a(t))\big]\mathrm{d}y_a(t) + \big[p_b(t) - f_b'(y_b(t))\big]\mathrm{d}y_b(t)\\ &\geq \frac{1}{\alpha(t)}\Big[f_a^{*'}(p_a(t))\mathrm{d}p_a(t) + f_b^{*'}(p_b(t))\mathrm{d}p_b(t)\Big]\end{aligned} \tag{A.10}$$

for all $t \in [0, T]$ where $f'(y(t))$ is the derivative of an operational cost function and $f^{*'}(p(t))$ is the corresponding Fenchel conjugate's derivative.

Lemma A.4.5 *The online pricing heuristics* (3.15) *and* (3.16) *are $\alpha_{c,d}$-competitive in welfare generated from the scheduling of charging and discharging power in the shared battery where*

$$\alpha_{c,d} = 2\ln\Big(\frac{6U_{c,d}}{L_{c,d}}\Big).$$

Proof of Lemma A.4.5: We will show that the pricing heuristics in (3.15) and (3.16) satisfy a *Generalized Differential Payment-Allocation Relationship* that handles both payments and remunerations of coupled resources such as charging and discharging power with parameter $\alpha_{c,d}$. Then the rest of the Lemma follows from Lemmas A.4.1, A.4.2, and A.4.3.

The proof follows similarly to that of Lemma A.4.4. Both the charging power and discharging power resources have zero-infinite step functions for their operational cost functions with the step occurring at the max charging power $\hat{P}_c$ and discharging power $\hat{P}_d$, respectively. Furthermore, each of the pricing functions (3.15) and (3.16) never allow $y_c(t)$ and $y_d(t)$ to exceed $\hat{P}_c$ and $\hat{P}_d$, respectively. Thus, $f_c'(y_c(t)) = 0$ and $f_c'(y_d(t)) = 0$. Next, the derivatives of the Fenchel conjugates are $f_c^{*'}(p_c(t)) = \hat{P}_c$ and $f_d^{*'}(p_d(t)) = \hat{P}_d$. The derivatives of the charging power pricing function (3.15) and the discharging power pricing function (3.16) are as follows:

$$\mathrm{d}p_c(t) = \Big(\frac{L_c}{6\hat{P}_c}\Big)\Big(\frac{6U_c}{L_c}\Big)^{\frac{y_c(t)}{\hat{P}_c}} \ln\Big(\frac{6U_c}{L_c}\Big)\mathrm{d}y_c(t),$$
$$\mathrm{d}p_d(t) = \Big(\frac{L_d}{6\hat{P}_d}\Big)\Big(\frac{6U_d}{L_d}\Big)^{\frac{y_d(t)}{\hat{P}_d}} \ln\Big(\frac{6U_d}{L_d}\Big)\mathrm{d}y_d(t).$$

After inserting $f_c'(y_c(t))$, $f_c^{*'}(p_c(t))$, $f_d'(y_d(t))$, $f_d^{*'}(p_d(t))$, $\mathrm{d}p_c(t)$, and $\mathrm{d}p_d(t)$ in (A.10), the relationship is as follows:

$$\Big(\frac{L_c}{6}\Big)\Big(\frac{6U_c}{L_c}\Big)^{\frac{y_c(t)}{\hat{P}_c}}\mathrm{d}y_c(t) + \Big(\frac{L_d}{6}\Big)\Big(\frac{6U_d}{L_d}\Big)^{\frac{y_d(t)}{\hat{P}_d}}\mathrm{d}y_d(t)$$
$$\geq \frac{1}{\alpha(t)}\left[\Big(\frac{L_c}{6}\Big)\Big(\frac{6U_c}{L_c}\Big)^{\frac{y_c(t)}{\hat{P}_c}} \ln\Big(\frac{6U_c}{L_c}\Big)\mathrm{d}y_c(t) + \Big(\frac{L_d}{6}\Big)\Big(\frac{6U_d}{L_d}\Big)^{\frac{y_d(t)}{\hat{P}_d}} \ln\Big(\frac{6U_d}{L_d}\Big)\mathrm{d}y_d(t)\right].$$

Now, using the assumption that the ratios of users' maximum valuation to minimum valuation for charging and discharging are equal, i.e., $\frac{U_c}{L_c} = \frac{U_d}{L_d} = \frac{U_{c,d}}{L_{c,d}}$, the relationship

can be simplified to:

$$\left[\left(\frac{L_c}{6}\right)\left(\frac{6U_{c,d}}{L_{c,d}}\right)^{\frac{y_c(t)}{\hat{P}_c}}\mathrm{d}y_c(t)+\left(\frac{L_d}{6}\right)\left(\frac{6U_{c,d}}{L_{c,d}}\right)^{\frac{y_d(t)}{\hat{P}_d}}\mathrm{d}y_d(t)\right]$$
$$\geq\frac{\ln\left(\frac{6U_{c,d}}{L_{c,d}}\right)}{\alpha(t)}\left[\left(\frac{L_c}{6}\right)\left(\frac{6U_{c,d}}{L_{c,d}}\right)^{\frac{y_c(t)}{\hat{P}_c}}\mathrm{d}y_c(t)+\left(\frac{L_c}{6}\right)\left(\frac{6U_{c,d}}{L_{c,d}}\right)^{\frac{y_d(t)}{\hat{P}_d}}\mathrm{d}y_d(t)\right].$$

The bracketed term that is shared on the LHS and the RHS represents the total payment and remuneration for charging/discharging at a given time t. To simplify this relationship further, there are 3 cases: 1) when the payment is greater than the remuneration and the bracketed term is positive, 2) when the payment is less than the remuneration and the bracketed term is negative, and 3) when the payment is equal to the remuneration and the bracketed term is zero.

