Czech Technical University in Prague
Faculty of Electrical Engineering

Department of Computer Science and Engineering

*Doctoral Thesis*

Jan Mrkos

# Dynamic Pricing of Electric Vehicle Charging using Markov Decision Processes

Ph.D. programme & Branch of study:

| | |
|---|---|
| (P2612) | Electrical Engineering and Information Technology |
| (2612V025) | Information Science and Computer Engineering |

Supervisor:

doc. Ing. Jiří Vokřínek, Ph.D.

Supervisor-Specialists:

doc. Ing. Michal Jakob, Ph.D.

Ing. Antonín Komenda, Ph.D.

February 2024

*Dedicated to my family, especially my wife. Thank you for your support and patience.*

# Acknowledgments

I would like to thank all my colleagues and co-authors who made this thesis possible. Special thanks belong to my supervisors Antonín Komenda, Jiří Vokřínek and Michal Jakob for their support and guidance throughout the years I spent as a Ph.D. student. I am also grateful to my colleagues from the smart mobility group, namely David Fiedler, Marek Cuchý and Martin Schaefer for their help, support and inspiring discussions. Furthermore, I would like to thank Karel Horák for inspiration and the LaTeX template, and Kateřina Helisová for her time and inspiring discussions.

# Declaration

I declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institution of tertiary education. Information derived from the published work of others has been acknowledged in the text, and a list of references is given.

In February 2024

# Abstract

We are currently witnessing a global energy shift from fossil fuels to renewable sources, with the electrification of transportation playing a crucial role in this transition. The simultaneous increase in electric vehicles (EVs) and the transition towards renewable energy sources are straining the electrical grid. Furthermore, the existing charging infrastructure and its associated business models are ill-prepared for the anticipated surge in EV adoption. Additionally, the adoption of EVs is impeded by the inconvenience and unpredictability of the charging process. The method of dynamic pricing for EV charging, as presented in this thesis, offers a promising strategy to manage the growing demand for EV charging and the consequent stress on the electrical grid. This approach also provides charging station operators with a more sustainable business model and enhances travel reliability for EV drivers.

Initially, we conduct a simulation study of various EV charging reservation products to identify those most beneficial for charging station operators and EV drivers. Our findings suggest that the greatest benefits for drivers stem from modifying their charging behavior by planning charging well in advance. Charging station operators can foster this planning behavior and boost their revenues by offering advanced reservations for charging sessions. This strategy proves particularly effective for long-distance travel requiring fast charging. Consequently, the rest of the thesis concentrates on the issue of dynamically pricing pre-planned EV charging sessions.

We begin by examining the dynamic pricing from the perspective of the charging station operator, who aims to maximize revenue. Through analyzing historical sales data, the operator can develop stochastic models of demand and customer behavior. Using these models, dynamic pricing can then framed as a sequential decision-making problem under uncertainty, formulated as a Markov Decision Process (MDP). In this MDP model, we apply several simplifications to render the problem tractable. For instance, we disregard the influence of a charging session reserved in the morning on the prices in the afternoon. Moreover, we incorporate a cyclical structure into the MDP state space to reduce the size of both the action and state spaces. These simplifications let us optimally solve this MDP model, demonstrating that dynamic pricing can surpass a flat-rate pricing baseline in terms of revenue and lead to a more efficient resource allocation.

Next, we address the limitations of the simplified model by considering pricing comprehensively, accounting for potential overlaps among all charging requests, both past and future. Additionally, to enhance the realism of the demand model, we replace the cyclical structure, which constrains its accuracy, with a Poisson process. This results in more general and realistic model, within which we also quantify the inherent approximation error of the discrete MDP. Given the new model's increased complexity, we develop a heuristic Monte Carlo tree search algorithm to solve it. By evaluating the solution's performance against multiple baselines, we demonstrate that our dynamic pricing method can efficiently and effectively solve realistically-sized charging station scenarios.

**Keywords** dynamic pricing, Markov decision process, electric vehicle charging

# Abstrakt

V současné době jsme svědky celosvětového přechodu od fosilních paliv k obnovitelným zdrojům energie, přičemž klíčovou roli v tomto přechodu hraje elektrifikace dopravy. Současný nárůst počtu elektrických vozidel (EV) a přechod na obnovitelné zdroje energie zatěžují elektrickou síť. Stávající nabíjecí infrastruktura a s ní spojené obchodní modely jsou navíc na očekávaný prudký nárůst počtu EV špatně připraveny. Rozšíření EV navíc brání nepohodlnost a nepředvídatelnost procesu nabíjení. Metoda dynamické tvorby cen nabíjení EV, kterou představujeme v této práci, nabízí slibnou strategii pro zvládnutí rostoucí poptávky po nabíjení EV a následného zatížení elektrické sítě. Tento přístup také poskytuje provozovatelům dobíjecích stanic udržitelnější obchodní model a zvyšuje spolehlivost cestování pro řidiče EV.

V této práci nejprve prezentujeme simulační studii různých forem rezervací nabíjení EV, která identifihuje ty jež jsou pro provozovatele nabíjecích stanic a řidiče EV nejvhodnější. Naše výsledky naznačují, že největší přínosy řidičům plynou z včasného plánování nabíjení předem. Provozovatelé dobíjecích stanic mohou toto plánování podpořit a zvýšit své příjmy tím, že nabídnou předběžné rezervace dobíjení. Tato strategie má nejlepší výsledky při cestách na dlouhé vzdálenosti vyžadujících rychlé nabíjení. Zbytek práce se tak soustředí na problematiku dynamické tvorby cen předem naplánovaného nabíjení EV.

Naše zkoumání dynamické tvorby cen nabíjení EV začínáme z pohledu provozovatele dobíjecí stanice, jehož cílem je maximalizace příjmů. Prostřednictvím analýzy historických dat o tržbách může provozovatel vytvořit stochastické modely poptávky a chování zákazníků. Pomocí těchto modelů lze pak dynamickou tvorbu cen rámovat jako problém sekvenčního rozhodování v neurčitosti, který formulujeme jako Markovský rozhodovací proces (MDP). V tomto MDP modelu používáme několik zjednodušení abychom ho dokázali vyřešit. Například, nezohledňujeme vliv dopoledních rezervací nabíjení na ceny v odpoledních hodinách. Navíc do stavového prostoru MDP začleňujeme cyklickou strukturu, abychom zmenšili velikost akčního i stavového prostoru. Tato zjednodušení nám umožňují optimálně vyřešit tento MDP model a ukazují, že dynamická tvorba cen může z hlediska příjmů překonat řešení s paušálními cenami a vést k efektivnějšímu rozdělování zdrojů.

V práci dále řešíme omezení zjednodušeného modelu, a to tím, že uvažujeme o tvorbě cen nabíjení EV komplexně a zohledňujeme potenciální překryvy mezi všemi rezervacemi nabíjení EV, a to jak minulými, tak budoucími. Navíc, abychom zvýšili realističnost modelu poptávky, nahrazujeme cyklickou strukturu, která omezuje jeho přesnost, Poissonovým procesem. Výsledkem je obecnější a realističtější model, v jehož rámci také kvantifikujeme chybu aproximace způsobenou diskretizací MDP. Vzhledem ke zvýšené složitosti nového modelu vyvíjíme k jeho řešení heuristický algoritmus na základě stromových Monte Carlo metod. Vyhodnocením výkonnosti našeho řešení ve srovnání s několika výchozími modely ukazujeme, že naše metoda dynamické tvorby cen dokáže účinně a efektivně řešit scénáře nabíjecích stanic realistické velikosti.

Klíčová slova dynamická tvorba cen, markovské rozhodovací procesy, nabíjení elektromobilů

# Contents

# Acronyms

**CS** Charging Station

**CSO** Charging Station Operator

**DSO** Distribution System Operator

**EMSP** Electric Mobility Service Provider

**ES** Energy Supplier

**EV** Electric Vehicle

**MCTS** Monte-Carlo tree search

**MDP** Markov Decision Process

**PDF** Probability Density Function

**PMF** Probability Mass Function

**VI** Value Iteration

# CHAPTER 1

# Introduction

After more than one hundred years of niche use, Electric Vehicles (EVs) seem on the cusp of displacing internal combustion engine vehicles in personal transportation. Improved fuel efficiency, environmental friendliness, and decreasing costs provide EVs with a competitive advantage. According to a Deloitte analysis, EVs have already reached cost parity with combustion engine vehicles in terms of total ownership cost, factoring in current subsidies [Ham+20]. These cost are anticipated to decrease further as the technology matures and automakers expand their EV offerings.

Therefore, the electrification of personal transportation is underway. However, hand in hand with the clear benefits of widespread deployment of EVs come many challenges. One of the most pressing problems is how to efficiently and cheaply distribute the energy from often unstable renewable sources to EVs. For cities and electric utilities, the widespread use of EVs may require large investments into infrastructure because large numbers of EVs could increase the peak load on the grid up to threefold [Fai09]. For EV drivers, the problem is the availability of well-maintained charging stations.

Dynamic pricing can serve as a method of load balancing, shifting peak loads to off-peak periods, thereby preventing infrastructure costs from growing. With the correct setup, this approach could reduce grid capacity requirements and necessary investments. Furthermore, dynamic pricing has the potential to boost revenues for charging stations, prompting increased investment in the development and maintenance of charging infrastructure.

In this thesis, we develop a dynamic pricing of reserved EV charging capacity using Markov Decision Process (MDP)s. The model envisions a future EV charging system where users can efficiently plan their route, balance trade-offs between charging speed and cost, and reserve a charging station in advance. The reservation would be for a future charging session and would remove uncertainty from EV user's trip. The price for each session would be determined based on anticipated demand, the condition of the charging station, and signals from energy distributors and suppliers. Varied prices across charging stations would incentivize users to select routes with optimal charging times and locations. Conflicting objectives among energy suppliers, EV drivers, grid operators and charging station operators would be resolved through the price and contracts. The charging station operators would utilize a MDP model to set prices that maximize charging station's goal, such as profitability.

Dynamic pricing is not always positively received by customers due to its potential to introduce uncertainty and complexity into the purchasing process. However, integrating a route planner with the reservation system can reduce uncertainties, making the process of selecting and using a charging station simple and efficient. This approach has been proposed by Cuchý, Jakob, and Mrkos [CJM24].

Currently, the practice of reserving a charging station for EVs in advance is uncommon. While there are exceptions[1], most charging station operators follow a first-come, first-serve basis with fixed rates.

This lack of reservations may be because reservations and dynamic pricing are *currently* not necessary. Industry estimates indicate that the number of EVs on the road in 2022 was still around 1% [22a]. With the current pace of development of charging infrastructure, this figure is low enough so for drivers to generally find available charging stations reliably when they look for them. Moreover, the current impact of EVs on the grid remains marginal. Most of the existing EV charging infrastructure was built relatively cost-effectively utilizing existing parking spaces and electrical infrastructure, resulting in relatively low costs for adding new charging stations.

Nonetheless, projections from the International Energy Agency suggest that electric vehicles composed 14% of all new car sales in 2022, a significant increase from 9% in 2021 and nearly ten times the figures from 2017 [23a]. Consequently, the anticipated rapid growth in the number of EVs on the roads is expected to strain existing charging infrastructures. As more charging stations are needed, the associated costs of installing them will likely escalate. This will likely make traveling with EVs more difficult, as drivers will have to compete for the limited charging capacity.

**Why EV charging reservations?** The concept of EV charging reservations has potential of significant advantages for different stakeholders, including drivers, charging station operators, and electrical grid administrators. For drivers, the ability to reserve charging slots in advance allows trips to be planned reliably and mitigates the risk of encountering fully occupied stations upon arrival. For charging station operators, reservations can lead to increased revenue by minimizing idle time. Additionally, as EV charging grows as a significant consumer of electrical energy, grid operators face the challenge of managing peak demand. For them, reservations can help by distributing demand more evenly over time.

However, there are drawbacks to consider. Currently, reservation services, such as those deployed by EVgo, require a fee, which is not pupular to the drivers. Moreover, much of the existing infrastructure lacks reservation support, due to incompatability of older protocols and hardware with the concept of reservations. There are also practical challenges to reservations, such as the possibility of an EV charging slot being occupied by

[1] Such as EVgo, who introduced charging station reservations for a flat fee in 2021.

an internal combustion engine vehicle, despite being reserved. These and other challenges have been identified in the ELECTRIFIC project [22b].

Given that the practical implementation of reservations in EV charging is still in its infancy and they play a crucial role in our dynamic pricing strategy, this work explores the feasibility of such reservations. We do this through a simulation model to assess the reservation use-case and evaluate its benefits for stakeholders in electromobility.

**Why Dynamic Pricing?** Reservations mitigate uncertainty for both EV drivers and charging station operators. Yet, the real advantage lies in dynamic pricing's ability to influence driver behavior that maximizes the utility of reservations for a Charging Station (CS) and grid operators, and by extension, for EV drivers.

Dynamic pricing is one of the proposed methods of changing customer behavior, specifically to discourage charging when the electrical grid is nearing maximum capacity [Rig+13; BG12]. Unlike other coordination strategies for charging [Ste+16; Bay+13; De +13; Ete+17], dynamic pricing can be deployed locally by the charging station operator, and does not require difficult-to-acquire system state information. It serves the charging station operator's interests and, through contractual pricing arrangements with the grid operator, can aid in reducing grid peak loads. Furthermore, dynamic pricing is a natural extension of existing EV charging practices.

Currently, EV charging typically uses fixed hourly or kWh rates, with additional fees. Incorporating dynamic pricing by adjusting these rates poses a challenge: true dynamic pricing complicates planning for EV drivers. On the other hand, announcing the price changes ahead of time compromises the system's adaptability, hindering the charging station operator's capacity to respond to system state variations. In the proposed system, EV charging sessions are priced in entirety when they are requested, allowing for session prices to accurately reflect the system's current state and anticipated demand. This approach allows charging station operators to price sessions based on the existing and prior requested sessions as well as other state information. Additionally, the reserved sessions have a known end time, allowing the charging station operator to schedule consecutive sessions predictably. For EV drivers, a reserved session's price is fixed, reducing uncertainty about trip cost. This dynamic pricing model is the focus of this thesis. This type of pricing is not equally useful for every charging station and every EV usecase, but it is particularly useful for fast charging stations around highways that are used for long-distance travel. Usecase at the opposite end of the spectrum would be charging at home or fleet charging, where the EV charging is much more flexible and other approches are more appropriate, such as [RFH14].

Our use-case for dynamic pricing is similar to the way that airline tickets are sold, where prices can fluctuate based on expected and actual demand throughout the selling period. Extending the analogy, charging sessions and airline tickets are perishable and cannot be stored for future use. The sale of a charging session also influences subsequent sales, a characteristic shared with products possessing a "network effect," such as multi-leg airline tickets, car rentals, or multi-day hotel stays. EV charging pricing is further complicated by the station's power capacity, which in the future may fluctuate depending on the electrical grid's state. This is similar to the dynamics of on-demand mobility platforms, where drivers have the flexibility to enter or exit the system. The successful application of dynamic pricing in these other sectors suggests its potential success for EV charging sessions, enhancing revenue and managing demand effectively.

**Why Markov Decision Processes?** The challenge of dynamically pricing EV charging is multifaceted. Prices must accurately reflect anticipated demand and the state of the CS, both of which are inherently uncertain and can fluctuate over time. Fortunately, CS operators have access to historical data, enabling them to construct a probabilistic model to predict demand and assess the state of the CSs. The task then becomes one of determining the pricing for a series of charging session requests, guided only by probabilistic forecasts of future requests. This is a sceneraio of sequential decision-making under uncertainty, and it can be modeled by the MDP framework.

MDPs offer a robust, well-established methodology for addressing sequential decision-making under conditions of uncertainty. MDPs have been successfully applied in various domains facing similar dynamic pricing challenges, including the dynamic pricing of electricity [LHZ18], on-demand transportation [Zuo+19], and cloud computing services [XL13]. In operations research, a closely related issue known as network revenue management, primarily inspired by airline ticket sales, similarly employs MDPs. However, in this context, the emphasis is on manipulating available capacity across different fare classes [SKS18].

## 1.1 Thesis Goals

The primary aim of this thesis is to develop a dynamic pricing model of EV charging that can accommodate future developments in the field of electric mobility. This aim is subdivided into three interconnected Objectives:

**Objective 1: Determine a reservation product for EV charging through multi-agent simulation.** The initial objective is to identify a suitable product for dynamic pricing, focusing on the reservation of EV charging capacity for future use by EV drivers. The reason is that the currently used first-come, first-serve model, priced by hour or kWh, will not be suitable for widespread adoption of EVs, particularly at fast-charging, highway stations. Our focus is on the usecase of long-distance travel with fast charging. The selected product must make planning predictable and reliable for EV drivers, aligning with studies indicating risk-averse tendencies among EV drivers [LSP17]. The product should also be beneficial to the charging station operators.

However, we shall also consider the practical challenges of implementing reservations, such as compatability of existing hardware and software solutions deployed at existing charging stations. Important aspect of of the product is that it should be applicable for use with *automated systems*, such as driver assitant technologies and route planners on the side of the driver [CJM24], and automated pricing and capacity management on the side of the charging station operator. Identifying a reservation product that is both widely beneficial and feasible to implement could have a significant impact on the future of EV charging.

Although reservations are not widely recognized in practice, most recent works covered by Saharan, Bawa, and Kumar [SBK20] in a survey on dynamic pricing of EV charging, describe pricing per hour or per kWh, with rates updated hourly, for example [Fan+21]. We propose that viewing charging as a holistic service — including parking, with defined start and end times — constitutes a crucial aspect of a dynamic pricing strategy that ensures predictability and reliability for EV drivers and charging station operators.

Since it is not feasible to implement a comparative study of different reservation types in practice, we aim to use a *multi-agent simulation* to evaluate the benefits of different reservation types for electromobility stakeholders. Once we determine the appropriate reservation product, we can use it as a foundation for the dynamic pricing model.

**Objective 2: Develop revenue maximizing MDP model of dynamic pricing of EV charging products that can be solved optimally.** While determining the right reservation product is fundamental to the success of the dynamic pricing model, it is the dynamic pricing method built on top of it that can provide the incentive structure for achieving equilibrium between the interests of the charging station operators, EV drivers and the electrical grid operators.

Many works on pricing of EV charging focus on maximizing social welfare or other similar objectives. In contrast, our focus on the revenue maximization of the charging station operator. This is because sustainability of the charging infrastructure is crucial for the widespread adoption of EVs. The second objective focuses on simplifying the problem of pricing EV charging into an MDP model that captures the essential aspects necessary for dynamic pricing. This model should be sufficiently straightforward to allow for optimal resolution using standard techniques. It will serve to both understand the problem and confirm the feasibility of employing MDPs for dynamic pricing of EV charging. Using MDPs or more generally stochastic dynamic program to model the problem of dynamic pricing of EVs is an often made choice due to the natural properties of the problem [Fan+21; LHG18]. Our goal here will be a dynamic pricing model of whole charging reservations that will be optimally solved on non-trivial problem instances.

Since our goal is to solve the problem optimally, we expect to simplify the problem to a degree that allows for optimal solution. However, this might limit the model's applicability to real-world scenarios.

**Objective 3: Develop a practical model of dynamic pricing of EV charging products that can solve realistically sized instances.** The practical applicability of the reservation product is important to us, and so is the applicability of and the dynamic pricing model. The second objective will provide us with a model that is optimally solvable, but it might be too simple to be usable in practically sized problem instances.

Therefore, the final goal is to create an MDP-based dynamic pricing model applicable in real-world settings. This model must be adaptable enough to incorporate future electric mobility innovations, like dynamic distribution costs, and capable of efficiently resolving realistically-sized problems. Additionally, the model should be fast enough to be used in real-time, or near real-time, scenarios, where the charging station operator needs to respond to the current state of the system. Also, the model should not require difficult to acquire system state information, as this would limit its applicability to real-world scenarios, or require excessive compute to train, retrain, or use.

## 1.2 Thesis Outline

The thesis opens with Chapter 2, where we introduce concepts and methods used in the rest of the thesis. This includes background on the electromobility ecosystems and its stakeholders. The literature review of reservations and dynamic pricing of EV charging is in corresponding chapters. In addition, we introduce the properties of the customer arrival process used in our demand model and a brief overview of MDPs and methods for solving them. In addition, we provide an overview of different pricing mechanisms and their applicability to EV charging.

In Chapter 3, we look at the problem of reservations in EV charging. We discuss the feasibility of reservations and associated technical challenges and propose a multi-agent simulation model for evaluating the use case of reservations. This chapter is primarily based on the paper [Bas+19], and is supported by results presented in [Bas+20].

The reservation product identified in Chapter 3 is used as a foundation for the dynamic pricing model in Chapter 4 that introduces our initial model of dynamic pricing using MDPs. This model is designed to be simple and concise, and we use it to familiarize ourselves with the problem. We managed to solve the problem optimally using standard MDP solution methods and evaluate it in simulation. This chapter is based on two conference publications, [MKJ18a] and [MKJ18b].

In the penultimate Chapter 5, we introduce an advanced MDP dynamic pricing model. This model addresses weaknesses and limitations of the model from Chapter 5 at the cost of increased complexity. To find the pricing policy for this new model, we need to employ a customized heuristic Monte-Carlo tree search (MCTS) solver. We evaluate the performance of the solver in simulation and compare it to the optimal solution and other baselines. This chapter is based on a an unpublished journal draft titled “Dynamic pricing of EV charging reservations with MDPs” that extends the work done in [MKJ18b], with preliminary results presented in [MB22].

Chapter 6 concludes the thesis and discusses possible future work.

# CHAPTER 2

# Background and Related Work

In this chapter, we will describe the background for the following chapters. We will start with the electromobility ecosystem and its stakeholders. Next, we will conceptualize and analyze reservations of EV charging and why we consider the type of reservations that we do, also providing related work references in the process. Then, we will discuss dynamic pricing of EV charging, motivations for dynamic pricing and related work in dynamic pricing of EV charging. Last, at the end of the chapter, we will provide technical background on methods and models used in this chapter, starting with the customer arrival process and budget distributions, followed by general description of MDPs and methods for solving them.

After more than a hundred years of niche use, battery EVs seem on the cusp of displacing internal combustion engine vehicles in personal transportation. The promises of the transition are great, including independence on the fossil fuels in transportation, cleaner cities and a reduction in $CO_2$ emissions.

However, the transition to electric mobility is not without its challenges and is still not certain. Book by Charette [Cha23] starts with a chapter titled "The EV Transition

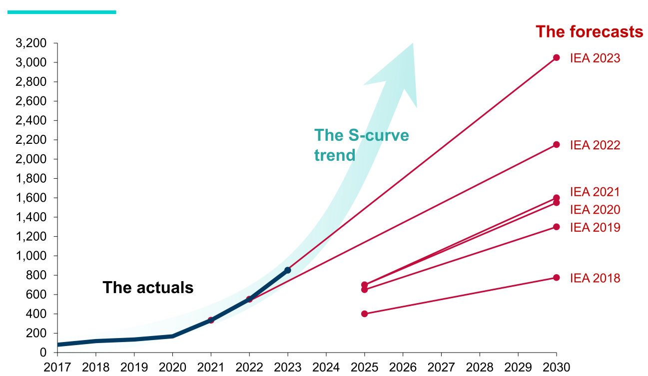


**Figure 2.1:** Development of battery demand forecats according to the International Energy Agency, as collected by [Sol24]. This illustrate the unpredictability of the future developments in the field of electromobility.

is Harder Than Anyone Thinks", which is a fitting title. This difficulty is caused by many unpredictable trends that lead to a lot of uncertainty. Reason for this is that the transition to electric mobility is part of a larger transition to carbon-free energy. The scale of this transition is unprecedented and stakeholders of this change include the automakers, minining companies, the energy producers, the grid operators and governments. In some form, this transition impacts everyone.

The stakeholders invoved in the transition are interconnected in a complex web of dependencies. For example, the automakers are dependent on the mining companies for the raw materials for the batteries. The EV users depend on the charging infrastructure and the grid operators. The grid operators depend on the energy producers and the government, who are at the same time involved in transitioning the energy sector to carbon-free, often intermittent, sources.

At the same time the future developments for each of the stakeholders are highly unpredictable. As an illustrative example, consider the production of lithium-ion batteries, an essential component of an EV. The availability of the primary resources is of great concern in the production, and "it can take five or more years to get a lithium mine up and going, but operations can start only after it has secured the required permits, a process that itself can take years [Cha23]". So in order to mine more lithium, decisions must be made well ahead of time. At the same time, the International Energy Agency, the foremost authority on global energy sector, has predicted in 2018 the demand for lithium-ion batteries in 2030 to be 0.78 TW h per year in it's conservative prediction based policies announced at the time [18]. In 2023, this conservative prediction for 2030 was above 3 TW h. At the same time, the most *optimistic* prediction in 2018 was 2.2 TW h which is still below the conservative prediction from 2023. This illustrates the perils of predicting the future for only one of the stakeholders, but similar situation plays out for most of them.

Report by Rocky Mountain Institute [Wal+23] identifies positive feedback loop between market scale, cost, and quality as the result for the exponential growth of battery sales: "As the battery market grows, unit cost keeps falling and quality keeps rising". What is more, cheaper and better batteries open up new markets, in different sectors and geogrephies, further driving growth. This may account for the difficulties of predicting the future of the battery demand.

Nevertheless, similar growth stories appear to play out in other relavant sectors. According to the Solar Energy Industries Association, an industry trade association, the costs of new photovoltaic installations in the U.S. have fallen by about 40% between 2010 and 2023 and new capacity was added with 24% annual growth rate [23d]. However, as solar is an intermittent energy source, the growth in photovoltaic installations is also driving the growth of the battery energy storage market.

The speed of change and the cross-sector interconectedness of the developments make not only for difficult predictions but also uncover unexpected challenges and obstacles to development. For example, utility planning and regulation have historically considere only utility-scale supply and delivery infrastructre. However, rising power generation from distributed sources, such as residential photovoltaics, is rising the cost of supporting grid infrastructure and drives a need for new regulations and planning tools [Cla+20].

The goal of this discussion was to illustrate the folly of authoritatively predicting the future of the electromobility. In this brief example, we have not touched on the difficulties caused by changing goverment incentives, international politics, trade wars and societal aspects of the transition that will make many skills obsolete while creating demand for vastly different ones [Cha23].

However, some obstacles to the electrified future are apparent already. According to the World Economic Forum [22c], the top two barriers to the development of electromobility today are primarily:

**Inadequate charging infrastructure** Insufficient charging infrastructure slows down adoption of EVs, as the users can not be certain to reliably charge their cars. This goes hand in hand with poor Charging Station (CS) maintenance and the lack of interoperability between different charging networks. Part of the reason for this is that the charging infrastructure is still in its infancy and the business models for charging stations are still being developed. Article published by McKinsey notes that, at least in the US, neither of the two most popular business models are profitable [Kam+22]. And continues "Making it profitable to sell public-charging services will probably be a prerequisite for building out a nationwide infrastructure". Dynamic pricing has the potential to improve the CS revenues and profitability, making the charging infrastructure more attractive for investment and expansion.

**Risk of grid overload** Grid overload is a problem already with increasing share of renewables in the electricty generation that will only be excasperated by the increasing number of EVs. To illustrate the gravity of the situation, fast charging stations today usually offer charging speed between 50 and 60 kW, but speeds go up to 350 kW. For a comparison, an average instantaneous power consumption[1] of a U.S. household is about 1.2 kW. The costs of upgrading the distribution network to cover just the night charging at maximal possible peak power intakes would be extreme, on par with building the grid for additional three times the number of households [Fai09]. However, such excessive upgrades will not be needed since the charging stations are not going to charge at full power all the time. However, some grid updates and smart-grid technologies will be required to cordinate the charging of EVs with the grid capacity. Dynamic pricing can be used here effectively to incentivize charging behavior that is beneficial to the grid, reducing the risk of grid overload.

For the dynamic pricing to realize its potential in increasing the CS profitability and reducing the risk of grid overload, the relationships between the CS and EV user on one side and charging station operator and grid operators need to be updated. Next, we

[1] Based on the 2015 https://www.eia.gov/tools/faqs/faq.php?id=97&t=3statistics of the U.S. Energy Information Administration.

discuss these as well as other stakeholders of the electromobility ecosystem and their relationships.

## 2.1 Electromobility Ecosystem

In this thesis, we focus on the dynamic pricing of EV charging which is a part of an electromobility ecosystem. The ecosystem consists of many stakehlolders, the ELECTRIFIC project [22b] has identified 20 relevant electromobility stakeholders and analyzed their various relationships [17]. However, this large number covers many different usecases that are not relevant to the dynamic pricing of EV charging.

> Our usecase for dynamic pricing is EV user charging at a public fast charging stations as part of long-distance travel.

Therefore, for our purposes, we will focus on the EV users, the charging station operators, and the grid operators.

**EV users** is anyone planning a long distance trip with an EV who is responsible for the charging of the vehicle. This includes the private EV owners as well private and commercial EV fleet users. The EV user is trying to optimize his cost and travel time from the trip, and in general maximize his convenience.

**Charging station operators** are central actors to the charging ecosystem. They are responsible for the operation, management, and maintenance of the charging stations. We are only interesetd in charging station operators who offer publicly available charging services. The charging station operator may or may not own the physical charging station and may or may not own the land and parking spot where the charging station is located.

**Electric Mobility Service Providers** are the companies that connect the EV users with the charging station operators. They can offer a range of services, such as aggregating third-party charging stations and making them available to EV users, unified payments and billing, charging station status updates, navigation, and customer support.

**Grid operators** are responsible for the operation, management, and maintenance of the electrical grid. *Grid operator* is an umbrella term for transmission system operators and distribution system operators, who have different responsibilities and roles. The transmission system operator is responsible for the long-distance transmission of electricity over the high (from 132 kV to 750 kV in continental Europe [23b]) voltage grid, while the distribution system operator is responsible for the local distribution of electricity over the medium and low voltage grids. In Europe, the transmission system operator is responsible for controlling the grid frequency and compensation of the grid imbalance.

The distribution system operator is not responsible for balancing power, but has an obligation to offer non-discrimatory access to the grid.

Examples of other relevant stakeholders who are not as closely invested in EV charging include the energy suppliers who buy and sell electricity, local and large energy producers, energy exchanges, EV fleet operators, manufacturers, governments and regulatory bodies.

Interactions between the different stakoholders are layered and complex. For example, large energy producers supply electricty the high-voltage transmission grid and sell it at the energy exchange or directly to energy suppliers. The distribution fee for using the local low-voltage grid for getting the energy to the final customer is charged by the energy seller to the customer. At the same time, the transmission operators can buy supplementary services from different actors to help balance the grid. Small local energy producers can also sell energy at the exchange or to energy sellers, but they are usually connected to a low-voltage grid and not directly to the transmission grid.