In case 1, the relationship simplifies to $\alpha \geq \ln\left(\frac{6U_{c,d}}{L_{c,d}}\right)$. In case 2, the relationship simplifies to $\alpha \leq \ln\left(\frac{6U_{c,d}}{L_{c,d}}\right)$. In case 3, the payment and remuneration fully cancel each other. As such, the Generalized Differential Allocation-Payment Relationship holds when choosing $\alpha = \alpha_{c,d} = \ln\left(\frac{6U_{c,d}}{L_{c,d}}\right)$. Because (A.10) holds for the charging/discharging pricing functions, cost functions, and Fenchel conjugates, the remainder of the proof follows from Lemmas A.4.1, A.4.2, and A.4.3. □

A.5 Proof for Proposition 4.3.1

Proof: In [140], it is shown that with probability $1 - \delta\sqrt{2}$ the mass of the true parameter never decreases below $\pi_0(\boldsymbol{\theta}^\star)e^{-\lambda|\mathcal{P}|}$ in the prior distribution during the entire learning process. As such, the desired reliability metric on the RHS of the constraints (4.16)-(4.18), i.e., $1 - \nu$, can be selected such that the constraints *must* hold for the

true parameter. Let $\pi^{\star}_{min} = \pi_0(\boldsymbol{\theta}^{\star})e^{-\lambda|\mathcal{P}|}$ be the minimum reachable mass of the true parameter in the prior distribution. Furthermore, we abuse notation and denote $\mathbb{P}_j^{safe} = \mathbb{P}_{\{\phi_c\}_{c\in\mathcal{C}}}\left[g_j\big(\mathbf{D}_\tau(\mathbf{p}_\tau)\big) \leq 0\right]$ as the probability that constraint j is upheld. Now, assuming the aggregator only has knowledge of the true parameter given by the prior distribution π_τ on day τ, the aggregator can calculate the probability of satisfying the constraint:

$$\sum_{\hat{\boldsymbol{\theta}}\in\Theta} \pi_\tau(\hat{\boldsymbol{\theta}})(\mathbb{P}_j^{safe}|\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}). \tag{A.11}$$

This can be split into two terms for the true parameter $\boldsymbol{\theta}^{\star}$ and all other parameters $\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\star}$:

$$\pi_\tau(\boldsymbol{\theta}^{\star})(\mathbb{P}_j^{safe}|\boldsymbol{\theta} = \boldsymbol{\theta}^{\star}) + (1 - \pi_\tau(\boldsymbol{\theta}^{\star}))(\mathbb{P}_j^{safe}|\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\star}). \tag{A.12}$$

Now, we can rewrite the probability assuming that $\boldsymbol{\theta}^{\star}$ has reached the minimum mass $\pi^{\star}_{min}$ in the prior distribution:

$$\pi^{\star}_{min}(\mathbb{P}_j^{safe}|\boldsymbol{\theta} = \boldsymbol{\theta}^{\star}) + (1 - \pi^{\star}_{min})(\mathbb{P}_j^{safe}|\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\star}). \tag{A.13}$$

Recall, the aggregator wants constraint j to hold with probability at least $1 - \mu$ for the true parameter $\boldsymbol{\theta}^{\star}$, so we can replace $(\mathbb{P}_j^{safe}|\boldsymbol{\theta} = \boldsymbol{\theta}^{\star})$ with $1 - \mu$. Furthermore, $(\mathbb{P}_j^{safe}|\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\star}) \leq 1$ and we replace it accordingly yielding:

$$\pi^{\star}_{min}(1 - \mu) + (1 - \pi^{\star}_{min}). \tag{A.14}$$

Now, we want this probability to be the minimum allowable probability across the prior π for constraint j to hold so we set it equal to the reliability metric:

$$\pi^{\star}_{min}(1-\mu) + (1-\pi^{\star}_{min}) = 1-\nu, \tag{A.15}$$

which yields

$$\nu = \mu\pi^{\star}_{min}. \tag{A.16}$$

Moreover, by selecting $\nu \leq \mu\pi^{\star}_{min}$ the aggregator ensures that constraint j will be upheld with probability at least $1-\mu$ for the true parameter $\boldsymbol{\theta}^{\star}$. (i.e., the total mass of the incorrect parameters $\boldsymbol{\theta} \neq \boldsymbol{\theta}^{\star}$ in the prior distribution π_{τ} can never be large enough to satisfy the constraint's inequality without the true parameter also satisfying the constraint). ■

Bibliography

[1] N. Tucker and M. Alizadeh, *An online scheduling algorithm for a community energy storage system*, *IEEE Transactions on Smart Grid* (2022) 1–1.

[2] A. Moradipari, N. Tucker, and M. Alizadeh, *Mobility-aware electric vehicle fast charging load models with geographical price variations*, *IEEE Transactions on Transportation Electrification* **7** (2021), no. 2 554–565.

[3] N. Tucker, A. Moradipari, and M. Alizadeh, *Constrained thompson sampling for real-time electricity pricing with grid reliability constraints*, *arXiv e-print arXiv:1908.07964* (2020) [arXiv:1908.0796].

[4] N. Tucker and M. Alizadeh, *An online admission control mechanism for electric vehicles at public parking infrastructures*, *IEEE Transactions on Smart Grid* (2019) 1–1.

[5] N. Tucker, G. Cezar, and M. Alizadeh, *Real-time electric vehicle smart charging at workplaces: A real-world case study*, 2022. arXiv preprint arXiv:2203.06847.

[6] A. Moradipari, N. Tucker, T. Zhang, G. Cezar, and M. Alizadeh, *Mobility-aware smart charging of electric bus fleets*, in *2020 IEEE Power and Energy Society General Meeting (PESGM)*, pp. 1–5, 2020.

[7] B. Turan, N. Tucker, and M. Alizadeh, *Smart charging benefits in autonomous mobility on demand systems*, *arXiv preprint arXiv:1907.00106* (2019).