However, to connect new local power source (such as rooftop solar-panels) to the grid, the grid operators need to approve the connection based on the maximum capacity of the local grid and maximum power supply of the local source, even though these maximums are only rarely reached. Ideally, these sources of energy could be used locally, avoiding the grid bottlenecks.

For connecting new local loads, such as a charging station, the situation is similar. The CS operator needs permission from distribution grid operator to install a charging station, who needs to make sure his distribution grid meets the maximal capacity offered by the charging station. Without this permission, the new CS can not be built. However, once the grid is sufficently built-up, the charging station operator today has no incentive to coordinate the EV charging with the grid operator.

The whole process is monitored and regulated by the national regulatory bodies, to make sure that the different actors are not abusing their position and to prevent blackuots, caused by mismatch between supply and demand and grid overloads.

However, the incereasing electricity production from distributed renewable sources and the increasing number of EV charging stations is putting a strain on the existing structures. Both consumption of energy in the charging stations and the production of energy from distributed, renewable sources is intermittent. Infrastructure buildup is innevitable with the energy transtion. However, for example, the requirement of sizing the grid capacity to the peak load and production could be relaxed in the smart grid future with dynamic contracts between stakeholders.

Then, CS operators would be faced with smaller installation costs and be incentivized to coordinate the charging with the grid operator, and could be rewarded for doing so. Dynamic pricing of EV charging would then be the tool with which the CS operator would incentivize the EV driver to charge at different times or places.

Notably, there are already ways for grid operators to control the state of the grid. For example, in Czechia, the grid operator requires local power generation units with power about 100 kW to be equipped with remote terminal units that allow the grid operator to remotely control and even shut down production [23c, p.7]. However, this is a contractual obligation and the producer is not informed ahead of time about the shutdown, nor incentivized to coordinate the production with the grid operator. While, at least in Czechia, the grid operators exercise these controls only rarely (couple of times per month) and only in some regions, the situation could change with the increasing number of distributed power sources and EV charging stations. Widely deploying a similar system for the EV charging would be unfortunate for the EV users and the CS operators, as it would unpredictibly reduce quality of charging services.

### 2.1.1 Types of Charging

There are two types of charging: destination or home charging and en-route charging. Destination charging is done at home or at work, where the car is parked for a longer period of time. En-route charging is done at public charging stations. For destination and home charging, slow chargers are usually sufficient, with power of around 3 kW, which translates into 10 or more hour charging for typical battery. For en-route charging, fast (7 kW to 22 kW) and ultra-fast (50 kW to 120 kW) chargers are much preffered. Fast chargers can fill[2] charge the battery in 4 to 6 hours, while ultra-fast chargers can charge the battery in one hour, and in 20-30 minutes to 80 % battery capacity.

The high charging speeds makes the fast chargers a desirable option for long-distance travel. Coincidentally, fast chargers require much stronger electrical grid and are more expensive to install and operate than slow chargers. With the expected fast growth of number of EVs, the demand for fast charging will likely outpace the rise supply. This creates stronger incentive for EV users to use and for charging station operators to implement reservations to allocate the limited charging capacity.

In other domains, such as air travel, reservations went through a similar process and are now a common part of mechanism for allocating capacity. At the start of the commercial passenger service, airline tickets were “like money” in a sense that passengers who bought a ticket could show up at the airport at any time. However, as the industry grew and became more competative, air-fares became a lot more complex through the practice of yield management. The requirement of reservations, as well as other constraints placed on ticket sales (such as advance purchase requirements, minimum stay requirements, overbooking etc.) made air travel much more efficient and affordable.

## 2.2 Reservations of EV Charging

This section aims to provide the necessary context to the concept of reservations in EV charging that explain the design of our choices in the rest of the thesis. After discussing the appearances of reservations in the EV charging literature, we start by analyzing the stakeholder requirements for reservations and classifying the types of reservations along two basic dimensions (Section 2.2.2). This analysis then serves as a

[2]Due to physical limitations, battery charging speed will vary based on the charge of the battery, with an empty battery being capable of quick charging up to about 80% of charge. The final 20 % are much slower.

basis for the design of reference charging reservation architecture that relies on the E-Mobility Systems Architecture (EMSA) framework [Kir+19]. This provides a structured understanding of the possible reservation systems within the electric mobility context (Section 2.2.3). The reference architecture also lets us identify important features of charging reservation systems, resulting in morphological analysis [NR15], which we use to categorize different reservation types. Finally, we showcase the use of the reference architecture by instantiating the ad-hoc reservation type (Section 2.2.3). Furthermore, we describe the application of this reservation system to real hardware in proof-of-concept experiments in Bavaria (Section 2.2.3).

Parts of the this section (especially Sections 2.2.1 to 2.2.3) rely heavily on expert contributions of the article co-authors first introduced in [Bas+19; Bas+20], others, such as the classification of reservation types in Section 2.2.2, were collaborative effort. The two publications on which this section is built also rely on the know-how and results obtained in the course of the ELECTRIFIC project (e.g., see publicly available report [17]), such as the stakeholder analysis (Section 2.2.2). The proof-of-concept of the reservation system on real hardware (Section 2.2.3) was also performed in collaboration with project partners from the industry.

### 2.2.1 Related Work

In the existing literature, contributions related to the charging of EVs can be classified into two major research clusters: charging scheduling and charging station selection, where both branches discuss reservations from different vantage points.

#### Charging Scheduling

The first cluster of papers prioritizes optimal EV charging schedules for parked fleets of EVs. These works aim to regulate the charging to minimize peak demand, cost, and balance the overall EV demand for electricity, with primary stakeholders being EV users and grid operators like Distribution System Operators (DSOs). The main objectives are (1) reducing power grid peak loads and (2) ensuring a sufficiently charged battery for subsequent trips. These strategies are typically termed “valley filling” or “smart charging”, with examples in [Wan+14; WWX15; YU16; Kha+18; BD18; Bas20]. A thorough review is in [MG15]. To our knowledge, none of these contributions proposes reservation systems to manage infrastructure more efficiently. Such a reservation system is crucial for demand-side management by both public and private stations, either through direct load control or capacity planning.

#### Charging Station Selection

The second research cluster centers on determining the optimal charging station in terms of proximity to a planned route and CS availability (e.g., utilization). These works, focusing on EVs on planned trips, aim to reduce waiting and trip times. The primary stakeholders here are the EV users and the CS operator, with key goals being (1) enhancing EV user experience and (2) maximizing CS operator profit. Contributions differentiate between short (city) and long (highway) trips. In highway contexts, [Yan+13; Yue+17] forecast waiting times at charging stations based on estimated queues, prioritizing stations with minimal utilization. Another strategy focusing on either minimal waiting or proximity for long trips is in [Rig+13]. In [CJM24], we use reservations implicitly when optimizing EV trip for cost and travel time. For urban settings, [Rig+13] employs a dynamic pricing model to mitigate congestion while maximizing profit.

Various authors propose reservation systems for both scenarios to refine charging station choices and minimize wait times. These require estimating expected delays (e.g., due to traffic, different charging profiles, CS queues) and exploring their transmission between the different participating agents. For example, vehicular networks are explored in [Kim+10; LPK11] with objectives to optimize user satisfaction and station utilization. The contribution of [Yue+17] emphasizes communication infrastructure (e.g., roadside devices to communicate CS state to EVs) to minimize both drive time and charging costs. [Liu+18] presents a decision model for reserving stations, taking into account traffic and charging resources, with the presumption of efficient communication networks. Real-world reservation system implementations are discussed in [Orc+18; Bas+19], leveraging the Open Charge Point Protocol (OCPP) for communication and reservation. One presents a centralized approach, while the other uses end-to-end communication.

Many prior contributions utilize proprietary designs, focusing on specific stakeholders like the EV user or CS/grid operators. In contrast, our focus in [Bas+19; Bas+20] was on introducing a generic reference architecture for EV charging station reservations, encompassing all critical stakeholders in electric mobility. We outline the goals of each stakeholder, identify standardized data models and communication protocols, and describe the foundational physical infrastructure. This results in a reference architecture that ensures interoperability, serves as a foundation for tailored implementations and enables the classification of reservation systems in electric mobility.

### 2.2.2 Reservations Analysis

We start our analysis of reservations by examining the requirements of both the EV charging station operators and the EV users, as derived from our literature analysis (as discussed in Section 2.2.1 and [17]). Then, we classify different potential reservation systems along the dimensions of reliability and planning, followed by morphological analysis of the reservation systems in terms of features different reservation systems can have.

#### Stakeholder Requirements

In our analysis, both electric mobility service provider (companies connecting EV users and third-party charging station operators) and charging station operator represent the supplier side, with the EV user the demand side (see Section 2.1 for description of stakeholders). The grid operator, here considered as a distribution system operator, has stakes in the operation of the reservation business. Detailed analysis of grid-side

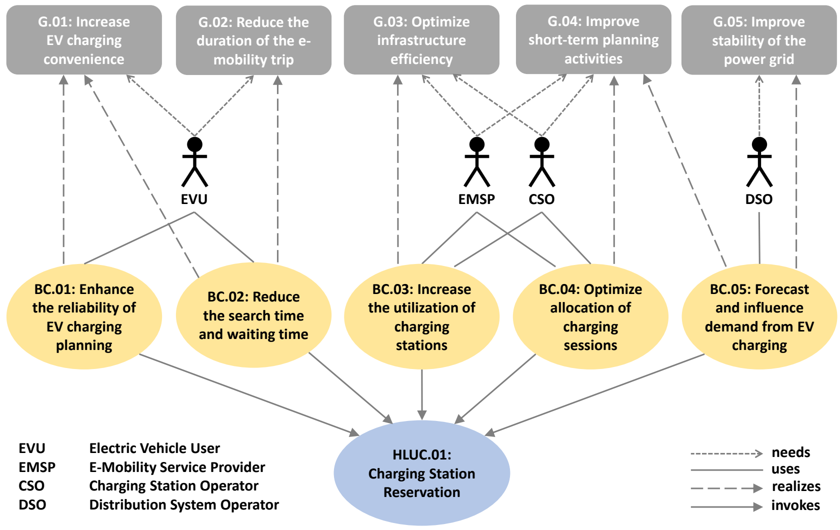


**Figure 2.2:** Charging station reservation business and requirements analysis, from [Bas+20], based on [17]

stakeholders is beyond the scope of this thesis, but it is included in Figure 2.2 to provide a comprehensive picture for the reference architecture.

Figure 2.2 illustrates the key stakeholders to the reservations and their goals (G), business cases (BC), and the high-level use case (HLUC), which is the charging station reservation system. For EV users, the main objective is to improve the user experience. That boils down to three things:

1. reducing trip duration,
2. making charging more convenient, especially by reducing waiting times,
3. reducing the travel cost, including the charging fees.

These goals can often be conflicting. For example, reducing trip cost may require a detour to a cheaper charging station, which increases trip duration. How EV users handle these trade-offs depends on the user's priorities. It is reminiscent of multi-objective travel options known in other transportation problems, where the options balance cheap, fast, and comfortable. Our proposed solution is the use of integrated EV route planners [CJM24], which can offer a concise selection of different tradeoffs.

While reservations can have a direct impact on costs and revenues, we don't include monetary costs in the analysis and instead focus on the time aspect (travel time, utilization). This is because the pricing of reservations is a complex topic in its own right and is the focus of Chapters 4 and 5. That said, for now, we assume there's a small fee to deter users from strategically reserving charging they have no intention of using.

While we don't explicitly address privacy, safety, reliability, and security, these factors largely depend on system and protocol implementations. Our proposed system lacks a central control point. Its decentralized nature strengthens resilience and security by avoiding single points of failure, preventing large-scale data breaches. Regarding privacy, unique vehicle identifiers are required for reservations but don't need to persist between reservations. Of course, payment considerations may discourage the use of ephemeral identifiers.

In [Bas+20], we have also identified *interoperability*[3] as an important requirement for reservation systems due to large amounts of heterogeneity in the electromobility ecosystems due to the large number of substitutive technologies and many different owners and operators of the infrastructure. While this aspect of reservations is important, it is not directly relevant to the topics covered in this thesis, and we, therefore, refer the reader to the original publication [Bas+20] for an in-depth discussion of the topic.

### Classification of Reservation Types

From the EV user perspective, CS reservation approaches can be categorized along two dimensions as illustrated in Figure 2.3: *Reliability* and *Planning*. To recognize the advantages that charging station reservation brings to supplier and user, we identified four different generic types of reservations suitable for the electric mobility context:

- Uncertain Ad-Hoc reservation.
- Guaranteed Ad-Hoc reservation.
- Uncertain Planned reservation.
- Guaranteed Planned (Full) reservation.

In most generally familiar contexts, reservations are thought about as planned and guaranteed, such as a cinema ticket or restaurant reservation. This means that the users can make their reservation well ahead of time (planned), and the service provider guarantees that the service will be provided at the agreed-upon time (guaranteed). Air travel is an example where things are not as simple. Passengers can be bumped from the flight when too many people who reserved their seats show up to board the flight. This practice is called overbooking, meaning the reservations are planned but not guaranteed, resulting in "planned uncertain reservations".

Importantly, in electric mobility, as we will discuss in Section 2.2.3, hardware and software limitations allow many existing CSs only to be blocked until the user arrives (ad-hoc reservation), and even then, the CS might be blocked by an internal combustion vehicle (uncertain reservation).

Figure 2.3 describes the four options. The "Planning" dimension refers to the planning capability and starting time of a reservation. More precisely, "Ad-Hoc" denotes that the reservation can be placed immediately, which blocks the corresponding charging connector of a given CS from being used by other EV users. Only one "Ad-Hoc" reservation can be placed simultaneously at one connector of a given CS. "Planned" indicates that the reservation's start time can be scheduled for some time slot in the future. This

[3] Defined as *"the ability of two or more systems or components to exchange information and to use the information that has been exchanged"* by IEEE [Ger+91].

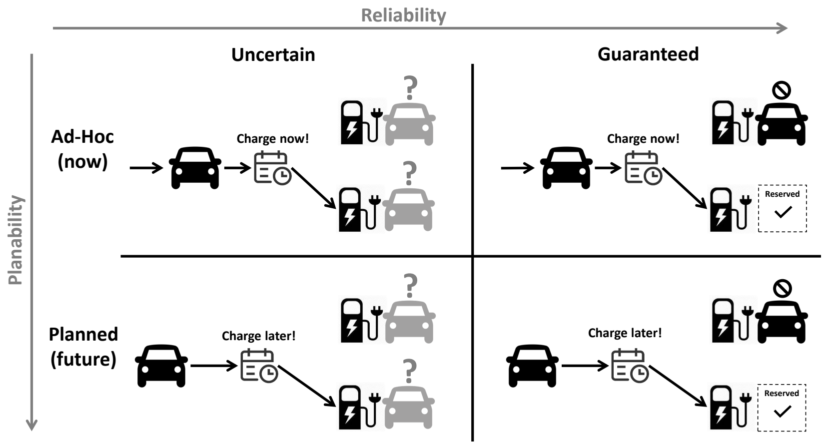


**Figure 2.3:** Classification of reservation approaches by reliability and planning capability. From [Bas+20].

reservation type allows for the CS connector to be used in the meantime. Furthermore, multiple "Planned" reservations can be scheduled for one CS connector by multiple drivers simultaneously. Note that the end time of the "Planned" reservation must be specified, while the end of the "Ad-Hoc" reservation does not.

For effective reservations, users need information on CS connector availability. The connector can either share only its current status or its expected future utilization based on existing future reservations. A connector broadcasting only its current status supports only "Ad-Hoc" reservations, as seen in Section 2.2.3. Full availability details enable advanced "planned" reservations.

The horizontal axis in Figure 2.3 denotes reservation reliability. It is a notable dimension for electric mobility since many CSs may face problems with fulfilling reserved charging obligations. "Uncertain" indicates reservations might not guarantee connector availability, such as when a parking space is obstructed by parked vehicles. "Guaranteed" reservations are viable when enforcement mechanisms, like penalties, sensors, or barriers, are present.

### Morphology of Charging Reservation Systems

In this section, we use the morphological analysis methodology [NR15] to identify important parameters for the design of reservation systems.

Table 2.1 presents parameters and their characteristics for CS reservations. These parameters were gathered through literature research and work done during the Electrific project. They can be understood as follows:

- *Enforceability:* This parameter decides if reservations guarantee charging or remain uncertain because, for example, the CS can be occupied upon user arrival. Implementing enforceable reservations typically requires the installation of additional devices, e.g., sensors or barriers.

| Parameter | Options | | |
|---|---|---|---|
| Enforceability<br>Level of charging certainty | No | Yes | |
| Planning<br>Specified start and end times | No | Yes | |
| Fee<br>Costs incurring from reservations | No | Fixed | Flexible |
| Data Availability<br>Availability of relevant information | No | Limited | Full |
| Roaming<br>Reservation across multiple operators | No | Yes | |
| Scheduling<br>Order in which reservations are processed | Policy | Priority | Auction |

**Table 2.1:** Charging station (CS) reservation system design decisions based on a morphological analysis. From [Bas+20].

- *Planning:* This parameter relates to whether reservation needs both start and end times specified. Without the "planning" parameter, the reservation starts immediately and ends when the user is done. With "planning," users negotiate the start and end time of the reservation with the operator.
- *Fee:* This indicates the costs associated with a reservation: "No" means it's free, "Fixed" means there's a static cost, and "Flexible" means the cost can vary.
- *Data Availability:* This concerns information about charging stations and their connectors. It could be no data, limited data (like the current availability of a charging station or connector), or full knowledge of the reservations queues, charging schedule, and power forecasts.
- *Roaming:* This determines whether users can reserve charging stations from different operators through a single service.
- *Scheduling:* This parameter specifies how reservations are processed. It could be based on policies (like first-come, first-served (FCFS)), assigned priorities (e.g., high or low), or "auctions," meaning any dynamic mechanism requiring active participation of the user.

By setting these parameters, we can create specific reservation systems. For example, a system configured with "Yes" for enforceability, "Yes" for planning, "Flexible" for fees, "Full" for data availability, "Yes" for roaming, and "FCFS" for scheduling would be the "full" reservation system. There are many possible configurations, but not all of them are practical. The choices may be influenced by the goals or hardware or software limitations.

In Table 2.2, we present a decision matrix [LM14] comparing instances of the four reservation types shown in Figure 2.3. The Table 2.2 compares the Guaranteed Ad-Hoc, Uncertain Planned, and Guaranteed Planned reservations against the Uncertain Ad-Hoc reservation baseline. The table shows that the Guaranteed Planned (full) reservation system has the most advantages over the baseline regarding features from Table 2.1. The “Fee” and “Roaming” features are considered orthogonal to the rest of the features that define the reservation type.

| **Criteria** | **Guaranteed Ad-Hoc** | **Uncertain Planned** | **Guaranteed Planned** |
|---|---|---|---|
| Enforceability | 1 | 0 | 1 |
| Planning | 0 | 1 | 1 |
| Fee | 0 | 0 | 0 |
| Data Availability | 0 | 1 | 1 |
| Roaming | 0 | 0 | 0 |
| Scheduling | 0 | 1 | 1 |
| **Sum** | 1 | 3 | 4 |

**Table 2.2:** Decision matrix analysis for the reservation service comparing the three concepts of Guaranteed Ad-Hoc, Uncertain Planned, and Guaranteed Planned reservations with the baseline of Uncertain Ad-Hoc reservations. 0 means as good as the baseline, $-1$ means worse, and $+1$ means better than the baseline. From [Bas+20].

### 2.2.3 Implementing Charging Reservation Systems and Proof of Concept

This section methodically analyzes the components of potential reservation systems and their interactions, identifying future and existing communication protocols that are involved. After this analysis, we identify the distinguishing features of possible charging station reservation types that form the morphology of CS reservations.

#### E-Mobility Systems Architecture

The modeling and engineering of the reference charging reservation system were done using the E-Mobility Systems Architecture (EMSA) [Kir+19] that was developed in the course of the ELECTRIFIC project [20b]. The EMSA is a three-dimensional architecture model (see Figure 2.4) consisting of five interoperability layers (Business, Function, Information, Communication, and Component), four domains, and six zones. The EMSA is based on the standardized Smart Grid Architecture Model (SGAM). The SGAM is a reference framework defined by EU Mandate M/490’s Reference Architecture working group for representing Smart Grid architectures [12]. The SGAM strongly focuses on engineering energy-related systems and provides certain benefits like the smart grid standards map; the EMSA is more suitable for systems engineering in electric mobility since it provides additional details on the respective domains.

In the development and analysis of the reservation system architecture below, EMSA allows us to separate concerns by utilizing architectural viewpoints represented by the layers. Additionally, it ensures interoperability during the design and development phases

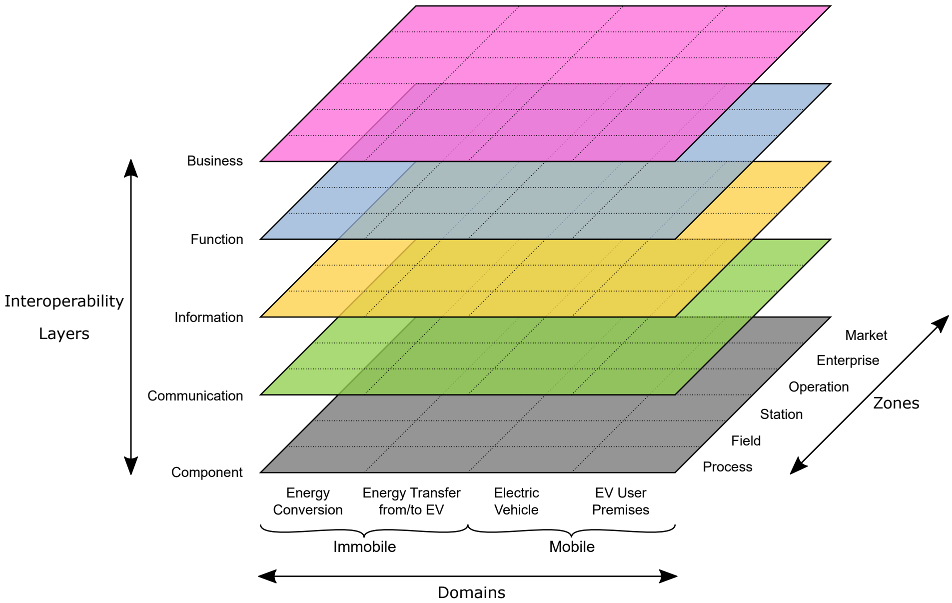


**Figure 2.4:** The E-Mobility Systems Architecture (EMSA) Model [Kir+19].

on multiple layers. The EMSA-based system model consists of the logical architecture (Business and Function layer) and the physical system architecture (Component, Information, and Communication layers).

### Logical and Physical Architecture

To fulfill the high-level use case "HLUC.01: Charging Station Reservation," which we derived from the business analysis (as depicted in Figure 2.2), we need a reservation system with various capabilities. This reservation system's logical structure comprises of four primary functional building blocks, as illustrated in Figure 2.5a.

Within "Function.01," the EV user requests specific CS information from their EMSP via a reservation service. The EMSP then returns this information and displays available CSs to the user through the reservation service's front end, such as a mobile application. The user then selects a CS connector for their reservation. This reservation request is relayed, optionally through a charging roaming service, from the EMSP to the CS operator, which subsequently reserves the chosen connector (referred to as "Function.02"). In the event of a successful reservation, the CS operator informs the EMSP, which, in turn, conveys the confirmation to the EV user via the mobile app. In cases where the reservation failed or was denied, the EV user also receives a notification.

Upon the EV user's arrival at the reserved CS, they authenticate themselves using their user ID, access the CS connector, and begin charging ("Function.03"). Additionally, grid operators can directly monitor and influence reserved charging sessions to facilitate smart charging strategies and implement demand-side management ("Function.04"). Note that this logical architecture intentionally excludes payment-related functions.

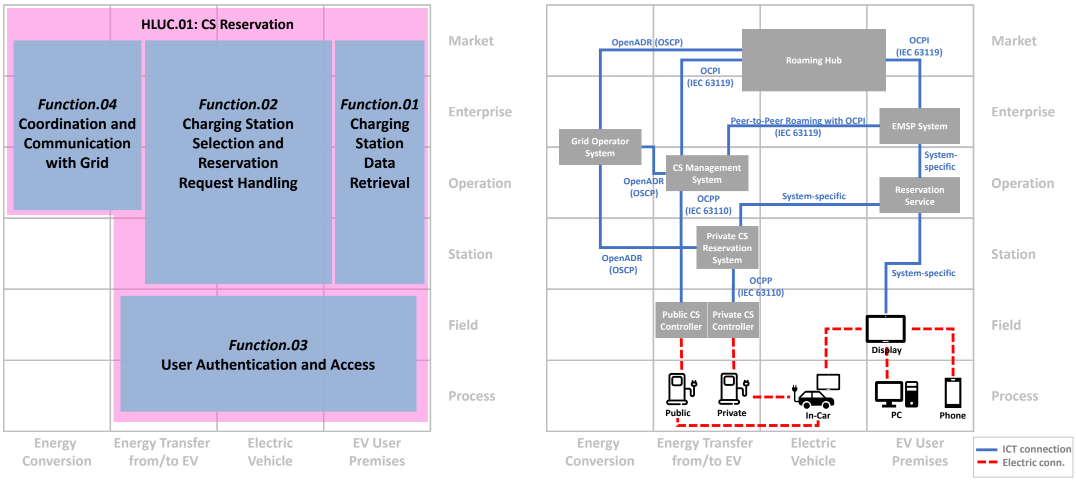


**(a)** Logical Architecture (Business and Function Layer in EMSA).

**(b)** Physical architecture (Component, Communication, and Information Layer in EMSA)

**Figure 2.5:** Reference architecture following EMSA for reservations consisting of logical (a) and physical (b) parts. From [Bas+20].

Since specific reservation types might omit certain steps in this sequential process, we have organized and modeled them as distinct, self-contained functions for clarity and flexibility.

The four functions in Figure 2.5a define the potential functionality of the CS reservation reference system. The physical system architecture consists of multiple components on the Component layer (see Figure 2.5b). The components from the *Field* zone upwards are mainly relevant for the actual software systems, these are connected via ICT (Information and Communication Technologies) connections (blue line). The charging stations, in-car systems, and the EV users' devices (e.g., mobile phones) have a physical, electric connection (dotted red line) to the other components, controllers, and displays. For details on the possible protocols and standards used in these connections, see [Bas+20].

### Uncertain Ad-Hoc Reservation System

The four reservation types shown in Figure 2.3, and described in terms of their features by Table 2.2, are the basic taxonomy of reservation types for electric mobility. During the Electrific projects, we have identified a number of limitations of existing charging station hardware and software that would make deployment of full or guaranteed reservations on existing hardware difficult and expensive. However, the uncertain ad-hoc reservation was determined to be feasible with only limited investment.

In terms of the parameters from Table 2.1, the system has "No" for enforceability, "No" for planning, "No" for fee, "Limited" for data availability (only the current status), "Yes" for roaming, and "first-come, first-served" for scheduling.

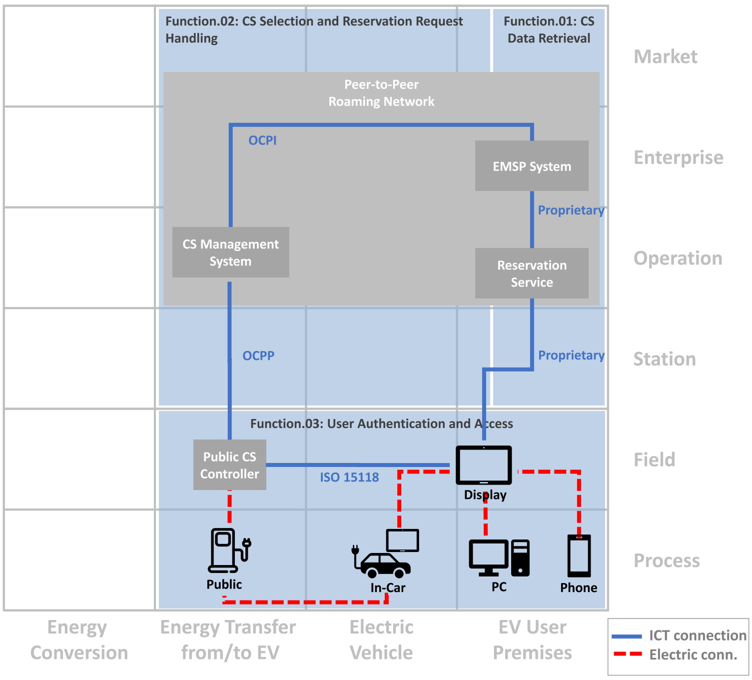


**Figure 2.6:** Reservation system instance with roaming, excluding private CS and grid communication. Modified from [Bas+20].

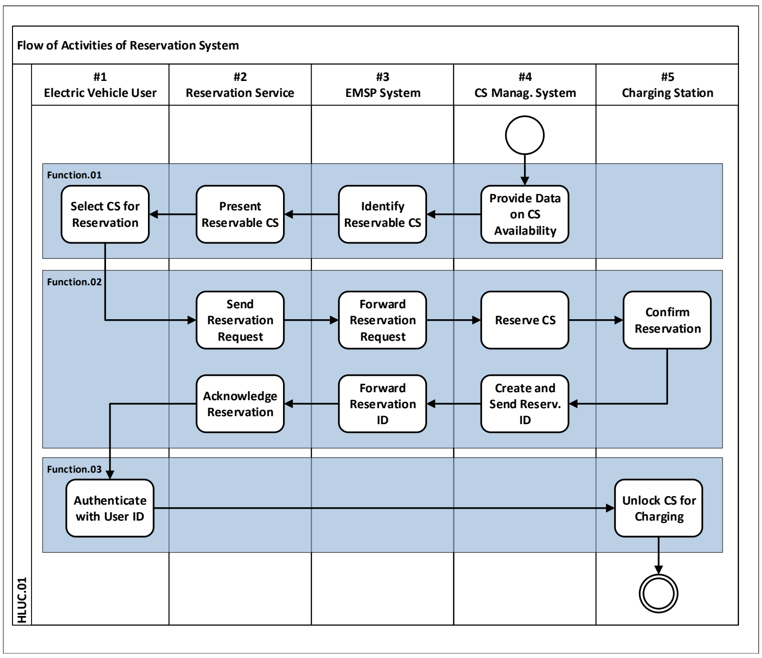


**Figure 2.7:** Operational thread of the reservation system instance modeled as a UML activity diagram. From [Bas+20].


**Figure 2.8:** The message showing on the monitor of the charging station that it has been reserved. From [Bas+19].

Figure 2.6 illustrates the components of the proposed system and their connections within the EMSA grid (see Section 2.2.3), relying on the open OCPI and OCPP protocols. Note the use of the peer-to-peer roaming network. While there was only one provider involved in the proof-of-concept, in [Bas+20], we argue that interoperability is an important aspect of reservation systems, and as such, we present this possible extension in Figure 2.6. The Figure 2.7 shows the operational thread of the reservation system instance, mapping the individual actions to the three major system functions from Section 2.2.3 to the high-level use case for reservations (HLUC.01 in Figure 2.2).