[8] N. Tucker, B. Turan, and M. Alizadeh, *Online charge scheduling for electric vehicles in autonomous mobility on demand fleets*, in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 226–231, 2019.

[9] N. Tucker, B. Ferguson, and M. Alizadeh, *An online pricing mechanism for electric vehicle parking assignment and charge scheduling*, *American Control Conference* (2019, *Printed.*).

[10] N. Tucker and M. Alizadeh, *Online pricing mechanisms for electric vehicle management at workplace charging facilities*, *Allerton Conference on Communication, Control, and Computing* (2018.).

[11] N. Tucker and M. Khanbaghi, *Jump linear quadratic control for energy management of a nanogrid*, in *2018 Annual American Control Conference (ACC)*, pp. 3171–3176, 2018.

[12] "Cumulative PEV Sales." `www.energy.gov/eere/vehicles/articles/fotw-1057-november-26-2018-one-million-plug-vehicles-have-been-sold-united`. *U.S. DOE*, Accessed: 20-Dec-2018.

[13] E. Veldman and R. A. Verzijlbergh, *Distribution grid impacts of smart electric vehicle charging from different perspectives*, *IEEE Transactions on Smart Grid* **6** (Jan, 2015) 333–342.

[14] W. Su, H. Eichi, W. Zeng, and M. Chow, *A survey on the electrification of transportation in a smart grid environment*, *IEEE Transactions on Industrial Informatics* **8** (Feb, 2012) 1–10.

[15] J. C. Mukherjee and A. Gupta, *A review of charge scheduling of electric vehicles in smart grid*, *IEEE Systems Journal* **9** (Dec, 2015) 1541–1553.

[16] M. A. Ortega-Vazquez, F. Bouffard, and V. Silva, *Electric vehicle aggregator/system operator coordination for charging scheduling and services procurement*, *IEEE Transactions on Power Systems* **28** (May, 2013) 1806–1815.

[17] Z. Luo, Z. Hu, Y. Song, Z. Xu, and H. Lu, *Optimal coordination of plug-in electric vehicles in power grids with cost-benefit analysis—part i: Enabling techniques*, *IEEE Transactions on Power Systems* **28** (Nov, 2013) 3546–3555.

[18] H. Zhang, Z. Hu, Z. Xu, and Y. Song, *An integrated planning framework for different types of pev charging facilities in urban area*, *IEEE Transactions on Smart Grid* **7** (Sept, 2016).

[19] N. Machiels, N. Leemput, F. Geth, J. V. Roy, J. Büscher, and J. Driesen, *Design criteria for electric vehicle fast charge infrastructure based on flemish mobility behavior*, *IEEE Transactions on Smart Grid* **5** (Jan, 2014) 320–327.

[20] G. Wang, Z. Xu, F. Wen, and K. P. Wong, *Traffic-constrained multiobjective planning of electric-vehicle charging stations*, *IEEE Transactions on Power Delivery* **28** (Oct, 2013).

[21] W. Yao, J. Zhao, F. Wen, Z. Dong, Y. Xue, Y. Xu, and K. Meng, *A multi-objective collaborative planning strategy for integrated power distribution and electric vehicle charging systems*, *IEEE Transactions on Power Systems* **29** (July, 2014) 1811–1821.

[22] A. Y. S. Lam, Y. Leung, and X. Chu, *Electric vehicle charging station placement: Formulation, complexity, and solutions*, *IEEE Transactions on Smart Grid* **5** (Nov, 2014) 2846–2856.

[23] Z. Liu, F. Wen, and G. Ledwich, *Optimal planning of electric-vehicle charging stations in distribution systems*, *IEEE Transactions on Power Delivery* **28** (Jan, 2013) 102–110.

[24] B. Ferguson, V. Nagaraj, E.C. Kara, and M. Alizadeh, *Optimal planning of workplace electric vehicle charging insfrastructure with smart charging opportunities*, .

[25] M. Gharbaoui, B. Martini, R. Bruno, L. Valcarenghi, M. Conti, and P. Castoldi, *Policies for efficient usage of an ev charging infrastructure deployed in city parking facilities*, *International Conference on ITS Telecomm* (2013) 384–389.

[26] V. Robu, E. H. Gerding, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings, *An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging*, *J. Artif. Int. Res.* **48** (Oct., 2013) 175–230.

[27] Z. Zheng and N. Shroff, *Online welfare maximization for electric vehicle charging with electricity cost*, in *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy '14, (New York, NY, USA), pp. 253–263, ACM, 2014.

[28] Q. Xiang, F. Kong, X. Liu, X. Chen, L. Kong, and L. Rao, *Auc2charge: An online auction framework for electric vehicle park-and-charge*, in *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems*, e-Energy '15.

[29] S. Stein, E. Gerding, V. Robu, and N. R. Jennings, *A model-based online mechanism with pre-commitment and its application to electric vehicle charging*, in *11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, 2012.

[30] M. Honarmand, A. Zakariazadeh, and S. Jadid, *Optimal scheduling of electric vehicles in an intelligent parking lot considering vehicle-to-grid concept and battery condition*, *Energy* **65** (2014).

[31] E. Bitar and S. Low, *Deadline differentiated pricing of deferrable electric power service*, in *IEEE 51st CDC*, pp. 4991–4997, 2012.

[32] A. Ghosh and V. Aggarwal, *Control of charging of electric vehicles through menu-based pricing*, *IEEE Transactions on Smart Grid* **9** (Nov, 2018) 5918–5929.

[33] H. Zhang, Z. Hu, Z. Xu, and Y. Song, *Optimal planning of pev charging station with single output multiple cables charging spots*, *IEEE Transactions on Smart Grid* **8** (Sept, 2017).