Thanks to the work done in the ELECTRIFIC project, we believe that the system described above could be deployed quickly and cheaply to many charging stations while satisfying the interoperability requirements. To test the first part of this hypothesis (quickly and cheaply), we have deployed the ad-hoc reservation system to several charging stations in a physical proof-of-concept (PoC) system described in the next section.

### Proof-of-Concept of Ad-hoc Reservations

During the ELECTRIFIC project [Eid+17], the electric mobility service provider E-Wald[4] implemented the uncertain ad-hoc reservation system in a field experiment. The experiment was performed at four E-Wald charging stations in Bavaria. The four charging stations of type Alfen ICU in four locations in Bavaria (Lam, Arnschwang, Parsberg, Deggendorf) were used and tested in the showcase. The test was successful for all four charging stations; the reservations were visible on the CS display and in the be.energised[5] CS Management System. EV users with a different RFID (userID) that was not used in the reservation could not charge during an active reservation. Unfortunately, displaying a

[4] https://e-wald.eu/

[5] https://www.chargepoint.com/en-gb/partners/overview

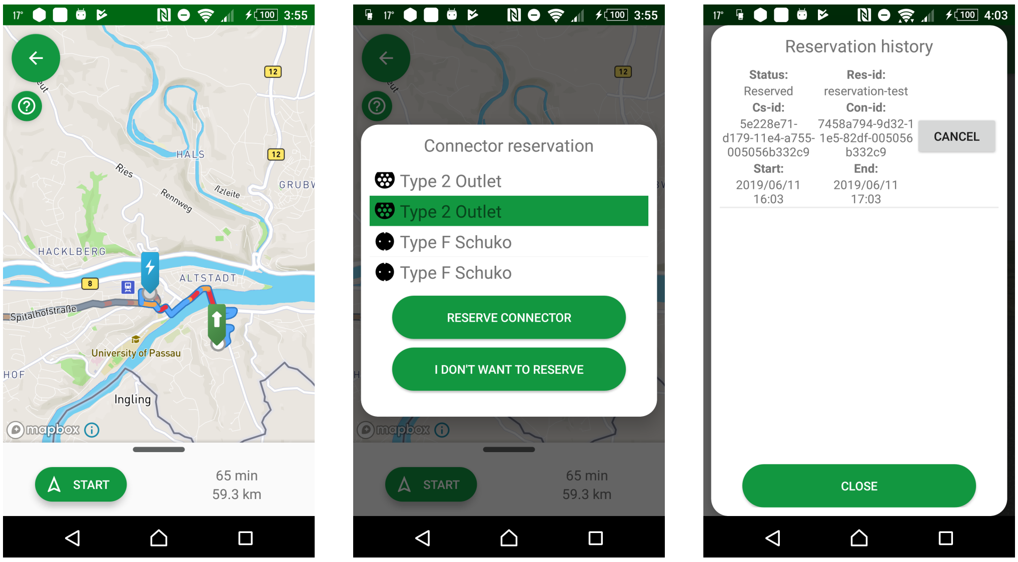


**(a)** Availability information of the charging stations

**(b)** Dialog box for reserving the corresponding charging station

**(c)** Dialog box for showing the summary of the carried of reservation

**Figure 2.9:** The different dialog boxes from the front-end perspective to reserve a connector of a given charging station. From [Bas+19].

custom reservation notification on the display did not work due to the technical limitations of the charging stations.

Implementing the reservation system in the proof-of-concept relied on the back-end reservation system service responsible for communicating with the CS and EMSP systems and the front-end mobile app used by the EV drivers.

The back-end of the reservation system solution, the Reservation Service (see Figure 2.6) relied on off-the-shelf computer (Ubuntu 16.04 with 4-core CPU and 24GB RAM). The back end communicates with the CSMS software to reserve the charging station. To realize the reservation, CSMS sends the request to the physical hardware (e.g., charging station) by means of the OCPP protocol. Once the back-end processes the reservation, a message is displayed on the charging station screen, informing users that the corresponding charger is reserved (Figure 2.8). The whole process is shown in Figure 2.7. The processing of the reservation is also communicated to the front-end application.

The main objective of the front-end for the reservation PoC is oriented towards providing EV users with the availability information about a given charging connector and its quick reservation process. To achieve this, the front end (e.g., mobile app) is designed such that the map displays only the available charging station(s). Figure 2.9a demonstrates the realized front-end such that the blue-marked icon of the CS indicates that the corresponding connector is available. Upon selecting the corresponding CS, a new dialog box pops up (see Figure 2.9b), demonstrating the CS relevant information, such as the connector types that can be reserved. The EV user then has the option of canceling the reservation or, in the other case, receiving a confirmation that the corresponding CS has been reserved successfully (see Figure 2.9c). If the reservation is not successful, the front end will display an error message, leaving the user with the opportunity to select another connector, CS, or skip the reservation process.

This proof-of-concept tests that all the components of the reservation system can be made to work with existing hardware. However, it does not determine whether such a system deployed at scale will benefit all or any of the stakeholders described in Section 2.2.2. This is a question for the simulation study in Section 3.1.

## 2.3 Dynamic pricing of EV charging

In this section, we provide a background on the dynamic pricing of EV charging and different functions on pricing of EV charging. This section is based on surveys done in [17; MKJ18a; MKJ18b; MB22].

### 2.3.1 Functions of Pricing of EV charging

Based on the classification of functions of pricing in the ELECTRIFIC project [17], we distinguish three possible functions of pricing of EV charging:

1. *Distribution of costs:* The price of charging services is primarily used to cover the operational and, in some cases, capital expenses of the charging station operators, grid operators who deliver the energy for charging, and energy producers who generate this energy.

2. *Long-term incentive:* The charging price can also be used to support long-term goals in electric mobility, such as subsidizing the cost of charging to increase EV adoption. Other uses could be to incentivize the use of renewable energy sources or slow home charging.

3. *Signal of resource scarcity:* The charging price can also be used to signal resource scarcity to promote efficient use of the charging resources. An example of this may be increasing the charging price when renewable energy sources are unavailable.

Today, the first and the second functions are primarily used in practice. The charging price is used to cover the costs of the charging station operators, and the cost of charging is subsidized by the government or EV manufacturers to incentivize EV adoption. However, as discussed at the start of Chapter 2, our focus is on the scenario of widespread EV use where the costs of charging are not subsidized to increase EV adoption as most vehicles will be electric.

In this future scenario, the third function comes to the fore since the price of charging can be used as an important signal of resource scarcity and to incentivize desired behavior change in the drivers in the short term. The scarce resources, in this case, could be the grid capacity, the availability of (renewable) power, the availability of charging connectors, or the availability of parking spaces. As discussed at the beginning of Chapter 2, most of these resources are not scarce today but will likely be in the future. These resources have one thing in common: they require the charging price to vary rapidly to incentivize behavior change.

In this work, we focus on the dynamic pricing of EV charging from the charging station operator's point of view, where we assume their goal will usually be the maximization of their profit. This should also mean an improvement in resource utilization over the common baseline of flat-rate pricing. This is often not the case today, as many CSs were built only as a tool to support a general EV adoption without much consideration of their profitability and future maintenance. As a result, the maintenance of many CSs is spotty. However, if the electrified future comes to pass, this is bound to change, similar to how gas stations are not built and operated by car manufacturers to increase car sales.

### Energy Pricing

As summarized in [17], variable energy pricing is not new. Energy generated from base-load power plants, such as nuclear and certain coal-fired facilities, tends to be cost-effective. However, these plants have long start-up times, making their downtime expensive. On the other hand, peaking power plants, which can swiftly come online when demand surges, like gas-fired plants, diesel generators, or battery storage, are relatively more expensive. Despite these variations in production costs, electrical energy is traditionally priced uniformly per kilowatt-hour (kWh).

This flat-rate pricing model does not incentivize users to reduce electricity consumption during peak demand periods, where costlier peaking power sources are utilized. The potential for substantial cost reductions by reducing demand during peak times is recognized [THL10]. However, the prevalent *accumulation meters* don't support variable energy pricing. This will require connected *smart meteres* that grid operators are currently introducing.

Simultaneously, while the widespread adoption of plug-in EV charging may not (at least initially) significantly strain overall power generation capacity, challenges can emerge in the distribution grid if the charging processes are not coordinated [LDL11]. This issue may be further compounded by the intermittent nature of renewable energy sources that strain the distribution grids already.

For the pricing of energy to the consumers, the International Energy Agency [Cro08, p. 61] recommends the following taxonomy for demand-side management pricing schemes:

1. *Time-of-use pricing:* The energy price is fixed to a pre-determined tariff for a given time period, e.g., a set of hours, and varies between time periods. This is the most common form of dynamic pricing today. Time-of-use tariffs reflect the average energy production and distribution cost during the given period.

2. *Real-time pricing:* The electricity prices are set in real-time based on wholesale electricity prices.

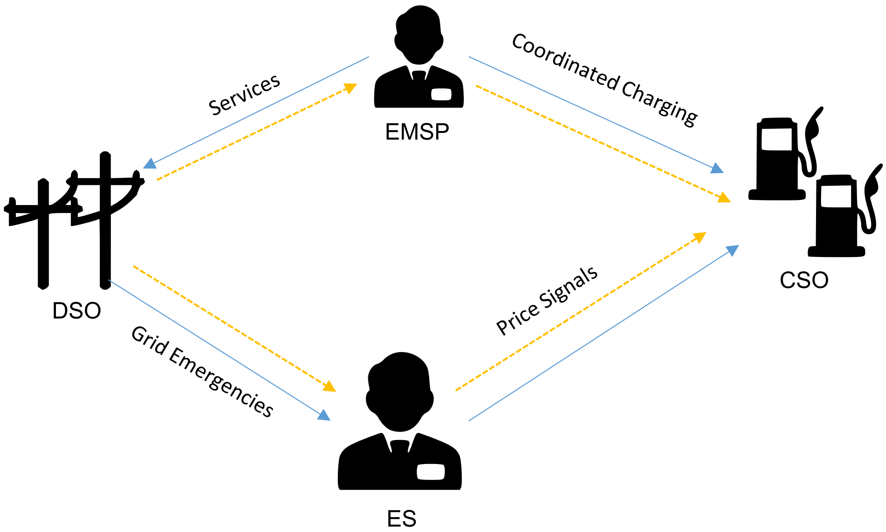


**Figure 2.10:** Example of different kinds of incentive structures from the Distribution System Operator (DSO) and the Charging Station Operator (CSO), either through the Energy Supplier (ES) or the Electric Mobility System Provider (EMSP). Adapted from [17]

3. *Critical peak pricing:* borrows from the time-of-use pricing the predetermined tariffs but reserves an option of replacing the normal rates with much higher peak rates during specific trigger conditions, such as severe grid instability.

The International Energy Agency's report further notes that since the responses to the pricing changes are voluntary, very high multiples (between 10 and 20) of the base tariff may be required to achieve significant demand reduction. However, the report highlights that technological solution that automatically switches loads on and off depending on the price (such as heaters or home EV chargers) can achieve much better results.

This supports our conclusions of Chapter 3 where we highlight the need for comprehensive driver assistant systems that plan the driver's route, including the charging stops, and can directly make reservations. A natural extension of such systems is to allow the drivers to choose not only between the different travel time options but between different cost options as well. Such a system is proposed in [CJM24], which develops a multi-criteria route planner for EVs that allows the drivers to choose between different routes based on different tradeoffs between cost and travel time.

### Pricing Relationships

While the pricing scheme taxonomy described above relates primarily to the sale of electricity to consumers, the same principles can be applied to multiple different actors in the electromobility ecosystem described in Chapter 2. The point here is that the end consumer using the electrical energy is paying not only for the energy itself but also for its delivery and other services. And just as energy generation has variable costs in space and time (e.g., due to renewables), so do the delivery and the supporting services.

Therefore, if the price of EV charging is to vary, it should reflect many of these sources of cost variability. The costs and their variable components can be grouped into three categories:

1. *Energy cost:* The cost of energy generation is the most obvious source of variability. Energy generation costs vary in space and time due to varying energy production costs from different sources. Today, renewables generally have much lower costs than fossil fuels, but their production is intermittent and depends on the weather.

2. *Energy delivery cost:* The different stakeholders to the delivery system pay each other Use of System (UoS) per kW or kWh. The contracts and payments in the distribution can get complicated. For example, the Distribution System Operator (DSO) pays the transmission system operator for the use of transmission lines, but the local distributed generation sources will pay the DSO for using the local distribution grid. The final user pays the distribution fees indirectly through the energy supplier. The variability in the costs of the DSO is caused by the varying load on the grid, which is caused by the varying load of the end users. Subsequently, the DSO has to stabilize the grid and provide reactive power management and voltage control.

3. *Charging costs:* This is primarily the cost of electricity used in charging the EVs, but it also includes the cost of the parking spot (which can vary in many cities already).

Not all the variable costs can be propagated through the system. For example, the distribution costs are heavily regulated and not subject to market forces. However, the auxiliary services provided by DSO can be priced, and the price can be propagated to the end user. To highlight how this can be done, see Figure 2.10 that shows two paths between the DSO and the Charging Station Operator (CSO) for propagating variable prices to the EV users.

The first path goes through the Energy Supplier (ES). ES has contracts with the DSO and CSO (or home charger owner) for energy delivery. DSO could set up contracts with ES for lower distribution fees in exchange for the ES employing demand-side management techniques to reduce or increase the load on the grid, as needed. ES can then pass the savings to the CSO or the home charger owner.

The second branch goes through the Electric Mobility Service Provider (EMSP). EMSP is a company that provides EV charging roaming services. DSO could buy auxiliary services for EMSP who could coordinate the charging of its users to meet DSO goals, passing on the savings to the CSO or the home charger owner.

So far, we have shown how the variable costs can be propagated to the CSO or the charger owner, but we have not discussed the pricing of charging services from the charger owner or the CSO to the EV user. There is a distinction to be made between slow chargers in homes, public slow chargers, and fast chargers:

**Privately owned (slow) charger** owners have a contract with the ES that determines the price, and EV charging is only a part of household electricity consumption. The contracted price is usually a flat rate or time-of-use tariff, but smart meters deployed by the ES or systems for coordinating charging could provide discounts to EV users. This is because privately owned EVs are not likely to be used all the time and spend much time plugged into the home charger. EV users are, therefore, flexible when they charge their EVs. A number of EV users can then be joined in a virtual power plant that sells this flexibility, stopping or slowing charging sessions when electricity is expensive (e.g., as in [RFH14]). A rudimentary variant of such a solution is the use of existing "night-current" time-of-use tariffs. Critical-peak pricing is also used in this context, for example, by the opt-in *Tempo Tariffs* in France [Tho22], where the price hike is a day in advance and the number of peak days in a year is limited.

**Public (slow) charging stations,** usually owned by the DSO or the municipality, are installed at the roadside, in lamp posts, etc. Today, the charging price is usually a flat rate or time-of-use tariff, but the charging station owner may opt to reduce the quality of service. For example, to help stabilize the grid. Same as with privately owned chargers, methods for coordination of charging can, in the future, be used to sell services to the DSO. However, some users of slow public charging stations will not be flexible in when they can charge. Apart from the charging fees, parking fees can be charged to the users, but EV users in public slow charging stations may not be generally expected to vacate the charging spot when finished charging.

**Dedicated (fast) charging stations** are used when the users are not flexible in time and need to recharge as fast as possible. Existing charging stations usually charge a flat rate per minute, with different rates for different charging speeds, or per kWh. The price of energy is currently not a large factor in the pricing schemes of the current charging stations as it is a resource that is always available for roughly the same price. However, in some places, the number of parking spots or the number of connectors is a limiting factor to how many customers a charging station can serve. Effectively, parking places, not electricity, are becoming a scarce resource in many charging stations. For this reason, some charging station operators resort to charging customers maximal tariff when the charging is finished, but the user has not yet unplugged the vehicle, and some, such as EVgo, have reservation systems for charging stations, as discussed in Chapter 3.

In this thesis, we propose an additional method of pricing EV charging. In conclusions of Chapter 3, we suggested a reservation system for fast charging stations as a solution for EV drivers to reach their destinations faster by planning and reserving their charging stops ahead of time, ignoring the problem of pricing the reservations. So, instead of pricing per kWh or hour per minute, we develop a pricing scheme that sets reservation prices, including charging. For reasons discussed in Chapter 3, namely, the benefit of full reservations for EV drivers in the fast-charging scenario, we focus on the pricing of fast-charging reservations. We focus on dynamic pricing of reservations in the sense of

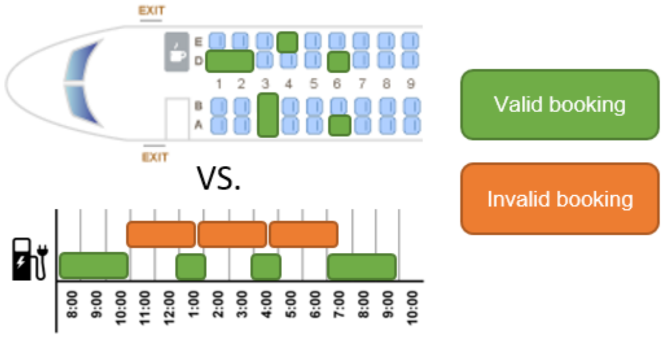


**Figure 2.11:** Difference between airline pricing and charging station pricing. Green rectangles show valid bookings. Seat bookings do not significantly affect bookings of other seats (except for large group bookings). On the other hand, a booking of short charging sessions and 1:00 and 4:00 block bookings of longer charging sessions are shown in red.

real-time pricing, that is, pricing that dynamically changes based on the system's current state.

While real-time pricing methods can have questionable effectiveness in terms of achieved effects without excessively large price changes [Cro08] when used without user-end technologies to manage the added complexity, we have already assumed that EV drivers will use sophisticated driver assistant systems that can recommend multiple options in terms of travel time. Such solutions exist [CJM24] and allow the users to plan trips with charging and offer multi-criterial tradeoffs between price and travel time. Equipped with such systems, users can enjoy the benefits of increased system efficiency without dealing with the complexity of real-time pricing.

### 2.3.2 Related Work

Dynamic pricing [AE08; MV99] is a technique to balance the load in various domains. It is studied in economics, revenue management, or supply chain management. In the field of smart mobility, where the system cannot be controlled centrally, dynamic pricing was proposed to improve the efficiency of taxi systems [Gan+13; GAM15] or power grid management for electromobility [Dei+11; Hay+15; VA16], balancing power load, power quality, and other grid-related metrics. These fields recognize dynamic pricing as a critical factor influencing buyers' behavior.

Dynamic pricing of charging services is a method that can potentially provide a cheap and robust alternative to expensive upgrades of our current infrastructure. However, the proposed applications focus on dynamic pricing primarily toward optimizing social welfare. Yet, in real-world situations, prospective charging station operators are often privately owned and, as such, not strongly incentivized to improve social welfare. Instead, private investors are concerned with the costs of installing and providing charging services and their financial returns[6]. Modeling charging station operators as self-interested agents instead of cooperative groups seems closer to future reality in many places worldwide.

*For this reason, in this work, we focus on maximizing the revenue of charging station operators and show that this can be done while still improving social welfare.* To this end, we propose MDP based dynamic pricing strategy for a self-interested charging station operator that aims to maximize its revenue while improving utilization of the limited available grid resources.

### Dynamic Pricing

From the revenue-maximizing standpoint, dynamic pricing can be viewed as a method of extracting additional revenue from services that are in high demand and that each buyer values differently, such as hotel rooms [Gu97] or airline tickets [SSL99]. For dynamic pricing of EV charging reservations to improve the system efficiency, it is therefore important that the variable costs in the power generation and distribution are projected onto the charging station operators, as exemplified in Figure 2.10. Only then can their self-interested, revenue-maximizing behavior improve the system as a whole.

For airfares and hotel rooms, the price changes based on the expected demand throughout the season, existing bookings, and the customer segment (business or tourist). Services such as airfares and hotel rooms have strict expiration deadlines, that is, the departure of the airplane and the arrival of a given booking day. Similarly, the EV charging resources in a given time window expire if there are no vehicles to use them.

With such a type of expiring services, the goal is to sell the available service capacity for a profit under the constraints given by their expiration and fluctuations in demand. Unused service capacity is a wasted profit opportunity. Maximizing revenue from these expiring services is the topic of revenue management [MV99].

Literature on revenue management, including dynamic pricing of hotel rooms and airline tickets, is classically not primarily concerned with competition of the service providers nor with competition among customers [CCX06; KJ17]. The primary concern is the profitable utilization of the sold resource (rooms, tickets). We also do not consider the game-theoretic aspects of the dynamic pricing problem of EV charging, neither on the supply side, such as presented by Xiong et al. [Xio+16], nor on the demand side (such as in [Ete+17]). [Ger+13] uses mechanism design to propose 2-sided market that takes driver preferences and opportunity costs of charging stations based on historical data to generate allocations and payment profiles in. [Ete+17] models the allocation of charging capacity as congestion game of greedy drivers and shows that the system can reach an equilibrium state with selected pricing strategies. In the simulation section of the paper, authors compare two pricing strategies and make pricing suggestions based on the number of customers in the system.

Main drawback of the competitive approaches is that they are tractable only in small instances or oversimplified versions of the problem. In our view, the deployment of EVs

[6] From the report "An Industry Study on Electric Vehicle Adoption in Hong Kong" by the Hong Kong Productivity Council (2014): www.hkpc.org/images/stories/corp_info/hkpc_pub/evstudyreport.pdf

is still limited, and there is no data to validate competitive aspects of dynamic pricing of charging services. Additionally, charging services from an e.g., fast highway charging station are not fully substitutable with another charging station from a competing operator as they most likely occopy different physical locations. While competition between charging stations might become important in the future, we focus on the problem of optimal pricing in an environment where other charging stations operate in virtual monopoly. Thus, we adopt a similar approach to the early work on airline revenue management that disregards the competitive aspects of pricing and focuses on the factors deemed more important at the time.

The common approach to the pricing of airline tickets, extensively studied, e.g., in [CCX06], is to determine the number of tickets that can be sold for a particular price to maximize the expected revenue. These decisions require the modeling of customers, which is based on parameters such as price sensitivity, seasonality of demand, and others. The airline can effectively accept and reject bookings by opening and closing the different fare classes on a flight. The optimal rule for accepting or rejecting air ticket bookings is as follows: *Is the profit from accepting a given booking higher than the expected profit from the seat that we could get later? If it is, confirm the booking (i.e., open the fare class). If it is lower, reject this booking (i.e., close the fare class).* Moreover, each *accepted* booking can influence following bookings as customers booking later can potentially not book seats at the same price. Additionally, if we include connecting flights and group bookings, pricing decisions can affect and be affected by other subsequent pricing processes. Because of this snowball effect, the complexity of the air ticket pricing problem (and similarly the charging service pricing problem) is, in general, intractable [MV99]. For this reason, the problem is often simplified (e.g., by assuming independence of demand between booking classes, allowing no cancellations, etc.) so that a solution can be found computationally.

An important distinction between airline pricing and charging station pricing is the coupling of the bookings in the case of charging services, which is not present when selling airline tickets. In the case of airline tickets, it is not particularly important which seat (in a given class) was sold as the booking of a single seat does not block the booking of surrounding seats (see Figure 2.11). However, booking a single time window can influence time windows before and after it, as it can block other charging opportunities.

We focus on developing a pricing strategy for a seller in an established market and not the design of the market. Nor do we consider optimization of the products themselves, such as optimization of the charging station placement [LLC14], proposal of airline routes [LSS14], rebalancing of rental fleets [Nou+15] or routing of EV vehicles [Sac+11; Bau+13].

In this thesis, we approach the charging station pricing strategy in a way similar to revenue management deals with the pricing of airline tickets. Both problems focus on expiring services where earlier transactions affect transactions that follow.

**Pricing of EV Charging**

In the electromobility domain, pricing is a relatively novel problem. Before the idea of demand-side management was part of the mainstream in EV charging research, most of the works on charging for electromobility focused on optimizing charging station placement [Fra+11; Wan+10; He+13; Xio+18]. Such approaches are only a seeming remedy in a changing environment where charging station placement is no longer optimal in the new environment. On the other hand, dynamic pricing of the charging services and its application to load balancing is robust to the dynamically changing situation in the infrastructure, demand, and energy costs.

Problems of EV charging described in the literature differ in their use and can be broadly divided into two categories [RRB15; Xia+16]. First is many-to-many, multiple EVs and multiple “fast” charging stations where users want to top up their battery quickly before continuing with their trip (as described, e.g., in [Che+16; Xio+16]). Second is many-to-one “destination” charging station where users are at the location for some other purpose and are not in a rush to leave, such as workplaces or homes (case covered for example by [RFH14; LNG15]).

In the first case, problems focus on limiting congestion, stabilizing the grid and maximizing utility. [Bay+13] focuses on the problem of stabilizing the grid with fast chargers and show substantial improvements in utilization of power when some customers are simply rerouted to other charging stations. [De +13] approaches the problem from a completely opposite side and avoids congestion by having EV drivers broadcast their intended charging locations who in turn update their foretasted waiting times for other drivers. [Xio+16] is an example of a work that focuses on pricing to improve social welfare in the model. The pricing problem proposed considered EV drivers’ travel patterns and self-interested charging behavior. The problem can be seen as a variation on sequential posted pricing [Cha+10] for charging stations. The proposed pricing minimizes the social cost of electromobility, but like many other articles on the topic, it would be hard to put to implement in practice. In [BMD12], the authors use the queuing theory methodology to study the performance of charging stations by dynamically changing the prices so that the overall throughput is maximized and the waiting time is minimized. The authors in [Bha+16] use the game theory methodology in general, specifically the Vickrey–Clarke–Groves (VCG) auction mechanism, to specify prices for charging sessions such that the social welfare function is maximized.

Problems of the second kind focus on the coordination of charging of multiple EVs without considering the traveling customer. [RFH14] coordinates charging using system of energy packets that ensures cars are charged in time without exceeding the grid limitations. [BG12] extends the coordination of charging with multi-modal routing to destination chargers in cities and [JSG14] focuses on the utilization of renewables in EV charging. Works such as [ARK12; AKR14] focus solely on balancing the grid using smart-grid approaches.

An additional branch of literature deals with vehicle-2-grid scenarios where EVs can store energy at one place or time and return it back to the grid at a later time [GA16; Lut+14].

Another line of research investigates pricing and allocation of charging resources in a centralized manner. Works such as [Zou+16; Xia+16; Ger+13] study allocation of EV charging resources in auctions from game-theoretical point of view. Authors present centralized as well as decentralized auctions and introduce mechanisms that have many useful properties, such as incentive compatibility or being close to social optimum. However, these models usually assume universal adoption by all EV drivers or CS operators, often ignoring the grid state. The current state of the EV charging infrastructure in Europe is fragmented, with many different owners and operators who often join one or more roaming networks that provide basic information sharing and payment processing capability.

Most of the works above consider charging station pricing as a tool for shifting driver demand in time and space to improve some global statistic, disregarding goals of the charging station owners. One of the exceptions is [Rig+13] that approaches the problem as a multi-agent system where charging stations use either simple congestion pricing formula to optimize combined charging station revenue and driver utility or central allocation mechnism to allocate charging. This differes from our approach to pricing, where we assume charging station to be a profit-maximizing agent who has only limited concern for the efficiency of the system as a whole.

More recent works, such as [DWH19; Shi+19], use MDP model of dynamic pricing in reinforcement learning. This approach has been most recently extended also to dynamic pricing of EV charging [Cui+23]. We systematically avoid the use of reinforcement learning in our work due to the large compute requirements to train the model and currently low availability of data to train the model on. Furthermore, while we ignore the competitive aspects of the pricing problem, we assume common refitting of our model to the latest data, which would be costly in the reinforcement learning approach.

[KZZ19] represents a mathematical optimization approach to dynamic pricing of EV charging. The main disadvantage of this approach is that it either does not scale well to large instances, or requires severe limitations on the problem formulation.

From the perspective of pricing realization, there are different types of contributions in the literature, categorized into offline and online approaches. The former method specifies charging prices for extended time periods (e.g., one day) based on some information related to the projected EV charging demand, such as the number of EVs to be charged during this period, their required charging amount, etc. On the other hand, online approaches specify charging prices for short periods and often update them. This is the approach that we are adopting in this thesis. In this respect, several contributions can be found in the literature. Like our approach, in [KKC16] the authors assume that the charging prices change dynamically, and the EV users are offered different prices on a session basis. The EV users can either accept or reject the proposed price. The authors also suggest that the CS operator has to pay some penalties in case the waiting time of the EV users exceeds a certain threshold. The proposed scheduling algorithm has the main objective of maximizing the profit of the CS operators. In [LR17], the authors also consider the problem of optimally allocating the charging stations' capacity to maximize the CS operators' profit. To this end, they propose a framework that changes the price of charging dynamically so that the EV users can either accept or reject the offered price. Consequently, the framework can also be used to minimize the number of rejections by EV users.

In this paper, we consider the dynamic pricing of EV charging using online and session-based techniques. However, unlike the contributions above, the underlying problem under study is formulated using the Markov Decision Process (MDP) methodology. Thanks to this, we manage to solve much larger problem instances thanks to the proposed MCTS method. To the best of our knowledge, this is the first attempt to apply MCTS to the dynamic pricing of EV charging.

## 2.4 Probability in Demand Modelling

In this section, we will introduce the basic concepts of probability theory that will be used in the rest of the thesis. First, we describe the properties of distributions we use to model customer behavior, followed by an introduction of the coin-flipping process and its limit as a model of customer arrivals. Finally, we will introduce the concept of demand function and its relation to the customer valuation distribution.

The goal of this introduction is to refresh the concepts to a reader who is familiar with probability theory, but maybe has not used the specific concepts for some time. For this reason, we will focus on the practical implications of the concepts and how we apply them and not on building up the theory behind them from the ground up. Additionally, since the naming of the concepts we use appears almost intentionally confusing, we will try to illustrate the theory on the simple experiment of repeated coin-flipping.

For a more thorough introduction that includes theoretical background, we recommend Pishro-Nik [Pis14], an online textbook on probability theory, and Gallager [Gal13], which focuses on random processes.

### 2.4.1 Coin Flipping

A simple coin flip is the most fundamental and intuitive experiment in probability theory. Also called a *Bernoulli trial*, it is a random experiment with two possible outcomes, heads or tails, and can be used to model any random experiment with two possible outcomes, gender of a newborn child, success or failure on an exam or, as we use it, arrival of a customer in an interval.

The probability of success (meaning heads, denoted 1) is denoted by $p$ and the probability of failure (meaning tails, denote 0) by $q = 1 - p$. For a typical coin, the probability of heads and tails is assumed equal, $p = q = 0.5$, but we assume a biased coin with $0 \leq p \leq 1$. The Probability Mass Function (PMF) of the random variable $X$ that

| Discrete Time | | Continous Time |
|---|---|---|
| Coin flipping, success probability $p$ | $\leftrightarrow$ | Poisson process with rate $\lambda$ |
| geometric($p$) | $\leftrightarrow$ | exponential($\lambda$) |
| $N_1 = \#$ of flips to success | | $T_1/\lambda =$ time to first point |
| $P\{N_1 = 1+k\} = p(1-p)^k$ | | $T_1$ has density $f_1(\tau) = \mathrm{e}^{-\tau}$ |
| for $k \in \{0,1,2\}$ | | for $\tau \geq 0$ |
| 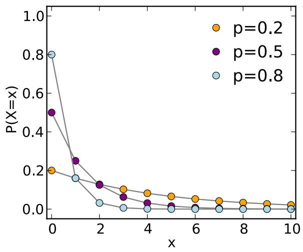 | | 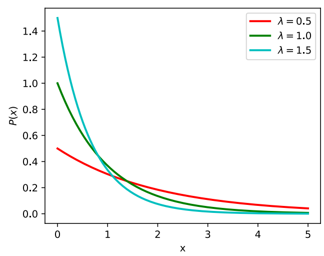 |
| binomial($n,p$) | $\leftrightarrow$ | Poisson($\lambda\delta$) |
| $X = \#$ of success. flips in $n$ | | $X = \#$ of points in $[\tau, \tau+\delta]$ |
| $P\{X = k\} = \binom{n}{k} p^k (1-p)^{n-k}$ | | $P\{X = k\} = \mathrm{e}^{-\lambda\delta}(\lambda\delta)^k/k!$ |
| for $k \in \{0,1,2\}$ | | for $k \in \{0,1,2\}$ |
| 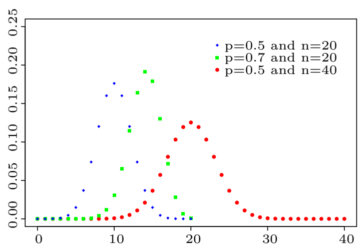 | | 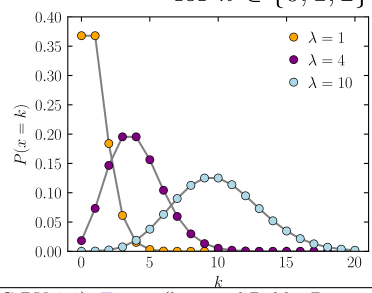 |

*Figures by EvgSkv (exponential, CC0), Skbkekas (geometric, Poisson, CC BY 3.0), Tayste (binomial,Public Domain).

**Table 2.3:** Overview of coin flipping concepts, based on [Pol14]*.

represents the outcome of a single coin flip is given by *Bernoulli distribution*:

$$P_X(x) = \begin{cases} 1-p & \text{if } x = 0 \\ p & \text{if } x = 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.1}$$

If we repeat the coin flip experiment $n$ times, we get a sequence of $n$ random variables $X_1, X_2, \ldots, X_n$. This sequence is called a *Bernoulli process* and is a fundamental building block of probability theory. This is a foundation of our customer arrival model, where each coin flip corresponds to a time interval with success representing a customer arrival in a given interval.

We can ask two questions about the number of successes in the sequence that result in two different random variables and associated probability distributions. Asking about the number of flips until the first success leads us to a Geometric distribution, and asking about the number of successes in a fixed number of flips leads us to a Binomial distribution.

### Geometric Distribution

Observing the repeated coin flips in the Bernoulli process, we can ask about the number of flips until we observe the first success. The random variable $X$ that represents this number is said to have *geometric distribution* with parameter $p$. The PMF of the geometric distribution is given by:

$$P_X(k) = \begin{cases} p(1-p)^{k-1} & \text{if } k \in \{1, 2, 3 \ldots\} \\ 0 & \text{otherwise} \end{cases}$$

Note that geometric distribution is supported on set $\{1, 2, 3, \ldots\}$ since it takes at least one coin flip to get a success. Alternatively, we can define geometric random variable $Y$ as the number of failures before the first success in a sequence of independent Bernoulli trials. The PMF of $Y$ is

$$P_Y(k) = \begin{cases} p(1-p)^{k} & \text{if } k \in \{0, 1, 2, \ldots\} \\ 0 & \text{otherwise} \end{cases}$$

with support $\{0, 1, 2, \ldots\}$ since success on first try means zero failures. In the demand model in Chapter 4, we use this variant of the geometric distribution, and it is also the variant shown in Table 2.3.

### Binomial Distribution

By defining $X$ as a number of successes in $n$ coin flips with the probability of success $p$, we get a *binomial distribution*. The PMF of $X$ is then given by:

$$P_X(k) = \begin{cases} \binom{n}{k} p^k (1-p)^{n-k} & \text{if } k \in \{0, 1, 2, \ldots\} \\ 0 & \text{otherwise} \end{cases}$$

Note that the binomial distribution is the distribution of the sum of $n$ independent Bernoulli random variables with parameter $p$.

The two distributions, geometric and binomial, describe two different aspects of the repeated coin-flipping process. The geometric distribution descriebs a number of intervals with no customers showing up (between intervals where customers did show up). The binomial distribution describes the number of customers arriving in a timespan spanning $n$ intervals.

### 2.4.2 Very Fast Coin Flipping

A Bernoulli process is a discrete set of events, successes or failures, that are assigned positions in a sequence. If we assign time to these events, e.g. flipping a coin every minute, we can use it to model events in time. In our customer arrival example, we model e.g. a minute with a successful trial as the arrival of a customer. An apparent drawback of this model is that we can't have more than one customer in one minute, which is bound to happen sooner or later in a real-life scenario. To address this shortcoming, we can increase the number of flips in a minute. To keep the expected number of successes per minute the same as before, we also need to decrease the probability of success $p$ proportionally. In a nutshell, by repeating this process, in a limit, we get a *Poisson point process*, a continuous-time stochastic process that models events in time.

Specifically, if we make $n$ coin flips in a minute with a probability of success $p$, we can increase $n$ and keep the expected number of successes per minute the same by keeping the product $\lambda = pn$ constant. If we set $n'$ to a higher value, the new process must have $p' = \frac{\lambda}{n'}$ for the expected number of successes to remain constant. In the limit, as $n'$ goes to infinity, we get a Poisson point process with parameter $\lambda$. $\lambda$ is called the intensity or rate of a Poisson process and by this informal derivation, it can be understood as the expected number of successes in the interval under consideration, e.g., a minute in this case. The following formal definition is by Pollard [Pol14, Chapter 10]:

> **Definition 2.1** (Poisson point process via Poisson distribution)**.** A *Poisson process* with fixed rate $\lambda \in (0, \infty)$ is a random mechanism that generates "points" strung along $[0, \infty)$ in such a way that
>
> 1. the number of points in any interval of length $\tau$ is a random variable with a Poisson distribution with parameter $\lambda t$.
>
> 2. the number of points landing in disjoint (= non-overlapping) intervals are independent random variables.

This definition showcases how the Poisson distribution appears in the Poisson process, and it has the same relationship as the connection between the Bernoulli process and the Binomial distribution. This also introduces Poisson process as a point process, stochastic object that consists of points randomly strewn along the real line.

Alternatively, Poisson process can be viewed as an example of a *counting process*, a stochastic process that counts the number of events in a given time interval, see e.g. the definition from [Pis14, Definition 11.1]:

> **Definition 2.2** (Counting process)**.** A random process $\{N(\tau), \tau \in [0, \infty)\}$ is said to be a counting process if $N(\tau)$ is the number of events that occured from time 0 up to and including time $\tau$. For counting process, we assume:
>
> 1. $N(0) = 0$
> 2. $N(\tau) \in \{0, 1, 2, \ldots\}$, for all $\tau \in [0, \infty)$
> 3. for $0 \leq \sigma < \tau$, $N(\tau) - N(\sigma)$ is the number of events that occur in the interval $(\sigma, \tau]$

The following definition defines Poisson process as a counting process and highlights the similarity with the Bernoulli distribution (compare the defintion below and Equation (2.1)):

> **Definition 2.3** (Poisson counting process via Bernoulli-like distribution)**.** A *Poisson process* with fixed rate $\lambda \in (0, \infty)$ is counting process $\{N(\tau), \tau \in [0, \infty)\}$ if the following conditions hold:
>
> 1. $N(0) = 0$
> 2. $N(\tau)$ has *independent* and *stationary* increments
> 3. for some small interval of length $\delta$, we have
>
> $$P\{N(\delta) = 0\} = 1 - \lambda\delta + O(\delta^2) \tag{2.2}$$
> $$P\{N(\delta) = 1\} = \lambda\delta + O(\delta^2) \tag{2.3}$$
> $$P\{N(\delta) \geq 2\} = O(\delta^2) \tag{2.4}$$

For description of independent and stationary increments, see e.g. [Pis14, Section 11.1.1].

Now, for the "ultra-fast coin flipping" (Poisson) process, we can consider the same two questions as we posed to our "regular" coin flipping process. First, what is the length of an interval between two successive points? This leads us to an exponential distribution. Second, what is the number of points in a fixed interval? Answer to this question results in a Poisson distribution.

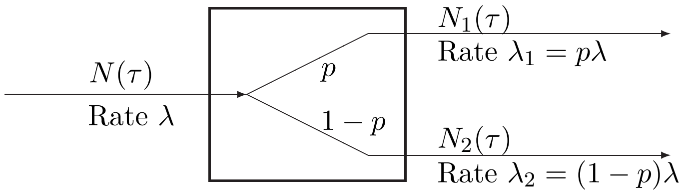


**Figure 2.12:** Splitting of a Poisson process into subprocesses by assigning each event from the original process into one of the subprocess. Adapted from [Gal13].

### Exponential Distribution

Exponential distribution is the continuous counterpart of the geometric distribution. It is parametrized by a single parameter $\lambda$ and the Probability Density Function (PDF) of the exponential distribution is given by:

$$f_X(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

The time between two points in the Poisson process with parameter $\lambda$ has an exponential distribution with parameter $\lambda$. As we described the Poisson process as a limit of a sequence of Bernoulli processes, we can accordingly derive the exponential distribution as a limit of a sequence of geometric distributions, through a similar mechanism. We do this by identifying the constant $\lambda$ with the product $pn$ for some $n$, setting $x = \frac{k}{n}$ and taking the limit $n \to \infty$.

### Poisson Distribution

Poisson distribution is often used to approximate the number of customer arrivals in a fixed time interval. For example, if we know that on average 10 customers arrive in an hour, we can model the number of customers arriving in a minute with a Poisson distribution with parameter $\lambda = 10/60 = 1/6$. The PMF of the Poisson distribution is given by:

$$P_X(k) = \begin{cases} \frac{e^{-\lambda}\lambda^k}{k!} & \text{if } k \in \{0, 1, 2, \ldots\} \\ 0 & \text{otherwise} \end{cases}$$

By our construction of the Poisson process, it is unsurprising that also the Poisson distribution can be viewed as a limit of Binomial distributions. Specifically, let $X$ be a random variable distributed according to binomial$(n, p)$ and assume that $n$ is large and $p$ is small and, importantly, $\lambda = np$ is constant. Then, the PMF of $X$ can be approximated by the PMF of the Poisson distribution with parameter $\lambda$. This is useful since the Poisson distribution is much easier to work with than the Binomial distribution due to the absence of the binomial coefficient.

### Combining Poisson Processes

In this thesis, we use the Poisson process to model arrivals of charging requests on a real timeline. However, our MDP models require arrivals in discrete timesteps, which could be modelled using Bernoulli process. The reason why we use the Poisson process and then discretize it, instead of using the Bernoulli process directly, is because of it's useful properties. Specifically, Poisson processes can be combined and split and remain Poisson processes, which does not hold for the Bernoulli process.

**Merging Poisson processes** Suppose $N_1(\tau)$ and $N_2(\tau)$ are two independent Poisson counting processes with rates $\lambda_1$ and $\lambda_2$. Then, the process $N(\tau) = N_1(\tau) + N_2(\tau)$ is also a Poisson counting process with rate $\lambda_1 + \lambda_2$ for all $\tau \in [0, \infty)$. This is a useful property that we use in the demand model in Chapter 5. We sketch the proof based on the counting process Definition 2.3.

Since $N_1(\tau)$ and $N_2(\tau)$ are independent, the increments of $N(\tau)$ will retain the properties of stationary and independent increments, as well as the property that $N(\tau) = 0$. To show that $P\{N(\delta) = 0\} = 1 - \lambda\delta + O(\delta^2)$, we use the independence of $N_1(\tau)$ and $N_2(\tau)$:

$$\begin{aligned} P\{N(\delta) = 0\} &= P\{N_1(\delta) = 0\}P\{N_2(\delta) = 0\} \\ &= (1 - \lambda_1\delta + O(\delta^2))(1 - \lambda_2\delta + O(\delta^2)) \\ &= 1 - (\lambda_1 + \lambda_2)\delta + O(\delta^2) \end{aligned}$$

Similarly, we show the property for $P\{N(\delta) = 1\}$ and $P\{N(\delta) \geq 2\}$. Different proofs based on different definitions of Poisson process can be found in [Gal13, Section 2.3]. The textbook also make an interesting observation that "when many independent counting processes (not necessarily Poisson) are added together, the sum process often tends to be approximately Poisson if the individual processes have small rates compared to the sum."

**Splitting Poisson process** We can also subdivide the combined process $N(\tau)$ to extract its constituent parts, $N_1(\tau)$ and $N_2(\tau)$. The splitting is illustrated in Figure 2.12. At each event in $N(\tau)$, we flip a coin with probability $p = \frac{\lambda_1}{\lambda_1+\lambda_2}$ to decide whether the event is assigned to $N_1(\tau)$ or $N_2(\tau)$. The proof that processes created by this procedure are indeed independent Poisson processes is sliglty more involved and can be also found in [Gal13, Section 2.3].

**Compound Poisson process** The splitting procedure that labels each event in the original process can be thought of as a Bernoulli process with parameter $p$, with each trial in the process associted with one event in the original Poisson process. If, instead of labelling the events, we used the outcomes of the Bernoulli process to increase the count of events in $N(\tau)$ by either 1 or 2, we would get a special type of Poisson process called *compound Poisson process*. The compound Poisson process is a Poisson counting process with "jumps", and the splitting probablities $p = \frac{\lambda_1}{\lambda_1+\lambda_2}$ and $1-p$ form the *jump size distribution*. For this reason, when we want to refer to a Poisson process that was created by combining multiple Poisson processes, we will sometime refer to it as a compound Poisson process.

While we show the splitting and combining of Poisson processes with only two subprocesses, the concepts hold for any fixed number of processes. This can be shown by repeated addition and splitting of processes.

### 2.4.3 Price Elasticity, Demand Function and Customer Valuation

In this thesis, we use two models of customer response to a price. First, we use an aggregate model of customer response that we called *price-demand function*. Second, we use an agent-based model where we assume that each customer has a *valuation* for the product and decides to buy the product if the price is lower than the valuation. We will show that these two models are dual to each other, meaning that they are equivalent in the sense that they give the same predictions about the customer behavior in our expectation-maximizing setting.

*Price elasticity of demand* is a term in economics. It mesures the local variation in demand for a product in response to a variation in price. It is a unitless value defined as the percentage change in quantity demanded in response to a percentual change in price. For an informal description of price elasticity of demand, see [Gal15], and [PPM08] for a textbook definition. For our discussion, it is sufficient to know that demand is said to be *elastic* if customers have a strong response to price changes, and *inelastic* if the response is weak.

*Demand function* is another term from economics [SF18, p. 46]. In this thesis, we refer to it as *price-demand function* to distinguish it from other demand-related objects. Demand function expresses the volume of products customers are willing to buy as a function of price. In our work, we use the price-demand function to model the aggregate response of customers to the price. The price-demand function $\mathcal{E}(h)$ is the probability that arriving customer will buy the requested product given that the price is $h$. Examples of price-demand functions are in Figure 4.5.

*Customer valuations* are the second way we can model the customer response to price. The assertion is that each customer has some internal budget or valuation $b$ for a product they are requesting. The budgets are distributed according to some probabilistic distribution $\beta$. If the price $h$ is lower than the budget $b$, the customer accepts the price and buys the product. Otherwise, the customer does not buy the product.

This simple agent-based model of customer behavior is equivalent to the price demand function. Consider the probability density function $\beta(h)$ of budgets in a population of customers. The price-demand function $\mathcal{E}(h)$ is then obtained as the integral of the budget distribution over the price, i.e. $\mathcal{E}(h) = \int_h^\infty \beta(h)dh$. This is a simple application of the cumulative distribution function of the budget distribution. Therefore, $\mathcal{E}(h) = 1 - F_\beta(h)$, where $F_\beta(h)$ is the cumulative distribution function of the budget distribution.

## 2.5 Decision Making in the Face of Unceartainty

In this section, we will introduce the basic concepts of MDPs, as well as some of the algorithms used to solve them. We only provide a basic overvivew of the concepts and algorithms, for a more thorough introduction, we recommend Mausam and Kolobov [MK16] whom we leverage throughout this section.

### 2.5.1 Markov Decision Process

Markov decision process is a broad mathematical framework for modelling decision making in stochastic environment. It's widely used in many fields, including economics, probabilistic planning, reinforcement learning, and operations research. The characterics of a domain that make it suitable for modelling as a MDP are:

- Sequential decision making - MDPs are useful when there is a sequence of actions to be made. "MDPs are a principled mechanism to trade off multiple competing objectives"[MK16], especially when reaching the different objectives has different costs and rewards.
- Fully observable states - the decision making agent fully knows the state of the environment. Partial observable MDPs are a generalization of MDPs that relax this assumption.
- Fair chance - the environment is stochastic, but neither the decision making agent, nor anyone else can influence the outcome of the stochastic process. The distribution that the outcome is generated from depends only on the agents' decision and the current state of the environment. This means that MDPs are not suitable for modelling adversarial environments, which is a common use case in game theory.

Next, we introduce one possible elementary definition of a MDP:

**Definition 2.4** (Markov Decision Process)**.** A *Markov Decision Process* is a tuple $(S, A, \mathcal{T}, R)$ where

- $S$ is a finite set of states,
- $A$ is a finite set of actions,
- $\mathcal{T}$ is a transition function $\mathcal{T}(s, a, s') = P(s'|s, a)$ that gives the probability of transitioning to state $s'$ when taking action $a$ in state $s$,
- $R$ is bounded a reward function $R(s, a, s')$ that gives the reward for taking action $a$ in state $s$ and transitioning to state $s'$,
- $\gamma$ is a discount factor $\gamma \in [0, 1)$.

In the literature, authors usually introduce MDPs in a similar fashion, specifying the set of states, actions, transition function and reward function. A solution of a MDP is a policy $\pi$, a function that maps states to actions.

The MDP develops from some initial state $s_0$ through actions taken by the agent. The agent chooses an action $a$ in state $s$ according to the policy $\pi$, which results in a transition to a new state $s'$ with probability $\mathcal{T}(s, a, s')$. The agent receives a reward $R(s, a, s')$ for the transition. The process then continues in the new state $s'$. The goal of the decision making agent is to find a policy that maximizes the expected reward, which is defined as the sum of consecutive rewards (indexed by $t$) discounted by the discount factor $\gamma$:

$$\mathbb{E}_{\mathcal{T}}\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(s_t), s_{t+1})\right] \tag{2.5}$$

Maximizing this expected linear additive utility function is a common choice for modelling the goal of the decision making agent. Applying the expectation to the sum results in the well known, recursive, Bellman equation that introduces the concept of value function of a policy:

$$V^\pi(s) = \sum_{s' \in S} \mathcal{T}(s, \pi(s), s') \left[R(s, \pi(s), s') + \gamma V^\pi(s')\right]$$

The value function $V^\pi(s)$ is the expected reward of following policy $\pi$ from state $s$ onwards. The value function of a policy is useful since it allows us to compare policies. A policy $\pi$ is better than a policy $\pi'$ if $V^\pi(s) \geq V^{\pi'}(s)$ for all $s \in S$. A policy $\pi^*$ is optimal if it is as good as or better than all other policies.

Under certain conditions, the optimal policy is guaranteed to exist. For example, in MDP defined by Definition 2.4 and using the linear additive utility, the optimal policy always exists and can be found by solving the Bellman equations.

However, Definition 2.4 is not the only way to define a MDP. In fact, there is a taxonomy of MDP definitions. The Definition 2.4 presents an *infinite horizon* MDP with a discount factor $\gamma$. The discount factor is used to model the fact that the agent prefers to receive rewards sooner rather than later, and setting $\gamma < 0$ guarantees convergence in Equation (2.5).

Another MDP variant is a *finite horizon* MDP, where the agent has a finite number of steps before reaching a terminal state. In this type of MDP, discount factor $\gamma$ can be set to 1 since the sum Equation (2.5) is finite. In this thesis, we deal primarily with finite horizon MDPs.

Both infinite and finite horizon MDPs are both subclasses of *stochastic shortest-path* MDPs, where the goal is to find the shortest path to some goal state. Both finite horizon and infinite horizon MDPs can be transformed to stochastic shortest-path MDPs by appropriately adding a goal state and modifying the reward function [MK16, Chapter 2.4.2]. This means that most methods for solving MDPs can be applied to both finite and infinite horizon MDPs unless they rely on specific properties of given class.

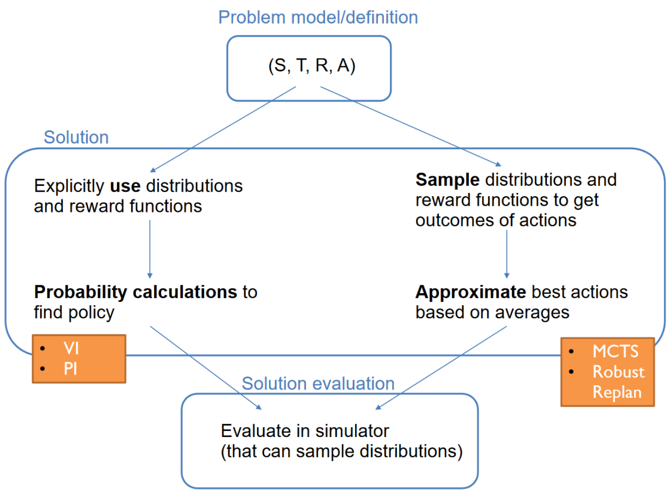


**Figure 2.13:** General approaches to solving MDPs. Policy Iteration (PI) and Value Iteration (VI) are offline algorithms that require explicit description of the environment. Monte Carlo Tree Search (MCTS) and Robust Replanning are online algorithms that only require a simulator of the environment.

### 2.5.2 Solving and Evaluating MDPs

Solving a MDP means finding an optimal policy $\pi^*$ that maximizes the expected reward. Since there are many different ways to define a MDP, there are also many different ways to solve and evaluate them. Figure 2.13 shows a general overview of the approaches to solving MDPs.

Solvers such as value iteration or policy iteration are offline algorithms, meaning the bulk of computation to find a policy must be done before the policy can provide any actions. However, once the calculation of policy is done, obtaining actions from the solution is fast. Additionally, it may often be feasible to obtain an optimal policy for given MDP.

Additionally, methods on the left side of Figure 2.13 require the enumeration of all states and their memory requirements can therefore be prohibitive for large MDPs. This also requires the definition of a MDP to explicitly include all transition probabilities and rewards. However, in certain situations, this might be impossible or impractical.

On the other hand, online algorithms on the right hand side of Figure 2.13, such as MCTS or robust replanning, are able to provide actions in near real-time by using heuristics to approximate the optimal policy. These methods are often used in situations where the MDP is too large to be solved offline or when the explicit definition of the MDP is not available. Instead, these methods can utilize a sampling simulator of the environment that can often be much easier to define and implement. The drawback of these methods is that getting an action from the algorithm can be computationally expensive each time and the action might not be optimal.

### 2.5.3 Value Iteration

**Algorithm 1** Value Iteration

1: $n \leftarrow 0$
2: $\forall s$: $V_0(s) \leftarrow 0$ ▷ *arbitrarily chosen initialization*
3: **while** running **do**
4: $\quad n \leftarrow n + 1$
5: $\quad$ **for all** state $s \in S$ **do**
6: $\quad\quad V_n(s) \leftarrow \max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s')[R(s, a, s') + \gamma V_{n-1}(s')]$ ▷ *Bellman backup*
7: $\quad$ Break if $V_n$ has converged
8: **return** $V_n$ and greedy policy $\pi^{V_n}$

Value iteration is the best known and most widely used algorithm for solving MDPs. The algorithm falls into the class of offline algorithms, meaning that it requires significant computation before the agent can start acting in the environment. The algorithm also requires fully defined MDP, meaning that the transition function and reward function must be known.

The algorithm is shown in Algorithm 1. It is an iterative algorithm that starts with an arbitrary value function and improves it repeatedly using the Bellman equation. Usually, the stopping criterion is based on the difference of two consecutive value functions, $\max_{s \in S} |V_n(s) - V_{n+1}(s)|$, but in practice, fixed number of iterations is often used. The algorithm returns the value function $V_n$ and the greedy policy. Greedy policy is the policy that chooses the action with the highest value in each state with respect to the last value function $V_n$, applying the Bellman equation again:

$$\pi^{V_n} = \arg\max_{a \in A} \sum_{s' \in S} \mathcal{T}(s, a, s')[R(s, a, s') + \gamma V_{n-1}(s')]$$

Value iteration comes in many flavors with different tradeoffs on theoretical soundness, speed and memory requirements. This tradeoff is usually made in one of the three places. First, the memory requirements can be reduced and many theoretical guarantees voided by storing value function in only one array and updating it in place (in line 6). This is called asynchronous value iteration. Second, the method can converge much faster with the use of a heuristic initialization of the value function (line 2). Finally, the third class of improvements can be made to the asynchronous value iteration by ordering the states on line 5 in specific way, based on the structure of the MDP. This can greatly reduce the convergence time at the cost of adding some preprocessing steps.

The popularity of value iteration is due to its simplicity and the fact that it is usually guaranteed to converge to the optimal value function and policy from any initialization. Drawback of value iteration is that it must enumerate all states in each iteration, which can be prohibitively expensive in large MDPs. However, as an offline algorithm, once converged, the resulting policy can provide actions in constant time.

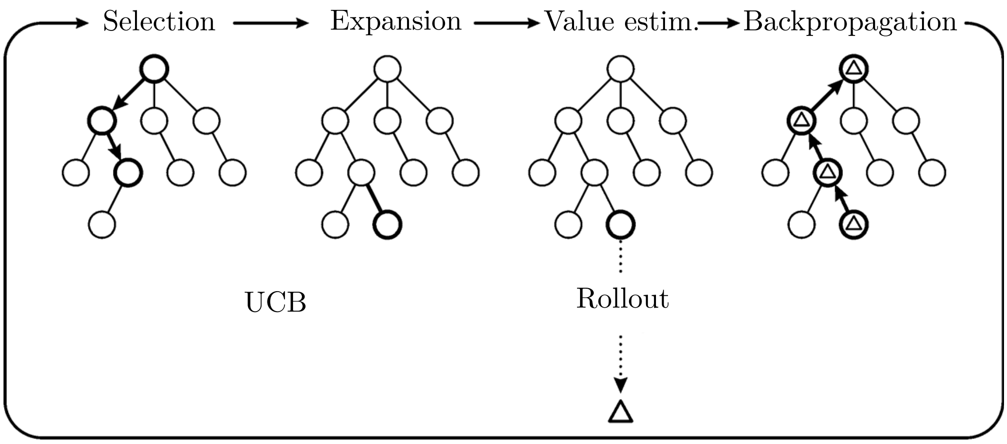


**Figure 2.14:** Overview of the stages of one iteration of MCTS algorithms. Adapted from [Bro+12].

## 2.5.4 Monte Carlo Tree Search

**Algorithm 2** General MCTS structure.

**INPUT:** Some MDP state $s$
**OUTPUT:** Action $a$ to be taken in $s$

Initialize root of the tree
**for all** $i = 0$ to iteration limit **do**
  Apply tree policy (traverse tree and expan a new leaf node) ▷ *Selection and expansion*
  Apply rollout policy (getting the approximate value of leaf node) ▷ *Value estimation*
  Backup values up the tree ▷ *Backpropagation*
**return** Most used action in the root

The main method we use to find actions in large MDPs is a variant of the MCTS algorithm. MCTS is a family of online heuristic search algorithms for solving problems on trees. This includes MDPs, partially observable MDPs and extensive form games.

Monte Carlo methods have been in use by the computer science community for a long time. As a method for a state-value estimation, the method became popular after the paper by Coulom [Cou06], which applies Monte Carlo evaluations to trees, coining the term “Monte Carlo Tree Search” and a paper by Kocsis and Szepesvári [KS06].

The crucial contribution of the latter was the application of the upper confidence bound (UCB1) algorithm [ACF02] initially used in bandit problems to handle the exploration/exploitation dilemma in the tree problem. The resulting UCT algorithm (Upper Confidence bounds for Trees) can quickly provide a good estimate of the optimal action while only evaluating small parts of the state space while converging to a theoretical optimum.

While the UCT algorithm has had the most success in playing games [Bro+12; Sil+16], it was originally ([KS06]) described as a method for solving finite-horizon MDPs.

The method relies on a simulator that samples results of actions applied to states. When using MCTS, the algorithm is rerun from every state where we need to make a decision. We show the high-level pseudocode of MCTS in Algorithm 2. Figure 2.14 illustrates the inner loop of the algorithm. In each run, first, the current state in the MDP is used to set up the root of the tree. Starting in the root, we use the tree policy based on UCB to select action (*selection*). Applying the action in the simulator, we get a new state. If the new state does not have a node in the tree already, we expand the tree by adding a new leaf node (*expansion*). If it does, we re-apply the tree policy to select the next action. This process is repeated until we reach a leaf node. In the leaf node, we perform *rollout*. Starting in the state corresponding to the leaf node, we apply the rollout policy to the simulator to quickly progress through the consecutive states of the MDP, only storing the collected rewards. Once we reach a terminal state or predefined depth, the rollout terminates and we use the collected rewards to estimate the value of the leaf node. This value is then propagated up the tree (*backpropagation*). After sufficient number of iterations, or after exhausting the time limit, the algorithm returns the most common action in the root node.