[34] N. R. Devanur and Z. Huang, *Primal dual gives almost optimal energy efficient online algorithms*, in *Proceedings of the 25th ACM-SIAM SODA*, SODA '14, (Philadelphia, PA, USA), 2014.

[35] W. Powell, *Approximate dynamic programming: Solving the curses of dimensionality*, 2007.

[36] D. P. Bertsekas, *Approximate dynamic programming*, 2008.

[37] M. S. Maxwell, M. Restrepo, S. G. Henderson, and H. Topaloglu, *Approximate dynamic programming for ambulance redeployment*, *INFORMS Journal on Computing* (2010).

[38] C. Novoa and R. Storer, *An approximate dynamic programming approach for the vehicle routing problem with stochastic demands*, *European Journal of Operational Research* (2009).

[39] X. Zhang, Z. Huang, C. Wu, Z. Li, and F. C. M. Lau, *Online auctions in iaas clouds: Welfare and profit maximization with server costs*, *IEEE/ACM Transactions on Networking* **25** (April, 2017).

[40] "NREL's PVWatts Calculator, [Online]." `https://pvwatts.nrel.gov/`,.

[41] M. Smith and J. Castellano, "Costs Associated with Non-Residential Electric Vehicle Supply Equipment, [Online]." `https://afdc.energy.gov/files/u/publication/evse_cost_report_2015.pdf`.

[42] M. Pavone, *Autonomous mobility-on-demand systems for future urban mobility*, . Publisher: Springer Berlin Heidelberg, doi=10.1007/978-3-662-45854-9_19.

[43] R. Zhang, K. Spieser, E. Frazzoli, and M. Pavone, *Models, algorithms, and evaluation for autonomous mobility-on-demand systems*, in *2015 American Control Conference (ACC)*, pp. 2573–2587, July, 2015.

[44] J. B. Greenblatt and S. Shaheen, *Automated vehicles, on-demand mobility, and environmental impacts*, *Current Sustainable/Renewable Energy Reports* **2** (Sep, 2015) 74–81.

[45] N. Agatz, A. Erera, M. Savelsbergh, and X. Wang, *Optimization for dynamic ride-sharing: A review*, *European Journal of Operational Research* **223** (2012), no. 2 295 – 303.

[46] M. Schreieck *et. al.*, *A matching algorithm for dynamic ridesharing*, *Transportation Research Procedia* **19** (2016) 272 – 285.

[47] A. Kleiner, B. Nebel, and V. A. Ziparo, *A mechanism for dynamic ride sharing based on parallel auctions*, in *Proceedings of the 22 International Joint Conference on Artificial Intelligence*, IJCAI'11, pp. 266–272, 2011.

[48] Alonso-Mora *et. al.*, *On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment*, *Proceedings of the National Academy of Sciences* **114** (2017), no. 3 462–467.

[49] N. Masoud and R. Jayakrishnan, *A decomposition algorithm to solve the multi-hop peer-to-peer ride-matching problem*, *Transportation Research Part B: Methodological* **99** (2017) 1 – 29.

[50] X. Tan and A. Leon-Garcia, *Autonomous mobility and energy service management in future smart cities: An overview*, *2018 International Conference on Universal Village* (Oct, 2018).

[51] R. Zhang, F. Rossi, and M. Pavone, *Model predictive control of autonomous mobility-on-demand systems*, in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May, 2016.

[52] T. D. Chen, K. M. Kockelman, and J. P. Hanna, *Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle and charging infrastructure decisions*, *Transportation Research Part A: Policy and Practice* **94** (2016) 243 – 254.

[53] T. D. Chen and K. M. Kockelman, *Management of a shared autonomous electric vehicle fleet: Implications of pricing schemes*, *Transportation Research Record* **2572** (2016), no. 1 37–46.

[54] M. Salazar, M. Tsao, I. Aguiar, M. Schiffer, and M. Pavone, *A congestion-aware routing scheme for autonomous mobility-on-demand systems*, in *European Control Conference*, 2019.

[55] M. Tsao, D. Milojevic, C. Ruch, M. Salazar Villalon, E. Frazzoli, and M. Pavone, *Model predictive control of ride-sharing autonomous mobility-on-demand systems*, IEEE, 2019-05. 2019 IEEE International Conference on Robotics and Automation (ICRA); Conference Location: Montreal, Canada; Conference Date: May 20-24, 2019.

[56] F. Rossi, R. Iglesias, M. Alizadeh, and M. Pavone, *On the interaction between autonomous mobility-on-demand systems and the power network: models and coordination algorithms*, *Robotics: Science and Systems XIV* (Jun, 2018).

[57] M. Salazar, F. Rossi, M. Schiffer, C. H. Onder, and M. Pavone, *On the interaction between autonomous mobility-on-demand and public transportation systems*, in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2262–2269, Nov, 2018.

[58] "California ISO OASIS Locational Marginal Prices (LMP), [Online]." `http://oasis.caiso.com/mrioasis/logon.do`.

[59] "California ISO Supply and Renewables, [Online]." http://www.caiso.com/TodaysOutlook/Pages/supply.aspx.

[60] "San Jose Heatmap, SherpaShare Rideshare Analysis, [Online]." https://www.sherpashare.com/heatmap.

[61] H. Ibrahim, A. Ilinca, and J. Perron, *Energy storage systems—characteristics and comparisons*, *Renewable and sustainable energy reviews* **12** (2008), no. 5 1221–1250.

[62] R. A. Huggins, *Energy storage*, vol. 406. Springer, 2010.

[63] Y. Brunet, *Energy storage*. John Wiley & Sons, 2013.

[64] M. Katsanevakis, R. A. Stewart, and J. Lu, *Aggregated applications and benefits of energy storage systems with application-specific control methods: A review*, *Renewable and Sustainable Energy Reviews* **75** (2017) 719–741.