The description presented above applies to the use of MCTS in both MDPs and deterministic games. However, there are some differences in the implementation. In games, the tree nodes correspond to states of the game. In MDPs, we have two kinds of tree nodes, one for states and one for actions, and these nodes follow each other every time. The state nodes are used to store the number of visits to given states. The action nodes are children of state nodes and store the count of uses of said action in given state, as well as the estimate of value of this state-action pair. The reason for this difference is that in case of MDPs, actions don't have deterministic outcomes, so we can't store the value associated with the action in following state.

CHAPTER 3

# Reservations for EV Charging Stations

Reservations are a common mechanism for allocating capacity in many domains. In this chapter, we build on our analysis of the concept of charging reservations in electric mobility in Section 2.2. The requirements and practical limitations described in that section resulted in a classification of different reservation types for electric mobility and proof-of-concept of one specific reservation type (Section 2.2.1). Based on this analysis of potential reservation systems and their feasibility, in this chapter, we simulate what a large-scale deployment of different reservation types would mean for the electric mobility ecosystem. This analysis results in recommendations for charging station operators on appropriate reservation types for different circumstances, as well as recommendations for other stakeholders in the electric mobility ecosystem.

This chapter is primarily based on the workshop paper [Bas+19] and also uses [Bas+20]. These works were done in collaboration with researchers from the University of Passau and the University of Mannheim, who provided the domain expertise and headed the technical analysis of reservation concepts provided in Section 2.2. With co-authors from Czech Technical University, we have performed the simulation studies and analysis. The research presented in this chapter also leverages the results and know-how we obtained in the course of the ELECTRIFIC H2020 project.

A reservation system for EV charging can mitigate the challenges of widespread EV deployment and benefit the various involved stakeholders, that is, EV users, electric mobility service providers, charging station operators, and grid operators, as described in Section 2.1. Concepts of EV charging reservations were proposed for the charging of EVs [Liu+18; Mej+16; Orc+18]. However, these concepts are not interoperable and do not consider the requirements of all the involved stakeholders. Additionally, there are currently many technical limitations on possible reservation systems. Specifically, most of the existing proposed solutions consider only a subset of relevant stakeholders in specific scenarios, do not satisfy interoperability requirements, and ignore the technical limitations of existing infrastructure.

For the success of reservation in the electric mobility context, it is required to identify and consider the needs and requirements of the above-mentioned stakeholders in EV charging. From our investigation described in Section 2.1, for practical applications using existing, legacy systems, only an *ad-hoc* type of reservation can be instantiated and implemented. This is due to the fact that most of the existing charging hardware does not support more advanced features.

Ad-hoc reservation is a limited form of reservation that immediately reserves the charging connector of a given CS for a specific user. The corresponding connector of the CS gets unlocked when this user arrives at the station. Additionally, only one ad-hoc reservation can be active at a time on a specific connector of a given CS.

Newly deployed charging stations can be equipped with more advanced systems that allow for more advanced types of reservations, such as the system operated by EVgo[1].

Reservations may become an essential part of allocating EV charging resources to EV users, particularly when the number of EVs exceeds the rate at which CSs are being added. Ad-hoc reservations are the only type of reservation possibly supported by most of the existing charging infrastructure. Therefore it is essential to comprehend the benefits and limitations of ad-hoc reservations.

In this chapter, we aim to answer our first research question from Section 1.1 and determine sutiable reservation product for EV charging. Based on our analysis in Section 2.2, this boils down to the following two questions: *(1) do ad-hoc reservations provide sufficent benefit to the different stakeholders to justify their deployment,* and *(2) are full reservations able to provide more benefits than ad-hoc reservations?*

We do this in Section 3.1 that presents the simulation study of the ad-hoc reservations and full reservations. This analysis aims to compare the different reservation approaches and demonstrate their usefulness to different stakeholders. First, we determine the instances of reservations to use and the types of simulated driver agents to use them. We also develop metrics to quantify how much the different reservation types satisfy the stakeholder's goals (Section 3.1.1). Next, we evaluate the scenarios we have designed in an agent-based simulator with charging and mobility on data from the Netherlands and Germany. In Germany (Section 3.2), we focus on reservations for slow-charging stations, while in the Netherlands (Section 3.3), we emphasize the use case of highway fast-chargers.

The Section 3.1 is based on my primary contribution in [Bas+19; Bas+20]. I was responsible for developing the simulation scenarios, performing the simulations and analyzing the results. My co-author from Czech Technical University, was responsible for the simulation engine[2] modifications and vehicle route planning algorithms used by the simulated driver agents.

---

[1] The EVgo reservations were rolled out in 2021, after the publication of our work. This deployment validates the findings described in this chapter.

[2] AgentPolis project

Finally, in Section 3.4, we present the conclusions from our simulation study and recommendations for the types of reservations to use in the context of electric mobility. Our main results are that:

1. The positive effect of ad-hoc reservations on electric mobility stakeholders is questionable and, under certain conditions, can be detrimental.

2. In high-demand situations, full-fledged planned reservations can bring significant time savings to EV users and allow the suppliers of EV charging to plan their operations better.

These results thus support our work on dynamic pricing of EV charging that utilizes full reservations in Chapters 4 and 5.

## 3.1 Simulation-Based Validation of Reservation Systems

The field tests of the reservation system in Bavaria desribed in Section 2.2.3 demonstrate the *technical feasibility* of Uncertain Ad-Hoc reservations, not their *viability*. This would require a wide-scale deployment with a large number of users. Hence, to evaluate the concept's viability, we conducted large-scale simulations. First, we used a set of real locations of Bavarian public slow-charging CS infrastructure and simulated EV user demand to measure the impact of ad-hoc reservations and other reservation types on the EV user experience. Second, we evaluated the impact of reservations on the fast-charging experience of EV users on the Dutch highway network.

In our simulations, we considered the perspectives of both the supply side (charging station operator and electromobility service provider) and the demand side (EV users), using the goals described in Section 2.2.2. For EV users, we focused on two things: making it more convenient to charge by reducing uncertainty, waiting times, and search for CSs (Goal G.01 in Figure 2.2) and reducing the trip duration (G.02). In our simulations, we measured overall trip duration, including driving, waiting, and charging time, to see if the users met those goals. We also factored in the extra time it takes for detours to reach a charging spot.

To check the impact on the charging stations (G.03 in Figure 2.2), we measured the duration of their use in hours. In each scenario, we assumed all charging spots had the same cost and speeds. In one scenario, we considered the fast highway charging (Section 3.3), and in the other, we looked at slow, roadside charging (Section 3.2). This simplified the decision-making process for our simulated user agents, whom we then could design as optimizing travel time instead of a bi-criterial combination of price and time.

Even though we didn't directly address some of the goals (G.01 and G.04 in Figure 2.2), we assert that reservations naturally help with them. Reservations let users plan their trips better, avoiding unexpected delays. The basic reservation types (as described in Section 2.2.2) let users plan a little bit further ahead, while the more advanced, reliable reservation types help users avoid unexpected problems, such as finding all the spots taken. Reservations also improve the short-term planning capabilities of the CS operator (G.04) by providing knowledge about future charging sessions and allowing for their optimization.

### 3.1.1 Simulation Set-Up

To evaluate the benefits from different reservation types (see Section 2.2.2) for EV users and charging providers (i.e., CS operator and EMSP), we used the AgentPolis[3] multi-agent simulation framework [Jak+12]. AgentPolis is a fully agent-based simulator built upon a discrete-event simulation core. In our use case, all agents were deterministic in nature (e.g., the same state-of-charge (SoC) level, the same EV user behavior, etc.), and so were all simulation runs for a given set of input parameters.

In the simulation case study, we spawned multiple EV user agents and CS agents who interacted with each other over a 24-hour period. CS agents offered charging to the users with different reservation types in each scenario of the case study based on an FCFS charging reservation strategy. To reduce the complexity of the modeling, in this simulation, the CS operator concurrently represented the role of the EMSP (in reality, this is a common case).

We modeled two types of EV users: *naive* and *prudent*, two possible extremes on the spectrum of possible driver behavior. The naive user does not plan his/her trip with charging in mind and reacts only when his battery runs low, whereas the prudent user plans his/her trip and tries to reserve charging whenever possible.

The simulations use real road networks and CS locations in the Netherlands and Bavaria. The Dutch scenario uses a small number of fast (highway) chargers, representing the highway, long-distance use-case for reservations. The Bavarian scenario contains many slow (22 kW) chargers, representing the urban, short-distance use-case.

In each location, we test three kinds of reservations (or lack of, see Section 2.2.2): (1) no reservations provide a baseline for comparison, (2) ad-hoc reservations emulating reservations successfully deployed in the field test (see Section 2.2.3), and (3) planned reservations, which at the time of publication of [Bas+19] and [Bas+20] were not yet deployed in the field, but today are offered by e.g. EVgo.

#### Reservation Configurations

For the simulation case study, we selected two (three) feasible reservation types according to the planning capability dimension from Section 2.2.2 and compared them against the initial scenario with no reservations:

1. No reservations: Users can not reserve a charging station, but the current availability status of the connector/charging station is broadcast to drivers.

[3] https://github.com/agents4its/agentpolis

2. Guaranteed Ad-Hoc reservations: Only one reservation is allowed at a specific connector of a given charging station at any time, as described in Section 2.2.2. By the morphology in Table 2.1 from Section 2.2.2, this instance of reservations can be described as having “Yes” for enforceability, “No” for planning, “No” for fee, “Limited” for data availability (only the current status), “Yes” for roaming, and “first-come, first-served” for scheduling.

3. Uncertain Ad-Hoc reservations: the same as guaranteed ad-hoc reservations, except it has “No” for enforceability. This is the same system as the ad-hoc reservation system described in Section 2.2.3 and evaluated in the proof-of-concept in Bavaria (Section 2.2.3) for their feasibility.

4. Guaranteed planned reservation: Users can make planned reservations at any time, and CS can accept multiple reservations for different periods at the same time. In terms of features from Table 2.1, this system is defined as having “Yes” for enforceability, “Yes” for planning, “No” for fee, “Full” for data availability (full occupancy information), “Yes” for roaming, and “first-come, first-served” for scheduling.

However, since the uncertainty of reservations only ever reduces the performance of reservation systems, we have focused on the two guaranteed reservation types to streamline the analysis, evaluating the uncertain reservations in only one supporting experiment in the Bavarian slow-charging scenario only. This scenario uses the prudent driver (see below) and planned reservations.

See Table 2.1 for all reservation types used in the simulation experiments and Table 3.1 for the implemented driver-reservation combinations.

### 3.1.2 Driver and Vehicle Model

In the simulation, each driver performs a single trip that will require charging. Drivers use a planner that is a derivation of an A* algorithm that minimizes travel time [CŠJ18]. Driver origin and destination were sampled from the set of charging locations, whereas the departure time was sampled from Gaussian distribution with mean at 8:00 AM and variance of 1 h.

We used two types of drivers in the simulation case study: *naive* and *prudent*. While both try to optimize their trip time, each driver does it in a different way. See Table 3.1 for the overview of the implemented driver-reservation combinations.

**The naive agent** Naive driver uses his/her EV in a similar way that many people use their combustion-engine vehicles, i.e., not planning their trip with refueling in mind. Instead, a *naive* agent searches for charging stations opportunistically when his/her battery state-of-charge (SoC) drops below a certain threshold, which differs for each simulation location. Then, the driver selects the closest charging stations in the direction of his travel and, depending on the scenario, sorts them by availability and distance and then travels to the closest, presumably available CS.

In the no reservations scenario, the *naive* agent has information about the current occupancy or vacancy status of all CSs but only knows the waiting time of a charging station it has visited. If the charging station is blocked upon arrival or 20 minutes before, the naive agent will look for some other vacant CS in range and go to that one instead. If there is no vacant CS in range, the agent will only drive to another CS in its range if the waiting time at its current CS is longer than 5 hours. When multiple CSs are available to the agent, it uses a simple distance-based heuristic to select a charging station not far from its route.

In the ad-hoc reservation scenario, when the naive agent selects a vacant CS, he makes an ad-hoc reservation to this CS. In the absence of vacant CSs, the user goes to the best reachable CS, possibly making ad-hoc reservations if another CS becomes available. Ad-hoc reservations can only be made a maximum $\delta_t = 20$ minutes before the requested charging starts. The CS appears occupied to other EV users when the ad-hoc reservation is made.

In Table 3.1, note that we have not evaluated the combination of naive driver and planned reservation, as that is identical to a naive driver with ad-hoc reservation in our implementation.

**The prudent agent** A Prudent driver uses a driver assistance system (e.g., a mobile or in-car app) in order to plan to charge before starting his/her trip and, if the scenario permits, reserves charging stations at the beginning of the trip as well.

In the no-reservations scenario, similar to the naive driver, the *prudent* agent checks the availability of the target CS 20 minutes before its arrival and, if it is occupied, tries to re-plan using some other vacant CS. If he is unable to re-plan, he goes to the originally selected CS and, similar to the naive agent, only moves to other occupied CS if the currently selected CS has a queue longer than 5 hours.

The prudent agent uses both Ad-Hoc as well as Planned reservations. When using the planned reservation, the prudent user has a full view of CS occupancy into the future, not just its current availability. As such, the planning agent plans all charging sessions in advance and reserves the corresponding CS before starting the trip.

When using ad-hoc reservations, the planning agent only knows the current availability of all CSs. The planning agent tries to place ad-hoc reservations on CSs he intends to use. If this originally planned CS is unavailable for reservation, he tries to re-plan from its current position. In case there are no available CSs, he sticks to the original plan, waiting at the planned CS (for waiting times below 5 hours).

### Vehicle Model

Vehicles in the simulation use a very simple discharge model of an EV, with a fixed range in kilometers and battery state-of-charge decreasing linearly with the traveled distance.

| | No Reservations | Ad-Hoc Reservations | Planned Reservations |
|---|---|---|---|
| Naive driver | ✓ | ✓ | X |
| Prudent driver | ✓ | ✓ | ✓[a] |

[a] This is the only combination where we simulated uncertain reservations in addition to the guaranteed ones. This was done in the slow-charging scenario only.

**Table 3.1:** Overview of simulated (guaranteed) reservation scenarios.

### Metrics

In the results, we use two main metrics based on the stakeholder requirements of Section 2.2.2 to compare the different scenarios:

1. Trip duration in hours—this includes driving time (including time searching for available CS), waiting time at charging station and charging time. This metric addresses EV user goals G.01 and G.02 from Figure 2.2.

2. CS utilization—number of hours a CS was in use on average in the simulated period (24 h). This metric addresses charging station operators' goal G.03 from Figure 2.2.

## 3.2 Bavarian Slow-Charger Scenarios

The slow-charging scenario consists of a set of 1804 public charging stations in Bavaria, whose locations (Figure 3.1) were taken from Lemnet[4]. To simplify the scenario configuration, we set each charging station's charging rate to 22kW.

The demand is formed by randomly sampling origin-destination pairs from the set of 1804 charging stations. We sample up to 10,000 origin-destination for 10,000 agents, but in most experiments, we do not use more than 3,000 drivers because a higher number of drivers leads to large amounts of charging station congestion (Figure 3.2). We use charging stations as origins and destinations because the density of charging stations is highest in large urban centers. Demand generated this way then simulates traffic between these centers. Departure time is sampled randomly from a truncated normal distribution, with the mean being around 8:00 AM with a variance of one hour, meaning we simulate a single "rush hour" period. For this simulation, we use vehicles with a range of 130 km.

### Uncertain Planned Reservations Scenarios

In addition to the guaranteed reservation scenarios, we evaluate the uncertain reservations in this scenario as well (see Table 3.1). This uncertain scenario is based on the planned reservations scenario with prudent drivers with one addition: when an EV user arrives at a CS, he may find that the charging station is blocked, for example, by a combustion vehicle blocking access to the parking spot of the CS. In the simulation, this event is

[4] https://www.lemnet.org/de

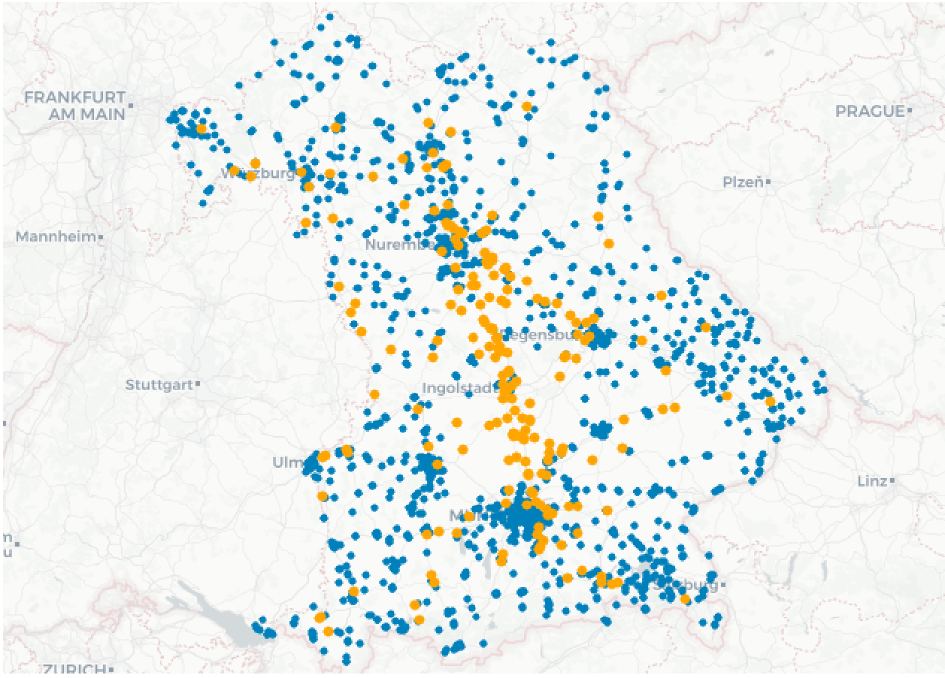


**Figure 3.1:** CS locations of reservation scenarios in Bavaria. Points in yellow are congested CSs in high-demand scenarios.

triggered randomly with probability $p_b$ every time an EV user arrives at a vacant CS. The blocking duration is sampled from a truncated random distribution with a mean of 2 hours and a variance of 1 hour.

Unlike the occupancy of chargers by other EVs, these blocking vehicles are not visible to the users before arriving at the charging station. The blocking vehicles also do not appear when there is a queue at the CS already. On the other hand, when one EV user discovers a blocking vehicle at some CS, this information is broadcast to all drivers. Once discovered, both naive and prudent user treat the blocking car as if it was another EV user.

### Results

In our first batch of experiments, we compare the different reservation types while under various amounts of demand. We first determine the maximum demand level (number of users in the simulation). We do this by running our simulation with an increasing number of naive users and guaranteed ad-hoc reservations. Starting from approximately 3000 users, the increasing congestion of the CSs causes an increasing ratio of users to be unable to reach their destination by the end of the day (see Figure 3.2). Assuming that in any realistic scenario, most users will reach their destination, we set the maximum demand in our following simulations to 3000 users. The minimum was chosen arbitrarily at 250 users.

In the next experiment, we compare the mean trip time of users and the utilization of charging infrastructure. First, we consider the case of guaranteed reservations (Figure 3.3).

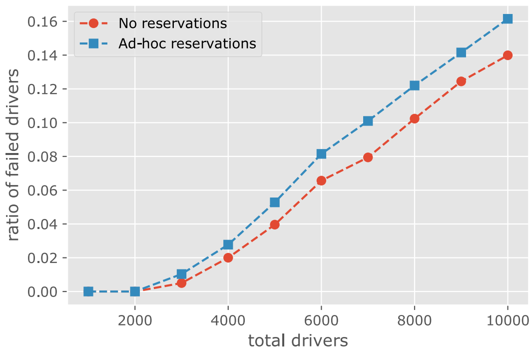


**Figure 3.2:** Number of naive users who failed to reach their destination with increasing demand.

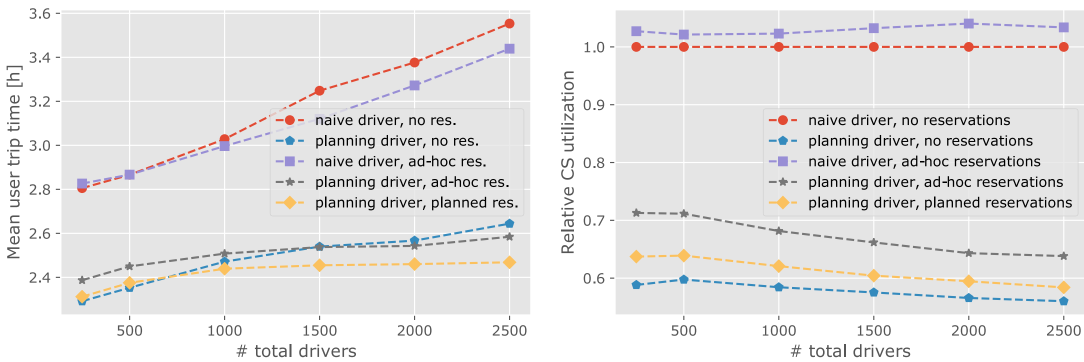


**Figure 3.3:** Mean user trip time and CS utilization of the guaranteed reservation types. The CS utilization is relative to the no-reservation and naive user baseline.

Here, we observe that the largest difference in trip time can be attributed to the user type, with a naive user taking much longer to complete his trip than a prudent user. However, the prudent user arrives at his destination with a much lower state of charge since he plans his trip to minimize time without consideration for the state of charge in the destination. On the other hand, the naive user charges whenever he runs under his pre-determined threshold state of charge, even when they could conceivably make it to the destination without charging. This leads to more detours, where he consumes charge, and also to a much higher state-of-charge at the destination and, therefore, much longer mean trip times.

However, for both types of users, reservations appear to reduce the average trip time when the demand increases. Surprisingly, for both agent types, in a scenario with low demand (250-500 users), no reservations perform better. This is caused by the fact that in these low-demand scenarios, the chance of finding a fast alternate route in case of an occupied CS is high, while reservations make CS unavailable even at times when no one is using them.

We calculate the mean CS utilization as the percentage of time the CS was occupied in the 24-hour period. Figure 3.3 shows the mean CS utilization across all CSs relative to

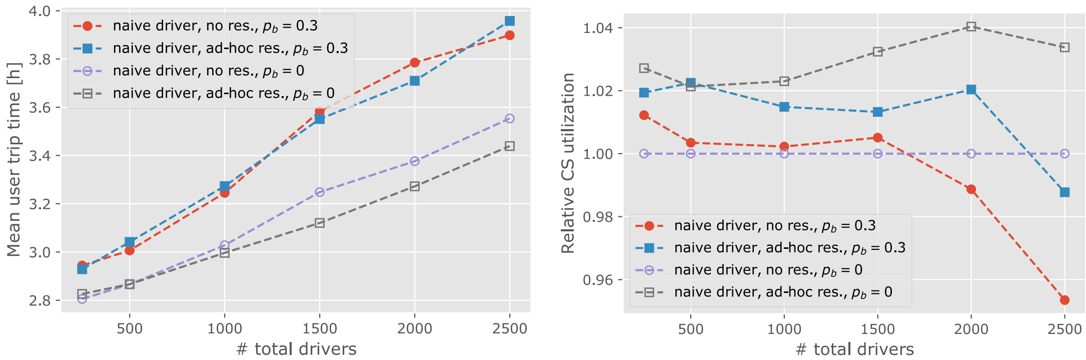


**Figure 3.4:** Mean user trip time and CS utilization of the uncertain reservation types. The CS utilization is relative to the no-reservation and naive user baseline.

the baseline scenario of naive user and no reservations. For both user types, reservations improve CS utilization. In the case of a naive user, the utilization is much higher than that of a prudent user. Again, this is because the naive user ends up charging a lot more, as described above. In our simulations, the mean CS utilization was always below 10% as a large number of CSs was far away from routes between main city centers (see Figure 3.1).

For the uncertain scenarios, we ran our simulations only with naive users. We compare the naive user with and without ad-hoc reservations in the guaranteed scenario to the uncertain scenario with $p_b$ (the probability of CS being blocked by some external car) being 30%. As Figure 3.4 shows, the reservations reduce the mean travel time in the guaranteed scenario but not so in the uncertain case. On the other hand, ad-hoc reservations consistently lead to better utilization of CS in both scenarios.

Additionally, we experimented with varying $\delta_t$ of ad-hoc reservations. We ran this experiment in the guaranteed scenario with the naive user and a demand of 1000 users. We varied $\delta_t$ from 0 to 60 minutes, and while the CS utilization increased and trip time reduced when increasing $\delta_t$ until approximately the 25-minute mark, increasing $\delta_t$ further brought no further benefit. However, the scale of maximum change of both average trip time and CS utilization was very small, about 1% of the absolute value.

Similarly, we experimented with varying the probability $p_b$ of a CS being blocked by some external car. We ran this experiment with the ad-hoc reservations using the naive user, and a demand of 3000 users. We varied the $p_b$ from 0% to 40% in 10% increments. Extrapolating the values, ad-hoc reservations seem only to make sense in scenarios with $p_b$ up to about 5%. For higher probabilities, reservations do not improve the average trip time. The deterioration of average trip time with $p_b$ going from 0% to 40% was approximately half an hour in an approximately 3.5-hour trip, the trend being approximately linear in $p_b$.

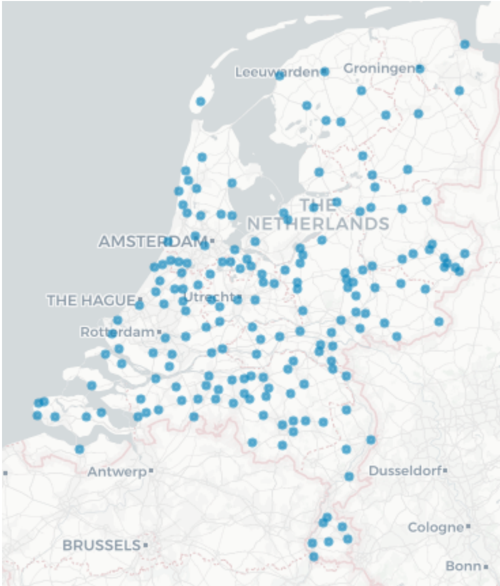


**Figure 3.5:** Charging location clusters for simulation case study.

## 3.3 Dutch Fast-Charger Scenarios

For the fast-charger case study, we selected the roadmap of the Netherlands, using open street map data[5] with 300,000 nodes and 800,000 edges (residential roads were removed). We selected the Netherlands as it has a dense and well-connected highway network, and its electric mobility infrastructure is well-developed. For charging station locations, we used open charge map data[6]. We clustered these locations to create a set of 150 fast chargers (50 kw), each with one charging slot in the most important EV charging centers (see Figure 3.5). In our experiments, we sampled between 10 and 2000 drivers using the infrastructure. Each simulation spans 24 hours. For the vehicles in this scenario, we assume a 200 km range with a 100 kWh battery and each vehicle starting with at least 30% state-of-charge. In this scenario, the naive driver starts to search for the next charging station when his battery level drops below 50% state-of-charge.

### Results

In our results, the mean trip duration of naive drivers is about twice that of prudent drivers. Because naive drivers look for charging when their state of charge drops below a threshold, they tend to arrive at their destination with a much higher state of charge than prudent drivers (who optimize time and, as a result, plan their trips with tighter margins). This means that naive drivers and prudent drivers are not comparable in terms of absolute trip duration. To make the two driver types comparable, we use the

[5] https://www.openstreetmap.org

[6] https://openchargemap.org

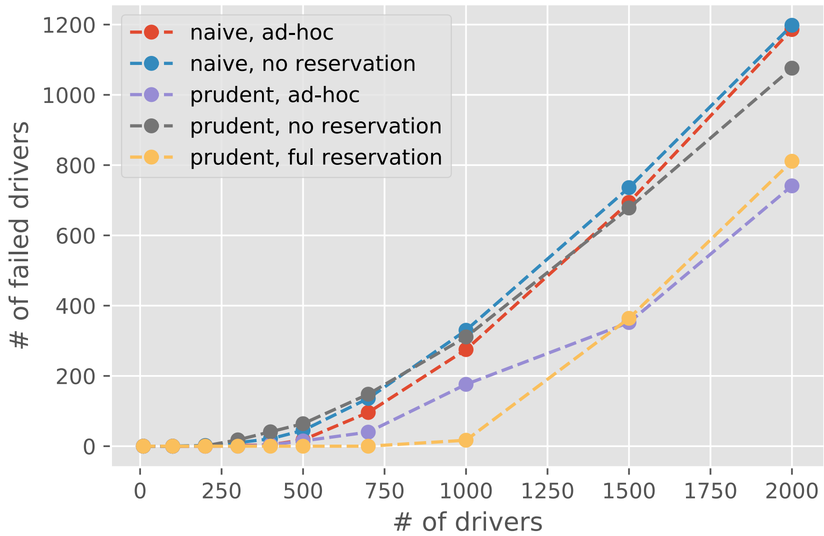


**Figure 3.6:** Number of failed drivers as demand increases for different reservation types.

*adjusted* versions of these variables obtained by subtracting the required time to charge this leftover energy from the total trip time and CS utilization.

We ran all simulation scenarios with different numbers of drivers (see Figure 3.6 for values). We found that the driver type and reservation method have a strong influence on the number of drivers that reach their destination successfully. More precisely, a driver fails if he/she runs out of charge en route or if he/she takes more than one day to finish the trip. By observing the results in Figure 3.6, we have decided to use the scenario with 200 drivers for comparison of four reservation types, as (1) the vast majority of all drivers finish their trips, and (2) the demand is high enough for reservations to make sense. In our experiments, it is apparent that the greater the demand, the bigger the benefit of reservations.

Note that in our experiments, we consider only the fixed value of battery capacity and charging speed. The relative results of different methods in Figures 3.7a and 3.7b would not change if we considered different battery capacities or charging speeds. This is because we select the scenario for comparisons as the case with saturated supply as per Figure 3.6.

Additionally, note that the error bars in Figure 3.7a,b are not confidence intervals as the results are from a single deterministic simulation run. Instead, they show the mean and standard deviation in the population of EV drivers, respectively, from the CSs used in the simulation run. Sensitivity to different driver origin/destinations and departure times is addressed in a separate sensitivity analysis.

In addition to this main experiment, we ran two additional types of experiments:

1. *Sensitivity analysis experiment* where we evaluated all scenarios with 10 different samples of driver origin-destination pairs to test the sensitivity of our results to geographical changes in demand. The sensitivity analysis suggests that the relative results in Figure 3.7a,b are preserved when the demand is sampled differently.