[65] R. Dai, R. Esmaeilbeigi, and H. Charkhgard, *The utilization of shared energy storage in energy systems: A comprehensive review*, *IEEE Transactions on Smart Grid* **12** (2021), no. 4 3163–3174.

[66] A. Paudel, K. Chaudhari, C. Long, and H. B. Gooi, *Peer-to-peer energy trading in a prosumer-based community microgrid: A game-theoretic model*, *IEEE Transactions on Industrial Electronics* **66** (2019), no. 8 6087–6097.

[67] D.-H. Park, J.-B. Park, K. Y. Lee, S.-Y. Son, and J. H. Roh, *A bidding-based peer-to-peer energy transaction model considering the green energy preference in virtual energy community*, *IEEE Access* **9** (2021) 87410–87419.

[68] A. S. Yahaya, N. Javaid, M. U. Javed, M. Shafiq, W. Z. Khan, and M. Y. Aalsalem, *Blockchain-based energy trading and load balancing using contract theory and reputation in a smart community*, *IEEE Access* **8** (2020) 222168–222186.

[69] M. Afzal, Q. Huang, W. Amin, K. Umer, A. Raza, and M. Naeem, *Blockchain enabled distributed demand side management in community energy system with smart homes*, *IEEE Access* **8** (2020) 37428–37439.

[70] L. Ma, N. Liu, J. Zhang, and L. Wang, *Real-time rolling horizon energy management for the energy-hub-coordinated prosumer community from a cooperative perspective*, *IEEE Transactions on Power Systems* **34** (2019), no. 2 1227–1242.

[71] N. Vespermann, T. Hamacher, and J. Kazempour, *Access economy for storage in energy communities*, *IEEE Transactions on Power Systems* **36** (2021), no. 3 2234–2250.

[72] D. Parra, M. Swierczynski, D. I. Stroe, S. A. Norman, A. Abdon, J. Worlitschek, T. O'Doherty, L. Rodrigues, M. Gillott, X. Zhang, *et. al.*, *An interdisciplinary review of energy storage for communities: Challenges and perspectives*, *Renewable and Sustainable Energy Reviews* **79** (2017) 730–749.

[73] S. J. Klein and S. Coffey, *Building a sustainable energy future, one community at a time*, *Renewable and Sustainable Energy Reviews* **60** (2016) 867–880.

[74] D. Kalathil, C. Wu, K. Poolla, and P. Varaiya, *The sharing economy for the electricity storage*, *IEEE Transactions on Smart Grid* **10** (2019), no. 1 556–567.

[75] N. Vespermann, T. Hamacher, and J. Kazempour, *Access economy for storage in energy communities*, *IEEE Transactions on Power Systems* **36** (2021), no. 3 2234–2250.

[76] D. Parra, S. A. Norman, G. S. Walker, and M. Gillott, *Optimum community energy storage system for demand load shifting*, *Applied Energy* **174** (2016) 130–143.

[77] T. Terlouw, T. AlSkaif, C. Bauer, and W. Van Sark, *Multi-objective optimization of energy arbitrage in community energy storage systems using different battery technologies*, *Applied energy* **239** (2019) 356–372.

[78] S. Van Der Stelt, T. AlSkaif, and W. van Sark, *Techno-economic analysis of household and community energy storage for residential prosumers with smart appliances*, *Applied Energy* **209** (2018) 266–276.

[79] R. Dai, R. Esmaeilbeigi, and H. Charkhgard, *The utilization of shared energy storage in energy systems: a comprehensive review*, *IEEE Transactions on Smart Grid* (2021).

[80] I. Lampropoulos, T. Alskaif, W. Schram, E. Bontekoe, S. Coccato, and W. Van Sark, *Review of energy in the built environment*, *Smart Cities* **3** (2020), no. 2 248–288.

[81] N. Komodakis and J.-C. Pesquet, *Playing with duality: An overview of recent primal? dual approaches for solving large-scale optimization problems*, *IEEE Signal Processing Magazine* **32** (2015), no. 6 31–54.

[82] J. Yao and P. Venkitasubramaniam, *Stochastic games of end-user energy storage sharing*, in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 4965–4972, IEEE, 2016.

[83] Y. Yang, G. Hu, and C. J. Spanos, *Optimal sharing and fair cost allocation of community energy storage*, 2021.

[84] S. Cui, Y.-W. Wang, Y. Shi, and J.-W. Xiao, *Community energy cooperation with the presence of cheating behaviors*, *IEEE Transactions on Smart Grid* **12** (2021), no. 1 561–573.

[85] J. Yao and P. Venkitasubramaniam, *Optimal end user energy storage sharing in demand response*, in *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 175–180, IEEE, 2015.

[86] E. Oh, *Reinforcement-learning-based virtual energy storage system operation strategy for wind power forecast uncertainty management*, *Applied Sciences* **10** (2020), no. 18 6420.

[87] A. Walker and S. Kwon, *Design of structured control policy for shared energy storage in residential community: A stochastic optimization approach*, *Applied Energy* **298** (2021) 117182.

[88] W. Tushar, B. Chai, C. Yuen, S. Huang, D. B. Smith, H. V. Poor, and Z. Yang, *Energy storage sharing in smart grid: A modified auction-based approach*, *IEEE Transactions on Smart Grid* **7** (2016), no. 3 1462–1475.

[89] H. Chen, Y. Yu, Z. Hu, H. Luo, C.-W. Tan, and R. Rajagopal, *Energy storage sharing strategy in distribution networks using bi-level optimization approach*, in *2017 IEEE Power & Energy Society General Meeting*, pp. 1–5, IEEE, 2017.

[90] J. Liu, N. Zhang, C. Kang, D. S. Kirschen, and Q. Xia, *Decision-making models for the participants in cloud energy storage*, *IEEE Transactions on Smart Grid* **9** (2017), no. 6 5512–5521.

[91] X. Tan, A. Leon-Garcia, Y. Wu, and D. H. K. Tsang, *Posted-price retailing of transactive energy: An optimal online mechanism without prediction*, *IEEE Journal on Selected Areas in Communications* **38** (2020), no. 1 5–16.

[92] W. Zhong, K. Xie, Y. Liu, C. Yang, and S. Xie, *Multi-resource allocation of shared energy storage: A distributed combinatorial auction approach*, *IEEE Transactions on Smart Grid* **11** (2020), no. 5 4105–4115.

[93] D. Zhao, H. Wang, J. Huang, and X. Lin, *Virtual energy storage sharing and capacity allocation*, *IEEE transactions on smart grid* **11** (2019), no. 2 1112–1123.

[94] N. Buchbinder and J. Naor, *Online primal-dual algorithms for covering and packing*, *Mathematics of Operations Research* **34** (2009), no. 2 270–286.

[95] N. Bansal, N. Buchbinder, and J. Naor, *Randomized competitive algorithms for generalized caching*, in *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pp. 235–244, 2008.

[96] N. Buchbinder, K. Jain, and J. S. Naor, *Online primal-dual algorithms for maximizing ad-auctions revenue*, in *European Symposium on Algorithms*, pp. 253–264, Springer, 2007.

[97] Z. Huang, *Sigact news online algorithms column 25: Online primal dual: Beyond linear programs*, *ACM SIGACT News* **45** (2014), no. 4 105–119.

[98] S. Anand, N. Garg, and A. Kumar, *Resource augmentation for weighted flow-time explained by dual fitting*, in *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pp. 1228–1241, SIAM, 2012.

[99] A. Gupta, R. Krishnaswamy, and K. Pruhs, *Online primal-dual for non-linear optimization with applications to speed scaling*, in *International Workshop on Approximation and Online Algorithms*, pp. 173–186, Springer, 2012.

[100] K. T. Nguyen, *Lagrangian duality in online scheduling with resource augmentation and speed scaling*, in *European Symposium on Algorithms*, pp. 755–766, Springer, 2013.

[101] N. R. Devanur and Z. Huang, *Primal dual gives almost optimal energy-efficient online algorithms*, *ACM Transactions on Algorithms (TALG)* **14** (2017), no. 1 1–30.

[102] Commercial and Residential Hourly Load Profiles for all TMY3 Locations in the United States [data set]. NREL. 2014. Retrieved from https://dx.doi.org/10.25984/1788456.

[103] Electric Schedule E-19. Medium General Demand-Metered TOU Service. Pacific Gas and Electric Company. Retrieved from https://www.pge.com/tariffs/assets/pdf/tariffbook/ELEC_SCHEDS_E-19.pdf.

[104] C. Eid, E. Koliou, M. Valles, J. Reneses, and R. Hakvoort, *Time-based pricing and electricity demand response: Existing barriers and next steps*, *Utilities Policy* **40** (2016) 15 – 25.

[105] V. Gomez, M. Chertkov, S. Backhaus, and H. J. Kappen, *Learning price-elasticity of smart consumers in power distribution systems*, in *2012 SmartGridComm*, pp. 647–652, IEEE, 2012.

[106] Z. Xu, T. Deng, Z. Hu, Y. Song, and J. Wang, *Data-driven pricing strategy for demand-side resource aggregators*, *IEEE Transactions on Smart Grid* **9** (2016), no. 1 57–66.

[107] P. Li and B. Zhang, *Linear estimation of treatment effects in demand response: An experimental design approach*, *arXiv preprint arXiv:1706.09835* (2017).

[108] P. Li, H. Wang, and B. Zhang, *A distributed online pricing strategy for demand response programs*, *IEEE Transactions on Smart Grid* **10** (2017), no. 1 350–360.

[109] K. Khezeli, W. Lin, and E. Bitar, *Learning to buy (and sell) demand response*, *IFAC-PapersOnLine* **50** (2017), no. 1 6761–6767.

[110] L. Jia, L. Tong, and Q. Zhao, *An online learning approach to dynamic pricing for demand response*, *arXiv preprint arXiv:1404.1325* (2014).

[111] K. Khezeli and E. Bitar, *Risk-sensitive learning and pricing for demand response*, *IEEE Transactions on Smart Grid* **9** (2017), no. 6 6000–6007.

[112] L. Jia, Q. Zhao, and L. Tong, *Retail pricing for stochastic demand with unknown parameters: An online machine learning approach*, in *2013 51st Allerton*, pp. 1353–1358, IEEE, 2013.

[113] D. Zhou, M. Balandat, and C. Tomlin, *Residential demand response targeting using machine learning with observational data*, in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 6663–6668, IEEE, 2016.

[114] Y. Li, Q. Hu, and N. Li, *Learning and selecting the right customers for reliability: A multi-armed bandit approach*, in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 4869–4874, IEEE, 2018.

[115] D. Kalathil and R. Rajagopal, *Online learning for demand response*, in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 218–222, IEEE, 2015.

[116] E. Dall'Anese, K. Baker, and T. Summers, *Chance-constrained ac optimal power flow for distribution systems with renewables*, *IEEE Transactions on Power Systems* **32** (2017), no. 5 3427–3438.

[117] R. Mieth and Y. Dvorkin, *Data-driven distributionally robust optimal power flow for distribution systems*, *IEEE Control Systems Letters* **2** (2018), no. 3 363–368.

[118] R. Mieth and Y. Dvorkin, *Online learning for network constrained demand response pricing in distribution systems*, *arXiv:1811.09384* (2018).

[119] Yi Wang, Qixin Chen, Chongqing Kang, Mingming Zhang, Ke Wang, and Yun Zhao, *Load profiling and its application to demand response: A review*, *Tsinghua Science and Technology* **20** (April, 2015) 117–129.