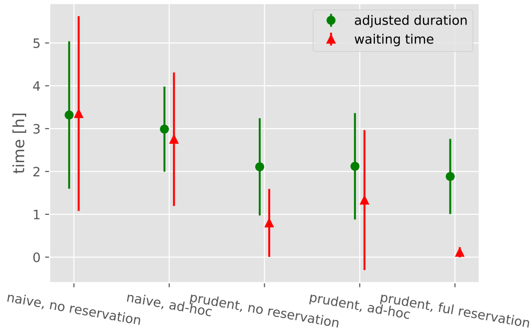


**(a)** Mean and std of *waiting times* (waiting at charging station for a connector to become available) and *adjusted trip times* (driving time and charging time) in the population of drivers in different scenarios. These results correspond to the EV user goals G.01 and G.02 in Section 2.2.2.

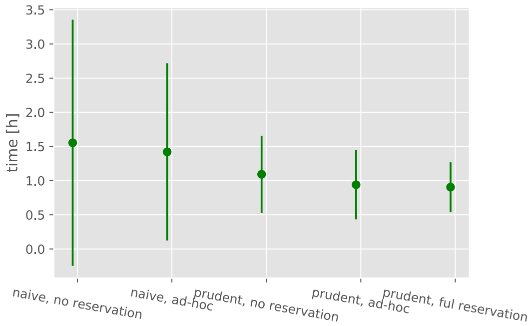


**(b)** Mean and std of *adjusted CS utilization* in the set of all charging station, in different scenarios (scenarios lasting 24 h). These results correspond to the charging station operator's goal G.03 Section 2.2.2.

**Figure 3.7:** Simulation results illustrating the mean waiting time of EV users (a) and mean utilization of the charging station (b).

2. *Baseline experiment* where we set all CS capacity to infinity to create a baseline for our main results. Compared with the main experiment, in the baseline, the adjusted trip duration of the naive driver is 12 min (0.2 h) longer than that of the prudent driver.

Figure 3.7a shows the mean of the adjusted trip duration in the different scenarios. The most important difference is between the driver types. According to our results, planning trip routes has a much bigger impact than choosing reservation types.

Surprisingly, ad-hoc reservations improve the trip time for naive drivers but not for the prudent driver (however, the prudent driver still performs better than the naive driver in all scenarios). In fact, ad-hoc reservations worsen the trip times of prudent drivers. This is because many prudent agents end up re-planning to the same charging stations when available slots open up and slow each other down. Notably, the prudent agents with planned reservations offer the lowest waiting time.

For utilization of charging stations, reservations and driver type seem to have no significant effect on the total (adjusted) utilization of charging infrastructure (Figure 3.7b) at the same demand level as previous figures. Without the adjustment, in the scenario with naive drivers, charging stations are utilized significantly more than in the scenarios with prudent agents as they arrive at their destination with a significantly higher state of charge.

## 3.4 Summary

Electric mobility provides excellent alternative transportation technology for our modern society with less environmental impact than traditional combustion-engine vehicles. In spite of their advantages, electric vehicles are not yet widely accepted and deployed, partly due to challenges such as (1) insufficient charging infrastructure and (2) proportionally long charging times. Those challenges lead to dissatisfaction of the EV users as well as inefficient utilization of the charging infrastructure.

Based on the analysis of stakeholder requirements, we have identified reservations about EV charging stations as a promising approach to address the above-mentioned two challenges of electric mobility. However, such a system has not yet been addressed wholistically in the literature. By analyzing the state of the existing charging infrastructure, existing communication protocols, and systems, we have developed an interoperable reference architecture for the reservation of charging stations.

This architecture can be instantiated into different reservation systems based on the needs and capabilities of the charging station operators.

In the process, we have identified that "full reservations" in the expected sense are not feasible with most existing CS equipment without significant upgrades. As an intermediate step towards full reservations, we have described ad-hoc reservations, a limited reservation variant, as supported by most of the existing infrastructure. We have demonstrated this capability in a field trial in Bavaria, where we implemented an ad-hoc reservation system for the E-Wald charging network.

However, the limits of the ad-hoc reservations don't make it clear that their benefits would outweigh the costs of their deployment to the whole EV ecosystem. To identify their benefit in relation to the baseline of no reservations and full reservations, we have developed a multi-agent simulation model of EV users and charging stations. We have used this model to compare the performance of different reservation types in two different locations with different charging infrastructures and user behavior.

We performed the first set of simulations on a road graph of Bavaria, Germany, involving slow chargers and drivers with relatively shorter ranges. The conclusions from this experiment were the following:

1. *Many CS are underutilized.* In all scenarios, the CS utilization was below 10%. While the simulation did not attempt to replicate existing travel patterns, the existing charging infrastructure was rarely built based on demand. In the simulated rush hour that modeled travel demand between urban centers, most CSs were rarely used. The conclusion here is that if EV reservations are to be deployed in the future, they should be deployed selectively to charging stations with high demand.

2. *Reservations don't make sense in low-demand situations.* This conclusion is unsurprising; again, we conclude that reservation systems should be deployed selectively to high-demand charging stations. In this sense, roadside and urban slow-charging is therefore not a good model for reservation systems since the network of slow-chargers is relatively more developed than the fast-charging network, and the demand for any individual CS is, therefore, lower on average.

3. *User behavior has a much larger impact on performance than the reservation system used.* The simulation showed that the travel time improvement from reservations

is negligent compared to the difference in driver behavior. Based on this finding, we conclude that driver assistant technologies that allow drivers to plan easily are an important piece in the puzzle of improving EV user experience. Conceivably, such systems could be integrated with reservation systems to provide a seamless experience for the user.

4. *Reservations improve trip times and reduce utilization in high-demand situations.* With increasing demand, we find that reservations do improve the trip times for drivers, and full reservations have a bigger impact than ad-hoc reservations. However, while full reservations benefit drivers at all demand levels, ad-hoc reservations increase travel times at lower demand levels. This is because ad-hoc reservations make CS unavailable even when no one uses them, while full reservations only make them unavailable when someone uses them.

5. *Uncertainty makes reservations useless* In our simulation of the uncertain ad-hoc reservations (referring to the situation where the driver, upon his arrival, finds the reserved charging location unexpectedly blocked by another vehicle), we find that reservations only make sense when the chance of the CS being blocked is low, and reservations make the whole system worse-off when reservations can't be guaranteed. Thus, we conclude that reservations should only be deployed when the reservations are enforceable, either by technical or legal means.

Overall, the conclusion from the Bavarian experiments is that reservations are not a silver bullet for improving EV user experience. We conclude that reservations should be deployed selectively to high-demand charging stations and that reservations should only be deployed when they can be guaranteed. Because the guarantees through technical means are likely a costly solution for roadside, slow-charging stations, we conclude that reservations make sense only for high-demand, fast-charging stations where the cost of the charging station and revenues are higher. This undermines the use-case for ad-hoc reservations as a widely deployable, stepping-stone solution.

However, the Bavarian set of simulations does not convincingly answer the question of what benefit the two reservation types bring to the EV ecosystem. Based on the limits of the above-described findings, as described in Section 3.3, we have developed a second simulation model for fast-charging stations in the Netherlands that addresses the shortcomings of the Bavarian model. Importantly, we "adjust" the large difference between the two different driver types in resulting travel times and CS utilization by accounting for charge "left-over" in the batteries of naive drivers at the end of their trips (see Section 3.2 for more comprehensive description). This adjustment is made on the assumption that all drivers can charge their vehicles to 100% state-of-charge at the end of their trip using a slow charger.

The conclusions from this fast-charger experiment were the following:

1. *Reservations don't greatly impact adjusted mean CS utilization.* As such, reservations are more useful in addressing the user experience (Goals G.01 and G.02 in Figure 2.2)

than improving CS utilization. The results here confirm the results from Bavaria when controlling for the overcharge by the naive drivers.

2. *By adjusting for end-of-trip charge, driver behavior is still more important than reservations.* This result reinforces the use case for driver assistant technologies, especially those integrating CS reservations seamlessly. Efficient assistant technologies can be an important tool for reducing fear of "range anxiety" and increasing EV adoption[RFK15].

3. *Ad-hoc reservations improve travel and waiting time for naive drivers but increase both travel and waiting times for prudent drivers.* The result from the Netherlands supplements the Bavarian results, where ad-hoc reservations improved travel times for both driver types at higher demand levels but not at lower demand levels. For the prudent drivers, the issue with ad-hoc reservations is that they actually don't help much with planning the trip and can make the situation worse in our population of prudent drivers due to their interactions with each other and reservations (see the discussion of results in Section 3.3). As a result, prudent drivers will end up waiting for a long time at charging stations off of their originally planned route. As such, we conclude that ad-hoc reservations are not practical as an intermediate solution toward full reservations.

4. *Full reservations in combination with prudent drivers provide significant improvements in travel times and waiting times compared to the baselines* Finally, our results show that when combined with full reservations, prudent drivers can achieve significant improvements in travel times and waiting times compared to the baselines. Full reservations enable the population of prudent drivers to distribute their charging sessions in time and space along their route to avoid waiting times completely. Therefore, our conclusion is that full reservations are the appropriate reservation type to pursue.

Overall, the final conclusion of our investigation of EV CS reservations is that full reservation systems should be considered for high-cost, high-demand charging stations. The intermediate solution of ad-hoc reservations, while attractive due to its low cost, is not a good stepping stone towards improving the user experience. However, the development of reservation systems should go hand-in-hand with the development of driver-assistant technologies that allow drivers to plan their trips, including charging, seamlessly. This is because driver behavior is a much larger factor in determining travel time and waiting time than the choice of reservation systems.

The full reservations, in combination with drivers who plan their trips, including charging in advance, can provide significant time savings to the drivers. For people who are used to internal combustion vehicles, this may require a shift in behavior since refueling can be done much faster with minimal constraints on refueling multiple vehicles in parallel. While planning their routes with charging reservations greatly increases the predictability of EV travel time and thus drivers' comfort, it might become even more important in the future when the number of simultaneously charging EVs will be putting significant strain on the electricity grid. In this case, the reservation system can be used to coordinate the charging of EVs in a way that helps maintain the stability of the grid.

The following Chapters 4 and 5 will construct a model of dynamic EV charging around the concept of full-reservations verified in this chapter. This pricing then can be used to propagate incentives in the EV ecosystem to encourage an equilibrium between the needs of all the stakeholders.

CHAPTER 4

# Dynamic Pricing of EV Charging with MDPs

In previous Chapter 3, we have concluded that full reservations are the most viable reservation type to use in EV charging. However, we have decided to ignore the pricing aspect of reservations, a feature that can play an important role in the adoption of reservations and can be used as a demand-side management tool to influence the behavior of EV drivers and maximize the revenue of charging station operators.

This chapter will investigate our simplified, but optimal revenue maximizing dynamic pricing model for EV charging based on two conference publications [MKJ18a; MKJ18b]. Both are first-author publications but rely on the work done during the ELECTRIFIC project.

First, we describe the problem of pricing EV charging as a deterministic multi-agent problem. By aggregating the agents' behavior and their hidden variables, we reformulate the pricing problem as a stochastic problem, which we simplify and solve with Markov Decision Processes. Finally, we evaluate the performance of the simplified optimal dynamic pricing strategy against several baselines on a dataset of EV charging.

## 4.1 Multi-agent Model of Pricing of Charging Services

In this section, we formalize the pricing of charging services as a multi-agent model. Our model describes a multi-agent system with a single charging station operated by one charging station operator and several EV drivers wanting to charge their vehicles. Within such a system, we focus on the problem of dynamic pricing strategy [AE08]. The single charging station operator uses the strategy to post prices that are then accepted or rejected by EV drivers.

We focus on the demand-response pricing strategy [AE08] for a single charging station. As input, we use a discretization of possible charging parameters (time, duration), current

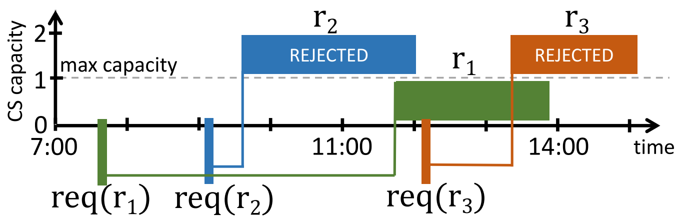


**Figure 4.1:** The requests for charging in the future arrive at the charging station in a sequence. Accepted charging requests and their duration affect whether later requests can be accommodated. If the charging station accepts requests as they arrive, $r_1$ arriving at $req(r_1)$ will block $r_2$ and $r_3$. If the charging station rejected $r_1$ instead, $r_2$ and $r_3$ could be accepted. Whether to accept or reject a charging request can be decided by comparing the value of the charging request with the expected value of the not yet allocated available capacity.

and historical data on the utilization of the charging station, and the expected price elasticity of the demand [Gal15] for charging services.

The goal of the pricing strategy described below is the maximization of charging station revenue within a particular time horizon. However, other optimization criteria are possible and investigated in the following sections to achieve different goals. For example, a publicly owned charging station that is not concerned with profits may attempt to maximize charging station utilization or minimize waiting times at the charging station.

### 4.1.1 Formalization

In our multi-agent model of pricing of charging in electromobility, we consider a finite set of EV driver *agents* $E = \{\phi_1, \phi_2, \ldots\}$ and one charging station operator as one additional agent $\Phi$ . The model is formally defined as a tuple $M = \langle c, d, E, \Phi \rangle$. $d = (r_1, r_2, \ldots)$ is demand expressed as a finite sequence of *charging services requests* $r_i$ sent by the EV driver agents to the charging station operator agent in a sequence.

Each request is defined by its $start(r_i)$ and $end(r_i)$ times of charging, time $req(r_i)$ when the request was issued and $char(r_i)$, the requested charge in kWh. We assume that $req(r_i) \leq start(r_i) < end(r_i)$ and $req(r_i) < req(r_{i+1})$. Each request is associated with exactly one driver agent $\phi$; i.e., denoting a set of requests associated with $\phi$ as $\phi(d)$, it holds $\bigcap_{\phi \in E} \phi(d) = \emptyset$ and $\bigcup_{\phi \in E} \phi(d) = d$.

For each time, we define *free capacity* $c : \mathbb{R} \to \mathbb{R}_0^+$ that determines the maximal charging capacity of the charging station in time. The free capacity $c$ represents the aggregate of all charging station constraints, such as power grid capacity or the number of available charging connectors.

The interaction of the EV driver agent and the charging station operator agent is modeled via requests $r_i$ from EV driver agent $\phi_j$ and prices $h_i$ returned by the charging station operator agent. Formally, $h_i = \Phi(r_i)$, where $\Phi$ is the *charging service pricing operator* of the charging station operator agent, $r_i$ is the charging service request and $h_i$ is the resulting proposed price. If the charging station operator agent can not accept a request, it sets $h_i = \infty$.

$\phi_1, \phi_2, \ldots$ denote the *decision processes* of the EV driver agents that determine whether the driver agent accepts the proposed price $h_i$. We write $\phi_j(h_i) = \top$ iff the $j$th EV driver agent is willing to pay price $h_i$ for the requested charging service $r_i$ and $\phi_j(h_i) = \bot$ otherwise ($\phi_j(\infty) = \bot$ always). Provided that $\phi_j(\Phi(r_i)) = \top$, the EV driver agent accepts the proposed price for charging request $r_i$, and we assume both the charging and payment eventually happen in the system. An execution of the model $M$ is a sequence of prices and decisions $(\langle h_1, \phi_j(h_1)\rangle, \langle h_2, \phi_{j'}(h_2)\rangle, \ldots, \langle h_i, \phi_{j'}(h_i)\rangle, \ldots)$ for all agents $\phi_j$ and their charging requests $r_i \in d$, such that $h_i = \Phi(r_i)$.

The goal of the charging station operator is to maximize its revenue by setting prices with $\Phi$. Given that the price $h_i$ of reservation $r_i$, the *revenue* $\rho$ at the end of the time horizon can be written as the sum of prices across all realized reservations (Equation (4.1)):

$$\rho(c, d, E, \Phi) = \sum_{\substack{r_i \in \phi_j(d) \\ \phi_j \in E}} \Phi(r_i) \mathbb{1}_{\phi_j(\Phi(r_i)) = \top} \tag{4.1}$$

We are looking for revenue maximizing pricing operator $\Phi^*$:

$$\Phi^* = \arg\max_{\Phi} \rho(c, d, E, \Phi) \tag{4.2}$$

The free capacity constraints the maximization:

$$\forall \tau \in \mathbb{R},\ c(\tau) \geq \sum_{\substack{r_i \in \phi_j(d) \\ \phi_j \in E}} \mathbb{1}_{\phi_j(\Phi(r_i)) = \top} \mathbb{1}_{start(r_i) \leq \tau < end(r_i)} \tag{4.3}$$

Here, $\mathbb{1}_{\phi_j(\Phi(r_i)) = \top}$ is the indicator function that equals 1 for requests accepted by the EV driver agent and $\mathbb{1}_{start(r_i) \leq \tau < end(r_i)}$ is the indicator function of the charging interval of request $r_i$.

Finding the maximizing pricing operator in Equation (4.2) is not a simple task. The pricing operator is part of the sum in Equation (4.1) involving many discontinuities and depends on the set of reservations $d$ through the customer decision processes $\phi_j$. We would have to start by modeling the customer processes to find the maximizing operator. It is impractical to model each customer independently, and it is almost impossible to do so accurately.

To make the problem tractable, we aggregate the behavior of a population of customer agents $E$ into the probability distributions that describe the behavior of customers together as $\mathcal{E}$, and the sequence of requests $d$ into a probabilistic demand model $D$. As such, we can no longer maximize the profit in absolute numbers. Instead, we maximize the expected profit (in the statistical sense):

$$\Phi^* = \arg\max_{\Phi} \mathbb{E}[\rho(c, D, \mathcal{E}, \Phi)], \text{ s.t. Equation (4.3)}$$

The framework of MDPs is the tool that deals with problems posed this way.

## 4.2 Optimal Dynamic Pricing as Markov Decision Process

The maximization problem given by Equation (4.2) is sequential in nature. The pricing function has to respond to the confirmed reservations as it can not exceed the capacity of the charging station (given by Equation (4.3)). Additionally, the charging station operator agent will generally not have exact knowledge of the individual $\phi_j$ or of the number of requests to be made in one day. However, the charging station can have a probabilistic model of the EV driver behavior and of expected demand. Thus, from the perspective of the charging station, revenue maximization is a MDP [Bel57].

We aggregate the charging station operator's understanding of the EV drivers' decision processes in $E = \{\phi_1, \phi_2, \ldots\}$ into the price-demand function $\mathcal{E}$. Given a request $r$ and a generated price $h$, price-demand function $\mathcal{E}(h)$ is the probability of an EV driver agent accepting the price $h$. Note that the description of customer buying decisions using the price-demand function is dual to the stochastic customer valuations of products, as described in Section 2.4.3. Arrivals of requests and their start times and end times are modeled using the demand model $D$.

Thus, we reformulate the maximization goal from Equation (4.2) as the maximization of the expected revenue across possible decision policies $\pi$:

$$\pi^* = \arg\max_{\pi} \mathbb{E}[\rho(c, D, \mathcal{E}, \pi)] \tag{4.4}$$

The expectation is with respect to the joined probability distribution of demand and probability of acceptance, and the maximization is subject to Equation (4.3). Note that unlike the pricing operator $\Phi$ in Equation (4.2), the decision policy $\pi$ is not a function of just the request $r$ but of the state of the system. Additionally, the decision policy does not work directly with the individual requests in $d$ and customer decision process $E$, but with the aggregate customer price-demand function $\mathcal{E}$ and demand $D$.

Based on the price-demand function, we define the MDP pricing strategy that determines prices offered by the charging station operator throughout the day. Generally, MDPs describe the development of a system through a set of discrete states. Therefore, we discretize time in our pricing model into $n$ disjoint charging timeslots, price into $\hat{h}$ price levels, and free capacity into $\hat{c}$ capacity levels.

Since we assume in our model that all charging sessions use the same electrical power, the maximal *free capacity* in the time interval is the upper limit on the number of concurrent charging sessions. Free capacity incorporates all charging station operator constraints, including the power grid capacity or the number of available charging connectors, into one number.

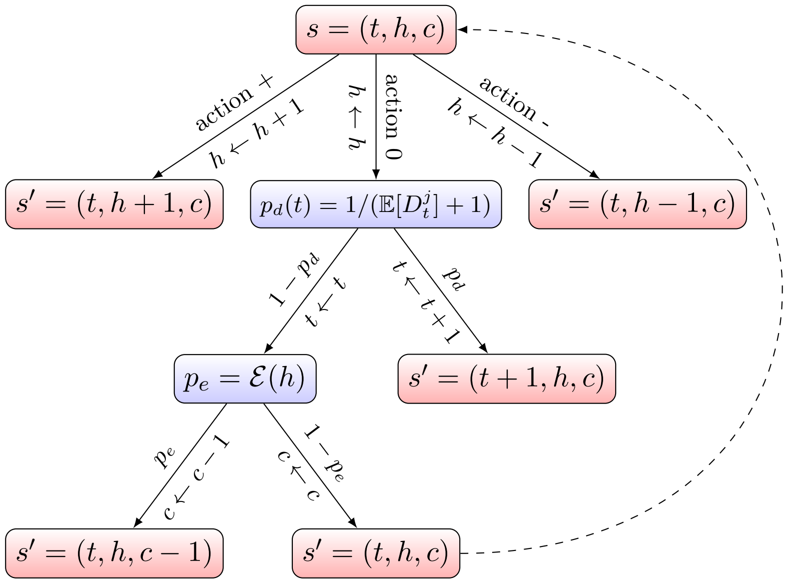


**Figure 4.2:** The structure of the transition function $\mathcal{T}$ for $j$th MDP in the split-MDP dynamic pricing strategy. From state $s$, the system can develop into one of the five possible states $s'$, one of them being the original state $s$. Given state $s$, the probability of getting to the next state $s'$ is given by multiplying the probabilities along edges. States are the decision nodes (in red), chance states are in blue, and contain the definition of the probability used on the edge.

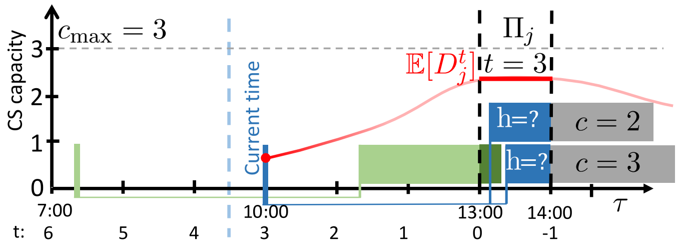


**Figure 4.3:** The figure shows the state of MDP $\Pi_j$ which is associated with the $j$th charging timeslot (13:00 to 14:00). The charging station shown has maximal capacity $c_{\max} = 3$. The green box represents an alredy reserved charging session, this reservation was made at 7:20. The expected demand for charging in $j$th timeslot between 10:00 and 11:00 is $\mathbb{E}[D_t^j]$. $t = 3$ signifies that the demand is from timeslot 3 intervals ahead of the timeslot $j$. The red curve illustrates the expected demand in the other charging timeslots. State of $\Pi_j$ is given by three values: $t = 3$, the *number of time intervals to charging* from current time, *current price* $h$ and *free capacity* $c$ that depends on the real-time overlap of the accepted charging requests (leftmost rectangle) with two examples of possible charging requests requested in the “current time” in interval 10:00 to 11:00.

### 4.2.1 Split-MDP Dynamic Pricing in Discrete Time Intervals

In our discretized model, for each timeslot, there are $\hat{h}$ price and $\hat{c}$ capacity levels that will inform the offered price for the next charging session in this timeslot. However, since we want to price reservations for future charging, the price of charging can also depend on how much ahead is the EV driver making the reservation, which can be up to $n$ timeslots before the charging (using the same time discretization of time as for the charging sessions). See Figure 4.3 for visualization of these parameters. As such, the state space in each timeslot has size $n\hat{h}\hat{c}$, and the branching factor is up to 5 (see Figure 4.2).

Because price or capacity change in any timeslot can potentially impact price in any other timeslot (as visualized in Figure 2.11), finding solution to Equation (4.4) means finding solution for $n$ timeslots together; that is in a state space with $n(\hat{h}\hat{c})^n$ states and branching factor at least $5^n$.

To avoid this combinatorial explosion, we find the optimal pricing policy for each timeslot independently through $n$ independent MDPs that maximize expected revenue in each time interval Equation (4.5). However, using this *split-MDP* model is a departure from Equation (4.4). Instead, the split-MDP provides $n$ independent solutions to

$$\pi_j^* = \arg\max_{\pi_j} \mathbb{E}[\rho(c, D, c, \mathcal{E}, \pi_j)],\ j \in \{1, \ldots, n\} \tag{4.5}$$

Combined together, summed maximums from Equation (4.5) are bound above by maximum from Equation (4.4). However, the tightness of this bound is difficult to determine without solving the joint MDP and also depends on the data used. We overcome this limitation in the next Chapter 5.

The pricing strategy, therefore, uses one MDP for each time interval. The solution to MDP for the $j$th interval is the pricing policy $\pi_j$.

The MDP is defined as a tuple $\Pi_j = \langle S, A, R, \mathcal{T}, s_0, S_T \rangle,\ j \in \{1, \ldots, n\}$. In each $\Pi_j$, $S$ is a finite set of *states*, $A$ is a finite set of actions; $\mathcal{T} : S \times A \times S \to [0, 1]$ is the transition function forming the transition model and giving a probability $P(s'|s, a)$ of getting to state $s'$ from state $s$ after action $a$; and a reward function $R : S \times A \times S \to \mathbb{R}$.

Starting in the initial state $s_0$, any action from $A$ can be chosen. Based on this action, the system develops as prescribed by $\mathcal{T}$ to the next state where another action can be applied. The reward can be received during the move based on the $R(s, a, s')$ function. For a detailed description of MDP models, see Section 2.5.

A state $s$ in our MDP is defined by triplet $(t, h, c)$ and is illustrated in Figure 4.3. Here $t \in \{0, \ldots, t_{\max}\} \cup \{-1\}$ denotes *the number of time intervals* to the execution of charging (0 denotes the hour of charging, $t = -1$ signifies one of the terminal states $S_T$ of the MDP). $c \in \{0, 1, \ldots, c_{\max}\}$ is the *current available capacity* in the time interval, i.e. how many more requests can be accommodated in the time interval, and $h \in \{0, 1, \ldots, h_{\max}\}$ is the current price level set in the time interval. The set of actions $A$ contains three actions, price $+1$, price $-1$, and no change to the price.

Because each accepted request reduces capacity in the timeslots by one, the reward function $R$ generates reward $h$ for any transition between states $s = (t, h, c)$ and $s' = (t, h, c-1)$ for all $t \in \{0, \ldots, t_{\max}\}$, and $h$ and $c$ in their domains.

The transition function $\mathcal{T}$ is based on the price-demand function $\mathcal{E}(h)$ and the expected number of requests $\mathbb{E}[D_t^j]$ for the charging in the $j$th interval that will be issued $t$ intervals before. This is illustrated with the red bar in the $j$th interval in Figure 4.3. Components of the transition function are given in Equations (4.6) and (4.7). Combining these components into a transition function is shown in Figure 4.2.

$$p_d(t) = \frac{1}{\mathbb{E}[D_t^j] + 1} \tag{4.6}$$

$$p_e(h) = \mathcal{E}(h) \tag{4.7}$$

The probability $p_d(t)$ of no more charging requests arriving into state $(t, h, c)$ is calculated from the expected absolute demand for charging in $k$th time interval $t$ intervals before the start of charging $\mathbb{E}[D_t^k]$. Equation (4.6) is obtained by modeling $D_t^k$ as having geometric distribution, arrival of charging request as Bernoulli trial (failure in Bernoulli trial meaning no more charging request $t$ intervals before charging) with failure probability $1 - p_d(t)$.

The geometric distribution (see Section 2.4.1) is the probability distribution of the number $Y$ of failures before first success in consecutive Bernoulli trials with success probability $q$ (this means 0 is in the support). $Y$ has expected value $\frac{1-q}{q}$. Identifying $Y$ with $D_t^j$, we get $q = p_d(t)$. For each $t$, the MDP remains in some state $s \in \{(t, h, c) | c \in \{0, 1, \ldots, c_{\max}\}, h \in \{0, 1, \ldots, h_{\max}\}\}$ until successful Bernoulli trial with $p_d(t)$ being the probability of success.

Arriving charging request has the probability of acceptance $\mathcal{E}(h)$. If the customer rejects the price offered by the pricing strategy, the MDP remains in the same state. This is illustrated in Figure 4.2 in the branch ending in $s' = (t, h, c)$.

There are few exceptions to the probabilities defined by Figure 4.2. One of them consists of the bounds of the domains of state variables. If the price is maximal resp. minimal, the action to increase resp. lower the price is not available. Similarly, when capacity is 0, $p_e(h) = 0$ as no additional request can be accommodated by the charging station operator in the given time interval. Finally, $t = -1$ denotes the set of terminal states $S_T = \{(t, h, c) \in S | t = -1\}$ where the MDP terminates.

The pricing strategy uses $n$ MDPs at once. For a charging request $r$, the free capacity $c_r^j$ in each MDP is calculated from the overlap (in continuous, not discretized, time) of accepted requests (green rectangle in Figure 4.3) and incoming request $r$ (blue rectangles in Figure 4.3). In every *timeslot* $I_j$ corresponding to $\Pi_j$, starting from the state $(t^j, h^j, c_r^j)$, we apply pricing policy $\pi_j$ repeatedly until the policy does not suggest price change. The final price of the request $r$ offered to the EV driver agent is the sum of the prices in timeslots the request overlaps. In the case of partial overlap with some timeslot, the price is proportional to the size of the overlap:

| CS Dataset Statistics | Mean | Std |
|---|---|---|
| Charging session duration | 0.726 h | 0.794 h |
| Charge per charging session | 6.72 kWh | 5.19 kWh |

**Table 4.1:** Summary statistics of the real-world charging data for the selected charging station with three charging points. The dataset contains charging sessions recorded over several weeks.