[120] J. L. Mathieu, S. Koch, and D. S. Callaway, *State estimation and control of electric loads to manage real-time energy imbalance*, *IEEE Transactions on Power Systems* **28** (2012), no. 1 430–440.

[121] J. M. Foster and M. C. Caramanis, *Optimal power market participation of plug-in electric vehicles pooled by distribution feeder*, *IEEE Transactions on power systems* **28** (2013), no. 3 2065–2076.

[122] S. Lin, F. Li, E. Tian, Y. Fu, and D. Li, *Clustering load profiles for demand response applications*, *IEEE Transactions on Smart Grid* **10** (March, 2019) 1599–1607.

[123] W. Zhang, J. Lian, C. Chang, and K. Kalsi, *Aggregated modeling and control of air conditioning loads for demand response*, *IEEE Transactions on Power Systems* **28** (Nov, 2013) 4655–4664.

[124] C. Chelmis, J. Kolte, and V. K. Prasanna, *Big data analytics for demand response: Clustering over space and time*, in *2015 IEEE International Conference on Big Data (Big Data)*, pp. 2223–2232, Oct, 2015.

[125] Y. Zhang, W. Chen, R. Xu, and J. Black, *A cluster-based method for calculating baselines for residential loads*, *IEEE Transactions on Smart Grid* **7** (Sep., 2016) 2368–2377.

[126] J. Kwac, J. Flora, and R. Rajagopal, *Household energy consumption segmentation using hourly data*, *IEEE Transactions on Smart Grid* **5** (2014), no. 1 420–430.

[127] M. Alizadeh, A. Scaglione, A. Applebaum, G. Kesidis, and K. Levitt, *Reduced-order load models for large populations of flexible appliances*, *IEEE Transactions on Power Systems* **30** (2015), no. 4.

[128] M. Alizadeh, A. Scaglione, R. J. Thomas, and D. Callaway, *Information infrastructure for cellular load management in green power delivery systems*, in *2011 SmartGridComm*, pp. 13–18, IEEE.

[129] T.-H. Chang, M. Alizadeh, and A. Scaglione, *Coordinated home energy management for real-time power balancing*, in *2012 IEEE Power and Energy Society General Meeting*, pp. 1–8, IEEE, 2012.

[130] M. Alizadeh and A. Scaglione, *Least laxity first scheduling of thermostatically controlled loads for regulation services*, in *2013 IEEE GlobalSIP*, pp. 503–506, IEEE, 2013.

[131] A. Krishnamurthy, Z. S. Wu, and V. Syrgkanis, *Semiparametric contextual bandits*, *arXiv preprint arXiv:1803.04204* (2018).

[132] D. J. Foster, A. Agarwal, M. Dudík, H. Luo, and R. E. Schapire, *Practical contextual bandits with regression oracles*, *arXiv preprint arXiv:1803.01088* (2018).

[133] T. Xu, Y. Yu, J. Turner, and A. Regan, *Thompson sampling in dynamic systems for contextual bandit problems*, *arXiv preprint arXiv:1310.5008* (2013).

[134] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, Z. Wen, *et. al.*, *A tutorial on thompson sampling*, *Foundations and Trends® in Machine Learning* **11** (2018), no. 1 1–96.

[135] D. Russo and B. Van Roy, *Learning to optimize via posterior sampling*, *Mathematics of Operations Research* **39** (2014), no. 4 1221–1243.

[136] S. Agrawal and N. Goyal, *Analysis of thompson sampling for the multi-armed bandit problem*, in *Conference on Learning Theory*, pp. 39–1, 2012.

[137] P. Auer, *Using confidence bounds for exploitation-exploration trade-offs*, *Journal of Machine Learning Research* **3** (2002), no. Nov 397–422.

[138] V. Dani, T. P. Hayes, and S. M. Kakade, *Stochastic linear optimization under bandit feedback*, .

[139] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, *Improved algorithms for linear stochastic bandits*, in *Advances in Neural Information Processing Systems*, pp. 2312–2320, 2011.

[140] A. Gopalan, S. Mannor, and Y. Mansour, *Thompson sampling for complex online problems*, in *International Conference on Machine Learning*, pp. 100–108, 2014.

[141] S. Amani, M. Alizadeh, and C. Thrampoulidis, *Linear stochastic bandits under safety constraints*, in *Advances in Neural Information Processing Systems*, pp. 9252–9262, 2019.

[142] A. Moradipari, S. Amani, M. Alizadeh, and C. Thrampoulidis, *Safe linear thompson sampling*, *arXiv preprint arXiv:1911.02156* (2019).

[143] V. Saxena, J. Jaldén, J. E. Gonzalez, I. Stoica, and H. Tullberg, *Constrained thompson sampling for wireless link optimization*, *arXiv preprint arXiv:1902.11102* (2019).

[144] A. Moradipari, C. Silva, and M. Alizadeh, *Learning to dynamically price electricity demand based on multi-armed bandits*, in *2018 IEEE GlobalSIP*, pp. 917–921, Nov, 2018.

[145] P. Andrianesis, M. Caramanis, R. Masiello, R. Tabors, and S. Bahramirad, *Locational marginal value of distributed energy resources as non-wires alternatives*, *IEEE Transactions on Smart Grid* (2019).

[146] M. E. Baran and F. F. Wu, *Optimal capacitor placement on radial distribution systems*, *IEEE Transactions on Power Delivery* **4** (Jan, 1989) 725–734.

[147] M. D. Sankur, R. Dobbe, E. Stewart, D. S. Callaway, and D. B. Arnold, *A linearized power flow model for optimization in unbalanced distribution systems*, *arXiv preprint arXiv:1606.04492* (2016).

[148] H. J. Liu, *Decentralized optimization approach for power distribution network and microgrid controls*. PhD thesis, University of Illinois at Urbana-Champaign, 2017.