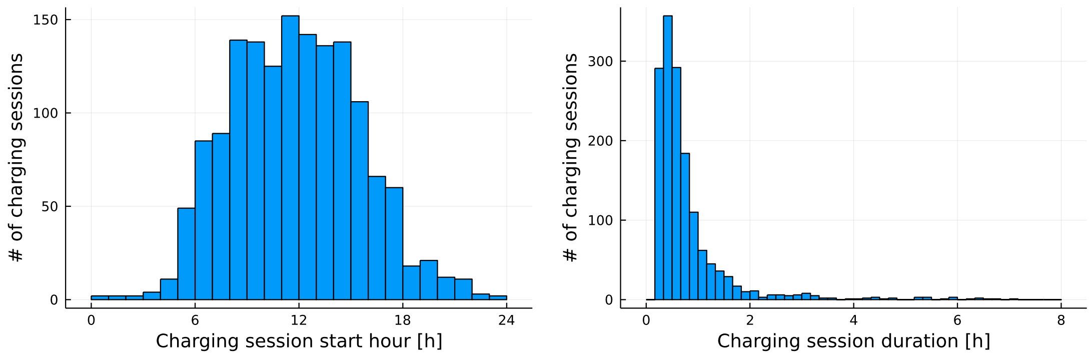


**(a)** **(b)**

**Figure 4.4:** Histograms show the start (a) and duration (b) of charging sessions contained in the considered dataset.

$$\Phi(r) = \sum_{j=1}^{n} \frac{|I_j \cap (start(r), end(r))|}{|I_j|} \Pi_j(r) \tag{4.8}$$

Note that $j$th MDP internally assumes that the EV driver is making a decision based on the price $h^j$ (the state variable in $\Pi_j$). Depending on the length of the charging request, this price is only part of the price $\Phi(r)$ offered to the driver by the pricing strategy. Thus, the length of the timeslot should be chosen, similar to the length of the average charging request.

## 4.3 Evaluation of the Optimal Split-MDP Dynamic Pricing Method

We base our evaluation of the split-MDP dynamic pricing algorithm on a real-world dataset. In this section, we first summarize the used dataset's statistics and describe the preprocessing we performed on the data. Then, we describe the experiments we conducted with the data and the results we obtained.

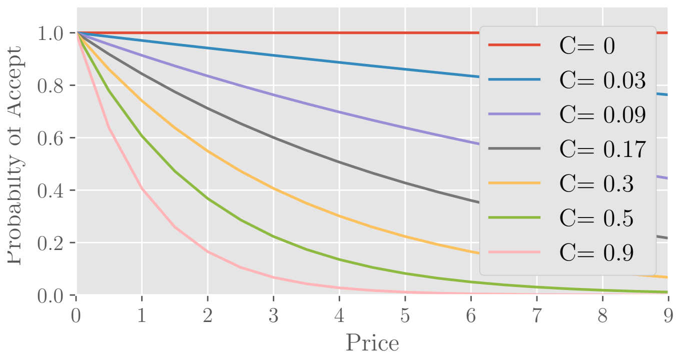


**Figure 4.5:** Price-demand functions $\mathcal{E}(h)$ for different values of the price elasticity parameter $C$.

### 4.3.1 Real-world EV Charging Data

The dataset contains information on charging sessions realized at charging stations of Mer Germany GmbH (formerly E-WALD, a partner in the ELECTRIFIC project)[1], one of the biggest EV charging station operators in Germany. The E-WALD dataset includes timestamps of the beginning and the end of each charging session, the status of the electricity meter at the beginning and the end of each charging session, and the anonymized identifier of a user who activated the charging session.

We start with preprocessing the dataset. In this step, we remove clearly erroneous data points, such as charging sessions with negative duration. We also merge some short consecutive charging sessions from the same customer, situation that often happens when the car is improperly connected to the charging station. After cleaning and preprocessing, the dataset contains 1513 charging sessions spanning over two years.

The summary statistics of the dataset can be found in Table 4.1. Figure 4.4 shows the histograms of the start time and duration of the charging sessions contained in the dataset. The dataset was collected at the charging station over a period of several weeks. In this period, the charging station averaged 2.53 charging sessions per day. With such low demand for charging, there were almost no conflicts in requested charging sessions. Thus, in our experiments, we randomly sample the dataset to generate single days with up to 60 daily charging sessions.

The particular charging station dataset does not contain any pricing information about the charging sessions for its three charging locations. However, E-WALD (similar to many other charging station operators) used only flat rate pricing in all their charging stations.

### 4.3.2 Split-MDP Pricing Strategy Implementation

To implement the MDP pricing strategy, we discretize a single day into 24 time intervals, each 1 hour long. As the real-world data was collected at a charging station with three

[1] https://e-wald.eu/

| Metric | Description |
|---|---|
| CS Revenue | Revenue of the charging station is the sum of prices of all charging sessions. Revenue is directly dependent on the selected pricing scheme $\Phi$ as given by Equation (4.1). |
| CS Utilization | Measured in hours. It is the added duration of all charging sessions the charging station realizes. This is a proxy of the social welfare of the EV drivers achieved through various pricing strategies. The higher the utilization, the more the EV driver charging demand was satisfied by the charging station. The definition of CS utilization $\mu$ is given by Equation (4.9). $$\mu(c, d, E, \Phi) = \sum_{r_i \in d} (end(r_i) - (start(r_i))\mathbb{1}_{\phi_i(\Phi(r_i))=\top} \quad (4.9)$$ |

**Table 4.2:** Description of the evaluation metrics.

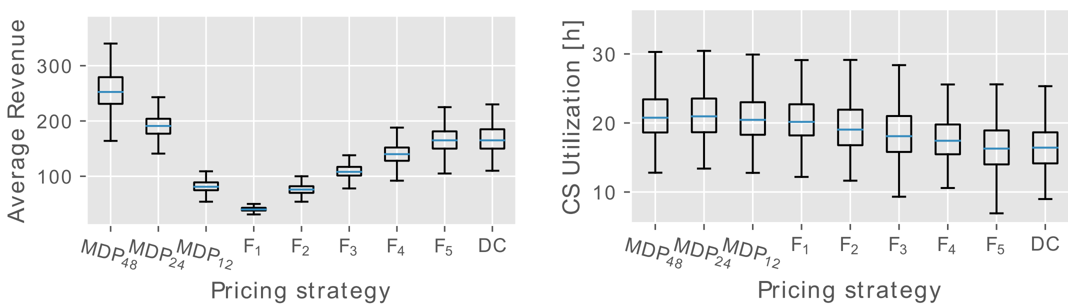


**Figure 4.6:** Experimental results using the evaluation metrics. Each boxplot is a result of 400 independent runs with 40 requests and price elasticity parameter $C = 0.03$ (see Figure 4.5). MDP lower index denotes the number of timeslots used for one day in the Split-MDP dynamic pricing strategy. DC denotes demand correlated strategy and $F_1$ to $F_5$ the different flat rate strategies.

charging slots, we consider our station to have three charging points in our experiments. That is, we use $c_{max} = 3$. This means that, at most, three charging sessions can be realized at any time.

The dataset contains information only about realized charging sessions[2]. In our electromobility model, EV drivers can book charging sessions ahead of time. We model this by randomly setting the request time $req(r)$ for each charging session $r$ in the dataset, with the request time drawn uniformly between 0 and 6 hours ahead of $start(r)$.

Values of $\mathbb{E}[D_t^j]$ are estimated from the dataset. For each time interval $j$ in the discretization of time, we calculate normalized histogram $H_t^j$ of request times $req(r^j)$ of request $r^j$ for which the charging interval $(start(r^j), end(r^j))$ is in $j$th time interval. Bins

[2]E-WALD, same as most of the other existing charging station operators, did not allow booking of charging services ahead of time

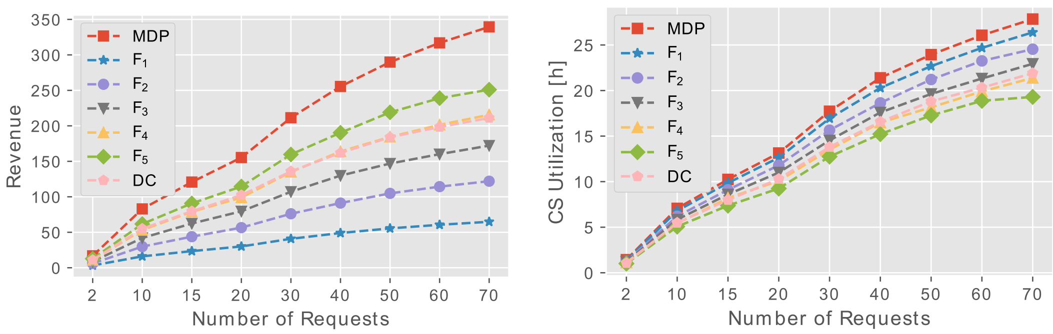


**Figure 4.7:** Performance of the Split-MDP dynamic pricing compared to the performance of the flat rate ($F_1$ to $F_5$) and the Demand Correlated (DC) pricing strategy during one simulated day using prices between 1 and 5 when varying the number of charging requests per day. The curves show the average values based on 400 runs with a random selection of booking requests from the real-world dataset and price elasticity parameter $C = 0.03$ in each run.

of the histogram are the intervals of the time discretization, and the normalization is done with the size of the dataset used. Denoting the *expected number of requests in a day* as $|D|$, we set $\mathbb{E}[D_t^j] = |D| \cdot H_t^j$.

Recall that $\mathcal{E}(h)$ is the probability of the EV driver agent accepting price $h$. Because we do not know the real price-demand curve for EV charging services and we cannot estimate it from available data, we define the *price-demand function* $\mathcal{E}(h)$ parametrically as $\mathcal{E}(h) = e^{-Ch}$. This is one of the standard demand functions in transport economic textbooks [McC01]. The different values of the price elasticity paramter $C$ and the corresponding shapes of price-demand functions are shown in Figure 4.5. The evaluation is done for multiple values of $C$. For $C = 0$, we talk about inelastic demand, as customers will accept any price. At $C = 0.5$, the demand is highly elastic as small changes to the price have a big effect on the customer's acceptance or refusal of the offer. For comparison, the price elasticity of demand for gas station services is usually described as relatively inelastic, meaning low values of the price-elasticity parameter $C$ [LP13]. When data on price elasticity of demand for EV charging becomes available, the parametric price-demand function $\mathcal{E}(h)$ can be replaced by a model learned from the data.

We use the same price-demand function for all drivers in each experiment. In reality, each person will respond differently to changing prices. However, using randomly selected $C$ for each user does not give us different results than using a single $C$ for all users. Using one $C$ for all users means aggregating the population's behavior. Thus, for simplicity, we use one value of $C$ for all users in the experiments described in Section 4.3.3.

Because of the compactness of the elements $\Pi_j$ of the split-MDP pricing, we can find the optimal solutin to each MDP $\Pi_j$ through the policy iteration algorithm [How60].

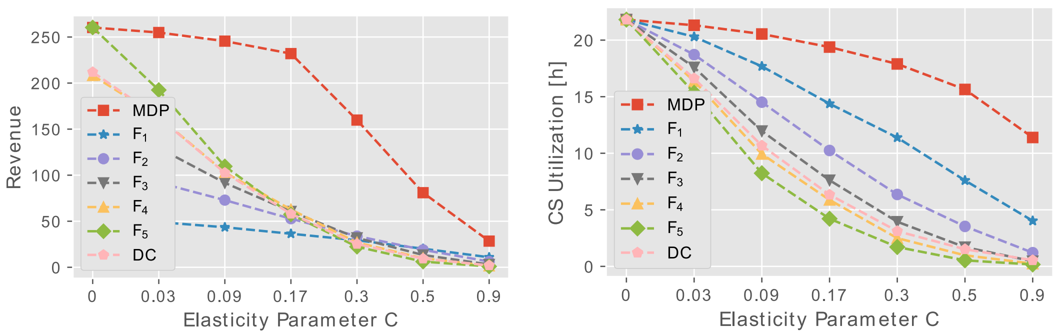


**Figure 4.8:** Performance of the Split-MDP dynamic pricing compared to the performance of the flat rate ($F_1$ to $F_5$) and the demand correlated pricing strategy (DC) during one simulated day using prices between 1 and 5 when varying the price elasticity parameter $C$. The curves show the average value of the metric based on 400 runs with a random selection of 40 booking requests from the real-world dataset in each run.

### 4.3.3 Experiments and Results

In our experiments, we compare the performance of the Split-MDP dynamic pricing strategy to the flat rate and demand-correlated pricing strategies. The flat rate pricing strategy calculates the charging request price similarly to the Split-MDP dynamic pricing strategy, except it uses the same price in all timeslots (see Equation (4.8)). The demand-correlated strategy sets prices based on the expected demand for each time interval. It sets the lowest price for timeslots with the smallest demand, while for timeslots with the highest demand, it sets the highest price. The goal is to equalize the number of sold charging sessions throughout the day.

In most experiments, we use $h_{max} = 5$ and five price levels, 1 to 5. The choice of the maximum price value is tied to the choice of price elasticity parameters. For the given choice of price levels and the average length of a charging session, we chose price elasticity parameters that generate price-demand functions of varied shapes over the domain of probable charging session prices (see Figure 4.5). The number of price levels is selected at 5 as a tradeoff between the fidelity of pricing and the size of the state space in the MDPs. We use the same price levels for the flat rate and demand-correlated strategies.

We use two metrics to compare the performance of different pricing strategies: charging station revenue and charging station utilization time. A detailed description of these metrics is given in Table 4.2.

In each experiment run, we randomly draw the sequence of requests $d$ of fixed length from the real-world E-WALD dataset. These requests are ordered by $req(r)$ and processed in parallel by an MDP for each interval. Each MDP discards some requests due to capacity constraints; the issuing EV driver agents refuse the price for some other requests. Metrics described in Table 4.2 are calculated from requests accepted by the charging stations with prices accepted by the EV driver agents. We perform 400 runs in each experiment and average the resulting metrics.

The runtime of the simulations and the solver implemented in Python is in the order of minutes on the Intel Core i7-3930K CPU @ 3.20GHz with 32 GB of RAM, with most of the time spent on pre-calculation of the policies using policy iteration algorithm (see Chapter 2) for the interval MDPs.

We fixed the price elasticity parameter for the first experiment at $C = 0.03$ and the number of requests to 40. In this experiment, we report the quartiles of the evaluation metrics in Figure 4.6. We can see that for given parameters, MDP dynamic pricing improves revenue. Furthermore, it also improves utilization. The same figure shows the effect of the length of the MDP timeslot on the results. Using shorter timeslots improves revenue but increases the variance of the results. For this reason, we use 24 timeslots per day in the rest of the experiments. Figure 4.6 also shows that the results obtained for the MDP dynamic pricing can be achieved reliably without increasing the variance of the observed metrics over the flat rate pricing.

The selection of the number of price levels was done empirically; the appropriate choice of price levels inherently depends on the concrete values of demand and price elasticity parameter. However, any sufficiently high number of price levels is enough to demonstrate trends in the dependence on price elasticity paramter and demand. Thus, in the following experiments we use only five price levels ($\hat{h} = 5$) for simplicity.

For the second experiment, we fixed the price elasticity parameter to $C = 0.03$ and varied the number of arriving customer requests from 2 to 70. Results of this experiment can be seen in Figure 4.7. As expected, increasing the number of booking requests increases revenue, utilization, and delivered charge for all pricing schemes. Unsurprisingly, the revenue-maximizing MDP strategy comes out ahead of other pricing strategies in terms of revenue. Moreover, MDP dynamic pricing outperforms the other pricing strategies in utilization across all numbers of booking requests. While the MDP dynamic pricing is closest to a flat rate with price 1 in utilization and delivered charge, it is closest to a flat rate with price 5 in revenue and average price per kWh.

For the third experiment, we varied the price elasticity parameter $C$ through values given in Figure 4.5. We fixed the number of requests in $d$, to 40. Results of the second experiment are shown in Figure 4.8. With the increasing elasticity, revenue and utilization decrease. MDP dynamic pricing again comes on top in both metrics.

In the second and third experiments, the revenue of the flat pricing strategy increases as we increase the flat rate. However, by increasing the number $h_{max}$ and the number of price levels above 5, we found that the flat rate revenue decreases after a certain threshold price level dependent on demand and elasticity. At all times, an MDP pricing strategy that uses the same $h_{max}$ achieves multiple times higher revenue.

Notice the utilization and delivered charge in the second experiment are the same for all pricing strategies when $C = 0$; when the demand is inelastic, customers always accept the offered price, and the charging capacity is distributed solely on a first-come, first-serve basis. The downslope trend of the utilization and delivered charge with increasing elasticity were to be expected, given the fixed number of 40 booking requests at an average duration of 0.726 h (the maximal theoretical utilization with three charging points would be $3 * 24$).

## 4.4 Summary

In this chapter, we have presented a Split-MDP dynamic pricing strategy that uses a collection of MDPs to adjust the prices of a charging station dynamically. We evaluated the method in simulation, using data from a real-world charging station augmented by parametric models of customer behavior not captured in the dataset. We have compared the proposed MDP dynamic pricing strategy with the baseline of the most commonly used flat rate pricing across a range of system parameters: the different price-demand functions and volume of demand for charging services. The revenue generated by the proposed dynamic pricing strategy was up to 5 times higher than any flat rate pricing method, with the relative revenue improvement increasing fast as the elasticity increased.

The simulation results show that the Split-MDP dynamic pricing can improve revenue compared to the baselines in all values of absolute daily demand and across all considered price-demand functions, except for completely inelastic demand. The relative increase in revenue is greater if demand is higher or if it is more elastic. While the proposed method maximizes revenue in its default configuration, it also improves the utilization of the charging station resources over the flat rate and demand-correlated baselines through the improved allocation of charging services.

We view the results of the method as promising in that even such a simple dynamic pricing strategy can so heavily outperform the baselines currently used in practice. However, the performed evaluation is limited by three factors:

1. The choice of baselines.
2. Assumptions made in the simulation (the parametric models used to model demand).
3. Abstraction used in the pricing model (the splitting of the MDP).

The baseline used in the evaluation, the flat rate per kWh, is today's most common form of pricing EV charging. However, in the taxonomy of pricing methods, it is the simplest form of pricing available. It is unsurprising that any form of dynamic pricing should outperform it.

The assumptions about customer demand made in the simulation also limit the evaluation. While we are unlikely to determine an accurate model of customer responses to changing prices from obtainable data, we can use more sophisticated demand models. For example, we can use a model where demand for a given timeslot changes in time.

Finally, splitting the MDP model as a method of addressing the large state space is limiting the optimality of the solution. In this abstraction, consecutive charging sessions don't influence each other's prices at all.

We will address these shortcomings in the following Chapter 5. First, we propose a full MDP model that captures the pricing problem more accurately and addresses the state space explosion using a heuristic instead of an exact solution to the MDP. We also extend the demand model to better capture reality and add baselines that let us compare the performance to optimal pricing strategies.

CHAPTER 5

# Generalized MDP model for Dynamic Pricing

The pricing MDP presented in Chapter 4 is oversimplified in several aspects. First, by using separate MDP for each hour (or other interval) of the day, the pricing solution for charging in an interval does not influence prices in other intervals. Similarly, the demand model for each interval is independent, with stationary probabilities of request arrivals.

In this section, we develop a model that addresses these shortfalls. We use a single MDP to generate pricing actions. Since such MDP has a much larger state space, we solve it with heuristic approaches. For the demand modelling, we replace the previous model with a more realistic model of Poisson arrivals which is often used in similar situations. For use in the MDP pricing, we develop a discretization of the Poisson process that allows us to use the MDP pricing and we quantify the discretization error.

This chapter is primarily based on a journal article draft titled “Dynamic pricing of EV charging reservations with MDPs” which studies the problem of allocating EV charging capacity using an improved dynamic pricing scheme than in Chapter 4. We focus on (1) maximizing the revenue of the CS operator and (2) maximizing the overall utilization of the corresponding charging station. Preliminary results were also published in [MB22]. To formulate the pricing problem, we apply the MDP methodology [Bel57].

To derive the optimal solution for the small instances of the MDP problem, we can use exact solution methods such as value iteration (VI), policy iteration, or integer linear programming. However, all these methods suffer from state space explosion problems due to the large-scale nature of the real-world environment. We use a MCTS heuristic solver to approximate the optimal pricing policy to remedy this problem. This is the first usage of MCTS in this kind of problem to the best of our knowledge. Consequently, we contribute to the body of research by applying the theory to the real-world problem of dynamic pricing of EV charging suitable for electric mobility.

Some of our key contributions are:

1. Novel model of dynamic pricing of EV charging problem using the MDP methodology;

2. A heuristics-based pricing strategy based on MCTS, which is suitable for large-scale setups;
3. Optimizations based on maximizing the revenue of the CS operators or the utilization of the available capacity;
4. Experimental results showing that the proposed heuristics-based approach is comparable to the exact methods such as Value Iteration. However, unlike those exact methods, the proposed heuristics-based approach can generate results for large-scale setups without suffering from the state-space explosion problem.

This chapter builds on the previous Chapter 5; therefore, we start by directly describing the new MDP model. Unlike in the previous chapter, we use domain-independent description (Section 5.1). The new model clarifies unstated assumptions made in the previous model by explicitly modeling demand as Poisson process. However, since the Poisson process is defined in continuous time, we need to discretize it to use it in the MDP model. The construction of the model is, therefore, followed by an analysis of the discretization error in Section 5.2. Since the new model is more complex in terms of the number of states, we follow with the description of the heuristic solver based on standard MCTS for MDPs, which uses speedup derived from the properties of the MDP (Section 5.3). Finally, at the end of the chapter in Section 5.4, we describe the experimental evaluation of the model and the MCTS solution method. Since the new model uses a different demand model, the solution is not directly comparable to the results from the previous chapter. However, we can still compare the performance of the MCTS solver to the exact solution methods. We also compare the performance of the MCTS solver to the flat rate baseline and an Oracle solution.

## 5.1 Pricing Problem Reformulation

In this section, we will start by summaraizing the improvements over the Split-MDP model from Chapter 4. This introduction is followed by quickly reintroducing the dynamic pricing problem, highlighting parts relevant to the updated MDP model, and doing so using domain-independent terminology. We will explain how we model the demand and the interaction protocol between customers and sellers. Then, we will present the problem as a MDP and show how the problem maps to the MDP building blocks. The transition function will receive the most attention, where we will use common demand assumptions to simplify and generalize the MDP transition function significantly.

### 5.1.1 Limitations of the Split-MDP Model

We ended Chapter 4 with the summary of the limitations of our Split-MDP model. In this section, we will summarize how we have improved on these limitations.

Starting with the limited abstraction of the MDP model in the form of multiplw “split” MDPs, we now use a single MDP to model the whole day. This means that the pricing decisions made in one interval influence the pricing decisions in other intervals. This is a more realistic model of the problem and addresses the main practical limitation of the previous model.

However, the use of a single MDP model means that we are dealing with a much larger state space where optimal solutions are much more difficult to find. To address this, we use a heuristic solver based on MCTS. This solver can find near-optimal solutions in a much shorter time than the exact solution methods.

This leads us to the baselines we use to compare our solution against. We still use a flat rate as a baseline. The flat rate is now set to maximize the objective based on training data. In smaller instances, we now use Value Iteration to find the optimal solution and compare our heuristic solution against it. Finally, we use an Oracle solution that knows the demand in advance and can, therefore, set the optimal price for each charging request. Unlike the optimal solution, the Oracle solution can be computed even in very large instances, and we use it to track the performance of our heuristic solver.

Finally, in this chapter, we significantly rework the demand model. The demand model described in Chapter 4 is represented in the MDP through the decision tree shown in Figure 4.2. This model is rather compact and efficient. In each timestep $t$, it runs a sequence of Bernoulli trials (coin flips) that represent arrivals of customer charging requests with success probability $1 - p_d$. When the trial is successful, the MDP remains in timestep $t$. With the first unsuccessful trial, the MDP transitions to timestep $t - 1$. The success probability in the model is based on the number of expected requests in timestep $t$, $\mathbb{E}[D_t^j]$, which is set from the observed data.

This model is unrealistic due to the independent treatment of consecutive timeslots of the charging station, $j$ and $j + 1$, as we have already mentioned. However, it also has another limitation. The model is designed so that the expected number of requests in each timestep $t$ (and timeslot $j$) observed by iterating the transition function (Figure 4.2) matches the historical demand data $\mathbb{E}[D_t^j]$. At the same time, the distribution of the number of requests in each timestep $t$ is fixed to be geometric by the design of the model and cannot be changed without fundamentally changing and complicating the model. However, in real-world, the number of requests in each timestep $t$ is not geometrically distributed.

The number of arrivals of customers in a time interval is usually modeled by the Poisson process. While in certain situations the geometric distribution can be a good approximation of the Poisson distribution, it is not the case in most situations. Geometric distribution has a non-increasing probability mass function, while the Poisson process attains maximum away from zero (see Table 2.3) for many values of its parameter. By using the model from Chapter 4, the MDP will, in many situations, effectively overestimate the probability of a low number of arrivals and overestimate the probability of a high number of arrivals in many timesteps.

The appropriate solution is to model the demand as a Poisson process, which is what we do. However, notably, the transition function in Figure 4.2 can not be modified to accommodate the Poisson process without significant changes to its structure. Namely, the loop between the root and one of the leaves has to be removed from the transition function. This allows for only one arrival in every timestep $t$. This means the number of timesteps in the new model will have to be much larger. The changes to the demand model and their implications are analyzed in detail in Section 5.2.

Because of these significant changes to the MDP model, we formulate the MDP again from the ground up in this chapter. This time, we use domain-independent language to describe the MDP. This allows us to describe the MDP in a more general way.

The next section briefly describes a dynamic pricing problem and a general dynamic pricing MDP formulation. The state in this MDP corresponds to a product request at a given time, given the current capacity of resources. The solution to this MDP is a policy that maps states to pricing actions. The optimal policy is a solution that maximizes expected revenue or efficiency, given the probabilistic description of demand.

### 5.1.2 Pricing Problem Description

In our work, dynamic pricing involves setting optimal or near-optimal prices online for different products as customer requests for these products arrive. The seller combines resources with predetermined capacities, such as seats in bus connections or the charging capacity of a set of charging stations, into products offered to customers, such as trips between cities with a transfer or half-day charging of an EV. Although the demand for these services is unknown beforehand, the seller has historical data about customer request arrivals and a model of customer responses to changing prices. The seller's goal is to price each arriving product request in a way that maximizes its objective function, considering demand uncertainty. The objective function can be maximizing either revenue or efficiency, such as for public utilities. This optimization occurs over a finite time horizon, after which the resources can no longer be sold, such as when the bus departs or the day ends.

Here, we give a minimalist description of the pricing problem by considering first the supply side that puts constraints on the seller and the demand side that prompts the seller's actions. The description is illustrated by Figure 5.1.

The supply side of the problem is formed by a set of $n$ resources that can be combined into $m$ products available for sale. Each product is represented by a vector $\boldsymbol{p} \in \mathbb{N}_0^n$, elements of which prescribe the number of individual resources used in the product. The availability of these products is constrained by the initial capacity of the resources $\boldsymbol{c}_0 \in \mathbb{N}^n$ and the lengths of selling periods of different resources $\boldsymbol{T} \in \mathbb{R}^+$ that determines the time after which each resource and product it is part of can no longer be sold.

In our case, the product is a charging session, and the different resources are charging capacities in different time slots.

The demand side of the problem is modeled by a non-homogeneous, compound Poisson counting process $N(\tau)$, $\tau \in (0, \max(\boldsymbol{T}))$ that models the arrivals of requests for different products in time and distributions of finite internal customer valuations for different products, $\{\beta_{\boldsymbol{p}} | \boldsymbol{p} \in P\}$. The customers accept the offered price for the requested product if it is below their internal valuation. Otherwise, they reject the offer and leave the system.

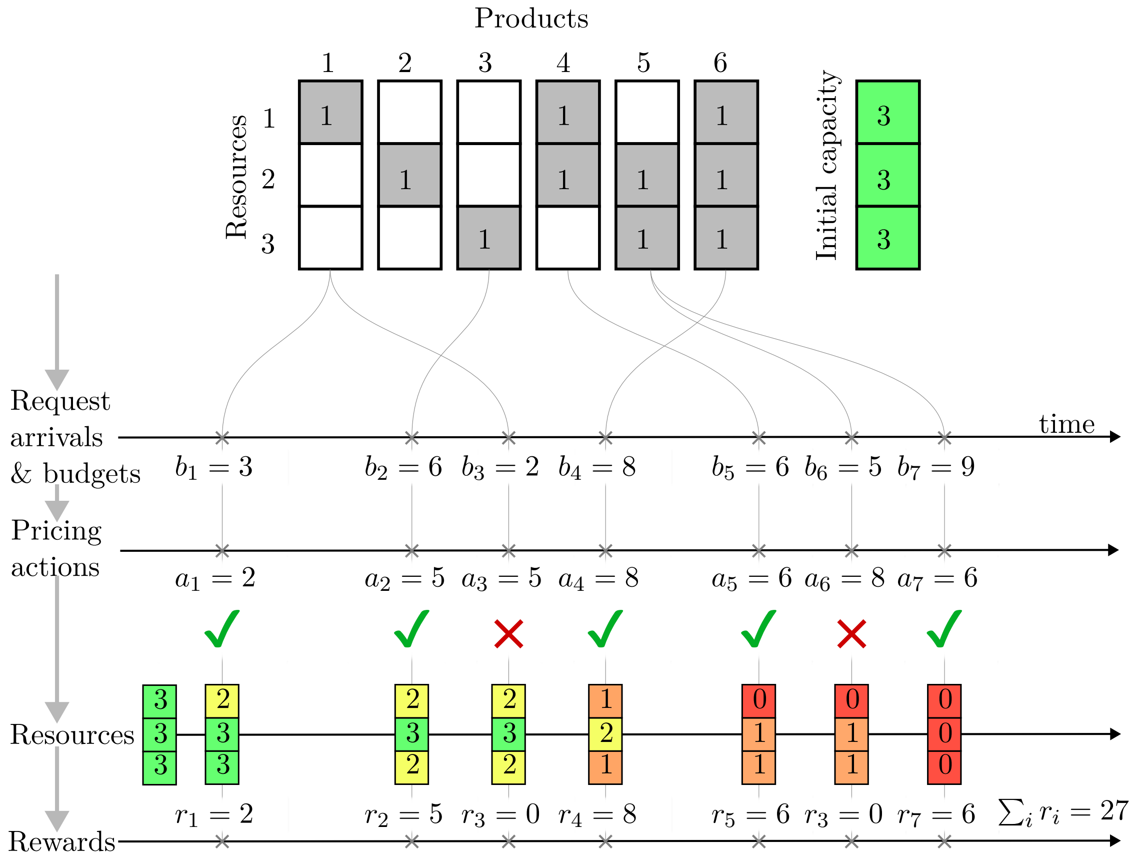


**Figure 5.1:** Illustration of the domain-independent description of dynamic pricing of EV charging. Products are vertical columns in the resource matrix on top. For example, product $[1, 1, 0]$ represents consecutive charging in the first time slot (resource) and second time slot out of the three. In the EV charging problem, we assume each product contains every resource (the charging slot capacity in a time interval) only once and that products contain only consecutive time intervals (i.e., there is no product $[1, 0, 1]$). Below the matrix, the *rquest arrivals and budgets* show customers requesting different products at different times and their budgets ($b_1$ to $b_7$). The next line below shows *pricing actions* $a_1$ to $a_7$ taken by the seller. When a customer accepts the price (i.e., when $b_i \geq a_i$, shown as green ✓), the seller accumulates a reward $r_i$, and his resource capacity is reduced by the product requirements. This is shown on the two bottom lines.