[149] H. Zhu and H. J. Liu, *Fast local voltage control under limited reactive power: Optimality and stability analysis*, *IEEE Transactions on Power Systems* **31** (2016), no. 5 3794–3803.

[150] P. Šulc, S. Backhaus, and M. Chertkov, *Optimal distributed control of reactive power via the alternating direction method of multipliers*, *IEEE Transactions on Energy Conversion* **29** (2014), no. 4.

[151] M. Farivar, L. Chen, and S. Low, *Equilibrium and dynamics of local voltage control in distribution systems*, in *52nd IEEE Conference on Decision and Control*, pp. 4329–4334, IEEE, 2013.

[152] M. Lubin, Y. Dvorkin, and L. Roald, *Chance constraints for improving the security of ac optimal power flow*, *IEEE Transactions on Power Systems* **34** (2019), no. 3 1908–1917.

[153] R. Mieth and Y. Dvorkin, *Distribution electricity pricing under uncertainty*, *IEEE Transactions on Power Systems* (2019) 1–1.

[154] J. Horrox and M. Casale, *Electric Buses: Lessons from Cities Pioneering Clean Transportation*, October 2019. U.S. PIRG Education Fund, Environment America Research & Policy Center, Frontier Group.

[155] "Greenhouse gases, Regulated Emissions, and Energy use in Transportation (GREET) Model." Argonne National Laboratory. https://greet.es.anl.gov/.

[156] *California transitioning to all-electric public bus fleet by 2040*, December 2018. California Air Resources Board. https://ww2.arb.ca.gov/.

[157] "Battery Bus Range - It's All in the Math." Mass Transit. https://www.masstransitmag.com/bus/article/12131451/battery-bus-range-its-all-in-the-math.

[158] M. Rogge, E. van der Hurk, A. Larsen, and D. U. Sauer, *Electric bus fleet size and mix problem with optimization of charging infrastructure*, *Applied Energy* **211** (2018) 282 – 295.

[159] A. Kunith, R. Mendelevitch, and D. Goehlich, *Electrification of a city bus network—an optimization model for cost-effective placing of charging infrastructure and battery sizing of fast-charging electric bus systems*, *International Journal of Sustainable Transportation* **11** (2017), no. 10 707–720, [https://doi.org/10.1080/15568318.2017.1310962].

[160] A. Houbbadi, R. Trigui, S. Pelissier, E. Redondo-Iglesias, and T. Bouton, *Optimal scheduling to manage an electric bus fleet overnight charging*, *Energies* **12** (2019), no. 14 2727.

[161] N. Qin, A. Gusrialdi, R. P. Brooker, T. Ali, *et. al.*, *Numerical analysis of electric bus fast charging strategies for demand charge reduction*, *Transportation Research Part A: Policy and Practice* **94** (2016) 386–396.

[162] Y. Wang, Y. Huang, J. Xu, and N. Barclay, *Optimal recharging scheduling for urban electric buses: A case study in davis*, *Transportation Research Part E: Logistics and Transportation Review* **100** (2017) 115–132.

[163] H. Chen, Z. Hu, Z. Xu, J. Li, H. Zhang, X. Xia, K. Ning, and M. Peng, *Coordinated charging strategies for electric bus fast charging stations*, in *2016 IEEE PES Asia-Pacific Power and Energy Engineering Conference (APPEEC)*, pp. 1174–1179, IEEE, 2016.

[164] T. Paul and H. Yamada, *Operation and charging scheduling of electric buses in a city bus route network*, in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pp. 2780–2786, Oct, 2014.

[165] M. Janovec and M. Kohání, *Exact approach to the electric bus fleet scheduling*, *Transportation Research Procedia* **40** (2019) 1380–1387.

[166] B. Kellison, "Workplace EV Charging is on the Rise." Wood Mackenzie, https://greentechmedia.com.

[167] Z. Ji, X. Huang, Z. Zhang, M. Jiang, and Q. Xu, *Evaluating the vehicle-to-grid potentials by electric vehicles: A quantitative study in china by 2030*, in *2020 IEEE PESGM*, pp. 1–5, 2020.

[168] B. Ferguson, V. Nagaraj, E. C. Kara, and M. Alizadeh, *Optimal planning of workplace electric vehicle charging infrastructure with smart charging opportunities*, in *2018 21st ITSC*, IEEE, 2018.

[169] P.-Y. Kong and G. K. Karagiannidis, *Charging schemes for plug-in hybrid electric vehicles in smart grid: A survey*, *IEEE Access* **4** (2016) 6846–6875.

[170] J. C. Mukherjee and A. Gupta, *A review of charge scheduling of electric vehicles in smart grid*, *IEEE Systems Journal* **9** (2015), no. 4.

[171] L. Richard and M. Petit, *Fast charging station with battery storage system for ev: Optimal integration into the grid*, in *2018 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–5, 2018.

[172] S. Wang, S. Bi, and Y.-J. A. Zhang, *A reinforcement learning approach for ev charging station dynamic pricing and scheduling control*, in *2018 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–5, 2018.

[173] B. Khaki, Y.-W. Chung, C. Chu, and R. Gadh, *Nonparametric user behavior prediction for distributed ev charging scheduling*, in *2018 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–5, 2018.

[174] Y. Xiong, B. Wang, C.-C. Chu, and R. Gadh, *Electric vehicle driver clustering using statistical model and machine learning*, in *2018 IEEE Power Energy Society General Meeting (PESGM)*, pp. 1–5, 2018.

[175] S. Powell, E. C. Kara, R. Sevlian, G. V. Cezar, S. Kiliccote, and R. Rajagopal, *Controlled workplace charging of electric vehicles: The impact of rate schedules on transformer aging*, *Applied Energy* (2020).