Realized demand takes the form of a sequence of timestamped product requests associated with hidden customer valuations of products, $b_i \sim \beta_{\boldsymbol{p}_i}$:

$$d = ((\boldsymbol{p}_1, \tau_1, b_1), (\boldsymbol{p}_2, \tau_2, b_2), \ldots) \tag{5.1}$$

The service provider needs to process these requests individually as they arrive. The protocol for a single interaction between the seller and customers is following:

1. An $i$th customer requests product $\boldsymbol{p}_i$ from the seller at time $\tau_i$.
2. If the request is feasible, that is, if $\tau_i \leq T_{\boldsymbol{p}_i}$ and the resource capacity is greater than the products resource requirements, $\boldsymbol{c} \geq \boldsymbol{p}_i$ in all components, then the seller proposes price $a \in A \subset \mathbb{R}$ to the customer. Otherwise, the seller rejects the request (equivalent to proposing an infinite price for the product).
3. The customer compares the price $a$ against their internal valuation of the product $b_i$, and if $a \leq b_i$, then the customer pays the price and buys the product. Otherwise, the customer rejects the price $a$ and leaves the system.

Crucially, the seller *does not* know the values of the customer's valuations $b_i$. Each customer requests a single product and leaves the system after accepting or rejecting the offered price.

### 5.1.3 MDP Model

After recounting the dynamic pricing problem, we will now develop the updated MDP model to determine the pricing policy $\pi$ that assigns price $a$ to the product-time pair $(\boldsymbol{p}, t)$ given some state of the seller's resources.

In Our MDP, defined as a 5-tuple $(S, \mathcal{T}, R, A, s_0)$, the seller starts in some initial state $s_0 \in S$. Each state captures the current timestep, what product is being requested, and how many resources are currently available. The seller offers a price $a \in A$ for the requested product, taking an action that results in a transition to a new state $s' \in S$. However, the transition is not deterministic because it is unknown whether the customer will accept or reject the price and what the next product request will be. By fitting the random demand process $N(\tau)$ and distributions of the customer internal valuations $\{\beta_{\boldsymbol{p}} | \boldsymbol{p} \in P\}$ to the historical data, we can estimate the transition probability $\mathcal{T}(s'|a, s)$, which determines the likelihood of reaching state $s'$ when taking action $a$ in the state $s$. The transition between states also generates rewards for the seller, determined by the function $R(s, a, s')$.

Several differences exist between this MDP model and the MDP described in Section 4.2.1. The differences are most easily seen by comparing the transition functions, shown in Figures 4.2 and 5.3:

1. The transition function in Figure 4.2 has a loop between the root and one of the leaves. This loop allows the MDP to remain in the same timestep for multiple

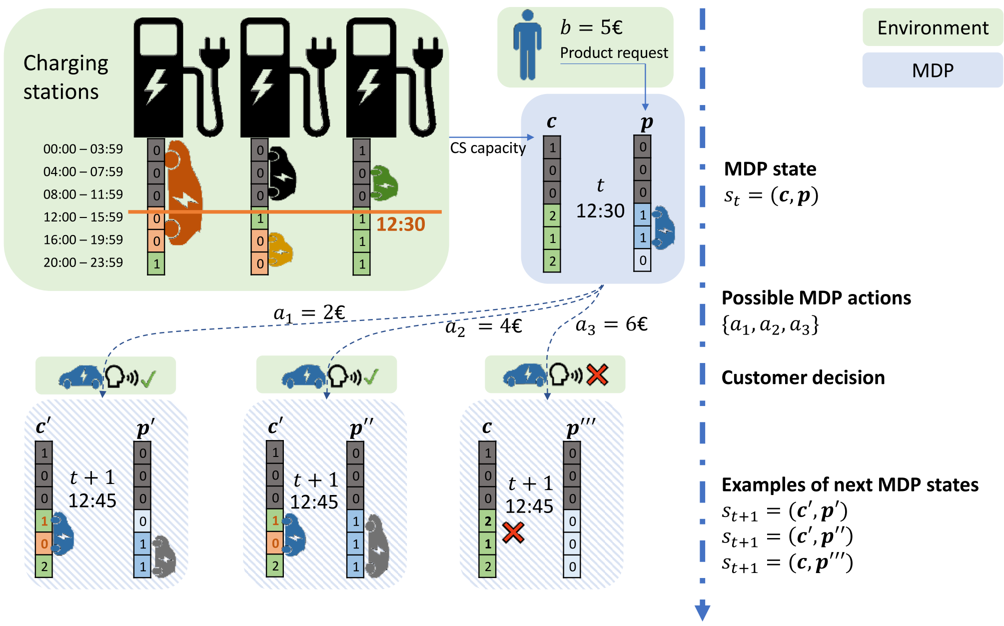


**Figure 5.2:** Illustration of the MDP states for the dynamic pricing of EV charging. Unlike Figure 5.1, this figure illustrates the expiration of resources after their selling period ends (grey in the capacity and product vectors). The blue squares represent the MDP states. At timestep $t$, the capacity of the charging station is expressed by the capacity vector $\boldsymbol{c}_t$. Elements of the vector represent available charging capacity in corresponding timeslots (time ranges in the green square). Possible charging session reservation request arriving since the previous timestep is expressed by the vector $\boldsymbol{p}_t$, with ones representing the requested timeslots. Based on the three state variables $\boldsymbol{c}_t, t, \boldsymbol{p}_t$, the pricing policy provides an action $a$, the price for charging, that the user either accepts (the first two states at the bottom) or rejects (the state on the right). The state then transitions into the next timestep (details of the transition function are illustrated by Figure 5.3). The accepted charging request leads to reduced capacity values. The next charging session reservation is entered into the new state. Note that the timesteps must have much finer resolution than the charging timeslots. The gray color shows past information regarding the charging capacity and session vectors $\boldsymbol{c}_t$ and $\boldsymbol{p}_t$, respectively.

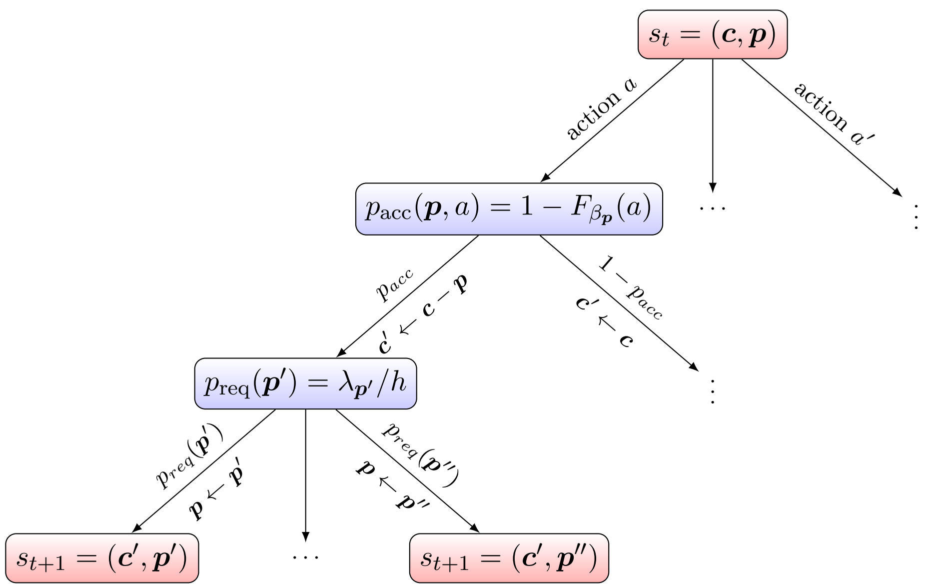


**Figure 5.3:** The structure of the transition function $\mathcal{T}$. Given state $s_t$, the probability of getting to the next state $s_{t+1}$ is given by multiplying the probabilities along the edges. States are the decision nodes (in red), and chance states are in blue and contain the definition of the probability used along the edge. Note that unlike the transition function shown in Figure 4.2, the transition function here does not contain the loop between the root and the leaf, and the system always develops to the next timestep instead of remaining in the same timestep for multiple iterations. We highlight this by having the timestep state variable in the lower index of state $s$ as $s_{t+1}$. Additionally, instead of using price as a state variable, which is increased or decreased through two actions, we use an action for each possible price. Lastly, the requested product in the next time step is added as a state variable.

iterations, making the representation more efficient. This is not the case in the transition function in Figure 5.3.

2. The tree in Figure 4.2 has the price as a state variable and only three price actions (increase, decrease, and unchanged). In Figure 5.3, the price *is* is an action. The action space is, therefore, much larger and covers all considered possible prices.

3. Finally, the states in Figure 5.3 have different structure. First, the requested product that is being priced is a state variable, which makes the state space much larger. Next, price is not a state variable. Finally, the meaning of the state variable $t$ is different. In Figure 4.2, the $t$ is a countdown, possibly reducing by one between consecutive MDP states. In Figure 5.3, the $t$ always increases by one with every transition.

Overall, the new MDP representation is less concise and results in a larger state space, making searching for solutions more difficult. The weaknesses of the previous model that lead to this added complexity are explained in Section 5.1.1.

In the new MDP, the *state space* $S$ consists of states $s = (\boldsymbol{c}, t, \boldsymbol{p})$ (we also use the notation $s_t = (\boldsymbol{c}, \boldsymbol{p})$). That is, the state is defined by the supply of all the resources $\boldsymbol{c}$ at time step $t$ and the product $\boldsymbol{p}$ being requested by some customer at time step $t$. By discretizing the selling period $(0, \max(\boldsymbol{T}))$ into $k$ time steps, we make sure the state space is finite. See Section 5.1.3 for justification of the discretization of time. Thus, $t$ refers to the timestep, the time interval index in the discretization of time. In contrast to the model in Chapter 4, the $t$ increases by one with every transition, so $s_t = (\boldsymbol{c}, \boldsymbol{p})$ is always followed by some $s_{t+1} = (\boldsymbol{c}', \boldsymbol{p})$.

The *action space* $A \subset \mathbb{R}$ is a set of possible prices that the seller can offer customers. This set can principally be continuous, but we assume the prices to form a finite set for our experiments. Tradeoffs made by the discretization are discussed in Section 5.1.3. One special price is the infinite price, which effectively forms the *reject* action since no customer will have infinite funds. This action is used when the seller lacks the resources to provide the requested product.

The *reward function* $R(s_t, a, s_{t+1})$ determines the reward obtained by transitioning from $s_t$ to $s_{t+1}$ by taking action $a$. If the seller's goal is revenue maximization, the reward is the price offered for the product. If it is efficiency, it is the sum of the product's resources, $\|\boldsymbol{p}\|_1$. In the case of revenue, reward has the value of the action $a$ if the customer accepts the offered price $a$ and 0 otherwise. Formally:

$$R(s_t, a, s_{t+1}) = \begin{cases} a, & \text{if } b > a \\ 0, & \text{otherwise} \end{cases}$$

Here, $a$ on the right-hand side of the equation is the value of the action $a$. $b > a$ means that the customers budget $b$ is greater than the price $a$, meaning customer accepts the price. Note that a successful sale implies capacity is reduced between $s_t = (\boldsymbol{c}, \boldsymbol{p})$ and $s_{t+1}$ from $\boldsymbol{c}$ to $\boldsymbol{c} - \boldsymbol{p}$ in $s_{t+1}$, which brings us to the definition of the transition function.

The *transition function* $\mathcal{T}(s_t, a, s_{t+1})$ is the most complex component of the MDP model. It determines the state $s_{t+1}$ the system develops into from state $s_t$ when the service provider takes action $a$. The transition function $\mathcal{T}$ is determined by two factors: the customer arrival processes $D(t)$ and the distributions of customer internal valuations $\{\beta_{\boldsymbol{p}} | \boldsymbol{p} \in P\}$. The structure of the transition function and how it combines these two components is shown in Figure 5.3.

In some state $s_t = (\boldsymbol{c}, \boldsymbol{p})$ (the root of the tree in Figure 5.3), the seller picks a price $a$ for product $\boldsymbol{p}$ requested by a customer with a hidden internal valuation of the product. Since the customer accepts the offered price only if his internal valuation modeled by a random variable $X \sim \beta_{\boldsymbol{p}}$ of the product $\boldsymbol{p}$ is greater than the offered price $a$, the probability of a customer accepting the offered price is given by the complementary cumulative density function $F_{\beta_{\boldsymbol{p}}}$ of $\beta_{\boldsymbol{p}}$ as

$$p_{\mathrm{acc}}(\boldsymbol{p}, a) = P(X > a) = 1 - F_{\beta_{\boldsymbol{p}}}(a) \tag{5.2}$$

This is shown in the second level of the tree in Figure 5.3. The budget distribution $\beta_{\boldsymbol{p}}$ could also be time-dependent, but we assume it is not for simplicity.

Independently of whether the product $\boldsymbol{p}$ is sold at time step $t$ in Figure 5.3, some product $\boldsymbol{p}'$ could be requested at time step $t+1$. The demand model $D(t)$ that determines the probability of a product request at time step $t$ is derived from the compound Poisson counting process $N(\tau)$ with rate $\lambda$. $N(\tau)$ counts the arrival of any request, and is created by merging independent, product-specific Poisson counting processes $N_{\boldsymbol{p}}(\tau)$ with rates $\lambda_{\boldsymbol{p}}$ into a single combined process. For the sake of explanation, for now, we assume these processes are homogenous. That is, the intensity of arrivals does not depend on time. Thus, the intensity of each product subprocess is $\lambda_{\boldsymbol{p}}$ and is constant, and from the properties of Poisson process (see Section 2.4.2), we have $\sum_{\boldsymbol{p} \in P} \lambda_{\boldsymbol{p}} = \lambda$.

However, in our MDP definition, we discretize the selling period $(0, T)$ into $k$ intervals. Assuming for now that each interval in the discretization has constant length $\frac{T}{k}$, we approximate the Poisson process with a discrete demand process $D(t)$, $t \in \{1, 2, 3, \dots, k\}$. $D(t)$ gives the probability of product arrivals in each interval. However, it allows for at most one product to be requested at any interval. $D(t)$ is a multi-class extension of the Bernoulli process with $|P| + 1$ possible outcomes, with $+1$ for no product request arriving in an interval. The probabilities of the different products at timestep $t$ in $D(t)$ are chosen in the following way:

$$p_{\mathrm{req}}(\boldsymbol{p}, t) = \begin{cases} \frac{\lambda_{\boldsymbol{p}}}{k}, & \boldsymbol{p} \in P \\ 1 - \sum_{j \in P} (\frac{\lambda_j}{k}), & \boldsymbol{p} = \emptyset \end{cases} \tag{5.3}$$

See Section 5.2 for details on how the discrete demand process with product request probabilities chosen this way behaves concerning the compound Poisson process $N(\tau)$ and what is the quality of this approximation.

We consider the approximation of $N(\tau)$ by $D(t)$ to be acceptable for a given number of timesteps $k$ and the expected number of requests $\lambda T$ under one condition. $D(t)$ can generate at most 1 request in any interval of the discretization. In contrast, $N(\tau)$ can generate any positive number of requests in *any* interval. Thus, for $D(t)$ to approximate $N(\tau)$ well, we want the expected number of requests assigned to different intervals between $D(t)$ and $N(\tau)$ to be low. Specifically, we want to pick $k$ so that the error term $\frac{Err_2(k,\lambda)}{\lambda}$ is small. For an explanation of $Err_2$ and justification of this approximation, see Section 5.2.

### Notes on the MDP Model

Having described the MDP model, in this section, we justify some of the choices made in its construction, starting with the discretization of time.

While continuous-time MDP formulation [GH09] is possible, it complicates the description of the problem, making it less intuitive and the solution more complex. The arrival time of customer product requests is continuous; the service provider can't influence these arrival times. However, they arrive as discrete events. Additionally, in pricing problems we are interested in, we assume that customers arriving at similar times will mostly exhibit similar types of behavior. That is to say that we expect demand to depend on time in a piecewise continuous fashion with a finite number of discontinuities, where the discretization can be made to match these discontinuities. Additionally, as discussed in the following Section 5.2, the discretization needs only needs to be fine enough for the probability of multiple requests arriving in one timestep to be low.

Next, let us consider the discretization of the action space. While our main solution method, MCTS, can accommodate continuous action spaces [Lee+20] and doing so could improve the sellers' objectives, we opt for finite action space. First, as we could find, most current service providers selling directly to customers offer discrete service prices. It seems a natural simplification to make in our case[1]. The second reason is technical. Where possible, we compare our heuristic solution to optimal baselines only applicable to finite action spaces.

The choice of the Poisson process to model arrival is a natural consequence of the so-called memoryless property that assumes that at any point, the time until the next customer request does not depend on how much time has passed since the last customer request. Fortunately, this assumption holds in many pricing problems since the assumed customer populations are large and customers act independently. For this reason, as well as its simplicity and useful properties, the Poisson process is a popular choice for modeling customer arrivals.

[1]Recent exception in the context of dynamic pricing in network revenue management is Lufthansa, which introduced continuous pricing in 2020 [20a]. The reasons service providers use discrete prices in dynamic pricing are part technical (systems in airline revenue management have historically used discrete prices), part psychological/marketing to simplify customer choices

## 5.2 Properties of the Demand Approximation in the MDP Model

Here, we formalize the definition of the discrete demand process used in our MDP and quantify how well it approximates the assumed Poisson demand process.

### 5.2.1 Convergence of the Discrete Demand Process

In this section, we will assume that the selling period $(0, T)$ is a unit interval $(0, 1)$, which is without a loss of generality through simple rescaling of the timeline. Next, let us describe the Poisson demand process obtained by combining the Poisson sub-processes for each product. We assume there is a Poisson counting process $N_{\boldsymbol{p}}(\tau)$ for each product $\boldsymbol{p} \in P$, defined by the rate $\lambda_{\boldsymbol{p}}$, generating arrival times of requests for that specific product. From the convenient properties of Poisson processes (see Section 2.4.2), the arrival times of all product requests can be considered as coming from a single compound Poisson process $N(\tau)$ with intensity $\lambda = \sum_{\boldsymbol{p} \in P} \lambda_{\boldsymbol{p}}$. The compound Poisson process generates arrivals of requests for any product.

However, since we discretize the time into $k$ timesteps, we have to approximate the arrivals generated by the Poisson process by the discrete Bernoulli process. The Bernoulli process is a sequence of Bernoulli trials, defined by the number of trials $k$ and the probability of arrival of any request in a single trial $p$. The approximation is based on the fact that the Poisson process is a limit of a sequence of Bernoulli processes created by keeping the product $kp = \lambda$ constant and taking $k \to +\infty$, as described in Section 2.4.2.

We can reconstruct the arrival process for individual products by assigning product indices to the arrivals according to the probabilities $\frac{\lambda_1}{\lambda}, \frac{\lambda_2}{\lambda}, \ldots, \frac{\lambda_m}{\lambda}$. Arrival in the Bernoulli process can then be assigned a product type by sampling a discrete distribution with probabilities $\frac{\lambda_1}{\lambda}, \frac{\lambda_2}{\lambda}, \ldots, \frac{\lambda_m}{\lambda}$, one for each product. We refer to the resulting object as *discrete demand process* with $m + 1$ outcomes, where $m$ outcomes correspond to the $m$ possible products and $m + 1$ outcome corresponds to no request being made. We call this process the discrete demand process $D^k_{\boldsymbol{\lambda}}(t)$ in $k$ time steps and demand intensities $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_m]$.

**Proposition 5.1.** *For given $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_m]$, let $D^k_{\boldsymbol{\lambda}}(t)$ be a discrete demand process with $k$ steps,$k \geq \sum_{i=1}^m \lambda_i$, and $m+1$ distinct possible values with outcomes $i \in \{1, ..., m, \emptyset\}$ occurring with probability:*

$$p_i = \begin{cases} \frac{\lambda_i}{k}, & i \in \{1, ..., m\} \\ 1 - \sum_{i=1}^m \frac{\lambda_i}{k}, & i = \emptyset \end{cases}$$

*Then, the sequence of the discrete demand processes $D^k_{\boldsymbol{\lambda}}(t)$ converges with $k \to +\infty$ to the compound Poisson process with arrival intensity $\lambda = \sum_{i=1}^m \lambda_i$ and discrete jump size distribution with probabilities $\frac{\lambda_1}{\lambda}, \frac{\lambda_2}{\lambda}, \ldots, \frac{\lambda_m}{\lambda}$.*

*Proof.* Consider the discrete demand process for a single product, which is a Bernoulli process $B^k(t, p_i)$ with $k$ steps and probability of success $p_i = \lambda_i/k$. This is a sequence of $k$ i.i.d. binary random variables with the probability of success $p_i$. We know (see Section 2.4.2) that with $k \to +\infty$, the sequence $B^k(t, \lambda_i/k)$ converges to a Poisson process with intensity $\lambda_i$. Combining the $m$ Poisson processes with $\lambda_i, i \in \{1, \ldots, m\}$ gives a Poisson process with intensity $\lambda = \sum_{i=1}^m \lambda_i$ with the required jump size distribution by the properties of combined Poisson processes.

Then, the issue is whether these individual Bernoulli processes combine into the required discrete demand process $D^k_{\boldsymbol{\lambda}}(t)$ and whether this process converges to the same Poisson process as the combination of their limits.

We show this using two Bernoulli processes consisting of i.i.d. random variables $X_i$ and $Y_i$, $i \leq k$. The success probability in the first process is $p_1$ and $p_2$ in the second process. These two processes can be combined into a new Bernoulli process with success probability $p = p_1 + p_2$, and a random variable $Z$ that determines the success class in the combined process with class probabilities $\frac{p_1}{p}, \frac{p_2}{p}$. However, unlike in the case of the Poisson process, the arrival classes are no longer independent since success in $X_j$ implies failure in $Y_j$. For separate processes, the event $\{X_j = 1, Y_j = 1\}$ for some $j$ has probability $p_1 p_2$. Nevertheless, since we have $p_1 = \frac{\lambda_1}{k}, p_2 = \frac{\lambda_2}{k}$, the probability is $P(X_j = 1, Y_j = 1) = \frac{\lambda_1 \lambda_2}{k^2}$. Since $P(X_j = 1, Y_j = 1)$ goes to 0 with $\frac{1}{k^2}$ while $p_1$ and $p_2$ go to 0 with $\frac{1}{k}$, the sequence of combined Bernoulli schemes approaches the desired Poisson process in the limit $k \to +\infty$. □

To summarize, we use a discrete demand process in the MDP model. The demand process consists of $k$ die-rolls with $m+1$ sided unbalanced die. Each die roll corresponds to a random arrival of a product request in some timestep. The $m$ sides of the die correspond to the different product requests made in different time steps, and the $m+1$ side represents no request done by the customer in the given step. The proposition means that the demand process defined this way approximates the assumed naturally occurring Poisson arrival processes that generate arrivals of $m$ different products on a real timeline.

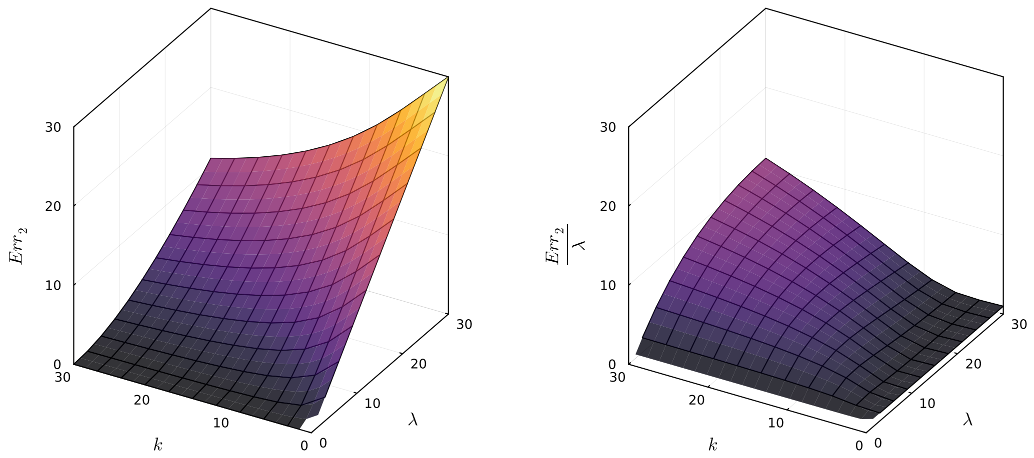


**Figure 5.4:** Visualization of the error term $Err_2(k, \lambda)$ (left) and the relative error $\frac{Err_2(k,\lambda)}{\lambda}$ (right) for different values of $k$ and $\lambda$. $Err_2$ is defined in Theorem 5.2 and represents the number of ignored events caused by the discretization of the Poisson process. With the number of timesteps $k$ approaching 0, the number of ignored events approaches the expected number of events $\lambda$. The relative error term $\frac{Err_2(k,\lambda)}{\lambda}$ (Equation (5.10)) represents the number of ignored events relative to the expected number of events.

### 5.2.2 Approximation Quality

The approximation quality depends on the value of $k$ that represents the number of timeslots the continuous selling period $(0, 1)$ is split into. However, the discrete demand process with the probability of any arrival $p = \sum_{i=1}^{m} \frac{\lambda_i}{k}$ is defined so that the expected number of requests in $k$ steps, which has binomial distribution and expected value $\mathbb{E}(Bin(k, p)) = pk$, *exactly* matches the expected number of arrivals from the Poisson process, $\mathbb{E}(Pois(\lambda)) = \lambda$, in the $(0, 1)$ interval. Therefore, the Poisson and the approximating discrete processes do not differ in this metric.

Thus, the approximation error is more nuanced. By definition, the discrete demand process allows only for 0 or 1 arrival in every timestep. However, the Poisson process can generate more than one arrival in any real interval. This means that the discrete demand model systematically underestimates the probability of more than one arrival (by effectively setting this probability to 0) while simultaneously overestimating the probability of exactly one arrival in every timestep of the discretization. This over- and under-estimation can be understood from the distribution of the number of arrivals in an interval of length $1/k$ in the Poisson process. We will denote this random variable $X$ and note that it is Poisson distributed[2]:

$$P(X = j) = \frac{(\frac{\lambda}{k})^j \mathrm{e}^{-\lambda/k}}{j!}$$

We get the approximation error through Taylor expansion of this term for different values of $j$ corresponding to the probabilities of $0, 1$ and 2 or more arrivals:

$$\begin{aligned}
P(X = 0) &= \mathrm{e}^{-\lambda/k} = 1 - \frac{\lambda}{k} + o(\frac{\lambda}{k}) \\
P(X = 1) &= (\frac{\lambda}{k})\mathrm{e}^{-\lambda/k} = \frac{\lambda}{k} + o(\frac{\lambda}{k}) \\
P(X > 1) &= o(\frac{\lambda}{k})
\end{aligned}$$

Therefore, the approximation error in one timestep is in $o(\frac{\lambda}{k})$ and gets smaller with $k$ increasing. However, we want to ensure not only that the error in one timestep is small but also that the accumulated error over all timesteps is acceptable.

To this end, we define two error metrics. $Err_1$ is the expected number of intervals in the discrete demand process where the Poisson process would have more than one arrival. That is, the expected value of a Binomial distribution $Bin(p, k)$ with $p$ being the probability of 2 or more arrivals from the Poisson process in an interval of length $\frac{\lambda}{k}$.

However, $Err_1$ only gives an expected number of *intervals* where we have a problem with more than one arrival, not the number of arrivals ignored by the discrete demand process, which could be higher. For this reason, we define $Err_2$, which gives the *expected number of arrivals missed by ignoring more than one arrival per timeslot.* By symmetry argument from $\mathbb{E}(Pois(\lambda)) = \mathbb{E}(Bin(k, \frac{\lambda}{k}))$, this is also the expected number of timeslots with one arrival added by the discrete demand process over the Poisson process. $Err_2$ is formally given as:

$$Err_2 = Err_1(k, \lambda)\, \mathbb{E}_{Pois(\lambda/k)}[X - 1 | X > 1]$$

where $X \sim Pois(\frac{\lambda}{k})$ is the Poisson distributed random variable that describes the number of arrivals in an interval of length $\frac{1}{k}$. $\mathbb{E}[X - 1|X > 1]$ is then the expected number of arrivals over one, conditional on more than one arrival.

The following proposition gives formulas for these two kinds of errors.

---

[2] Unfortunately, the terminology here is historically misleading. In the discrete case, we have a sequence of Bernoulli distributed random variables (sequence of coin tosses), which is called a Bernoulli process. The process's number of successes (heads) is distributed according to a Binomial distribution. However, in the continuous case, we have a sequence of exponentially distributed random variables that form a Poisson process, and the number of arrivals in the process that is Poisson distributed.

**Proposition 5.2.** *Let $D^k_{\boldsymbol{\lambda}}(t)$ be a Bernoulli demand process that converges to a Poisson process with intensity $\lambda$ on a unit interval when the discretization of time into $k$ intervals is refined in the limit $k \to +\infty$.*

*Then, the expected number of intervals in which the Poisson process will register more than one arrival is*

$$Err_1(k, \lambda) = k - (k + \lambda)\mathrm{e}^{-\lambda/k} \tag{5.4}$$

*The sum of the expected number of arrivals over one summed across all intervals in the discretization of* $(0, 1)$ *is*

$$Err_2(k, \lambda) = \lambda\mathrm{e}^{-\lambda/k} + (\lambda - k)(1 - \mathrm{e}^{-\lambda/k}) \tag{5.5}$$

*Proof.* Equation (5.4) is easily seen from the definition. The expected value of $Bin(p, k)$ is $kp$, substituting the probability of two or more events in Poisson as $p = 1 - P(X = 0) - P(X = 1) = 1 - \mathrm{e}^{-\lambda/k} - \frac{\lambda}{k}\mathrm{e}^{-\lambda/k}$ and simplifying the expressions immediately provides the result.

To show Equation (5.5), we need to express the conditional expectation

$$\mathbb{E}_{Pois(\lambda/k)}\left[X - 1 | X > 1\right]$$

of a Poisson distributed random variable $X \sim Pois(\frac{\lambda}{k})$:

$$\mathbb{E}\left[X - 1 | X > 1\right] = \sum_{j=0}^{+\infty} jP(X - 1 = j | X > 1) = \sum_{j=0}^{+\infty} j\frac{P(X - 1 = j \wedge X > 1)}{P(X > 1)} \tag{5.6}$$

$$= \frac{1}{1 - P(X = 0) - P(X = 1)} \sum_{j=2}^{+\infty}(j - 1)P(X = j) \tag{5.7}$$

$$= \frac{1}{1 - P(X = 0) - P(X = 1)} \left(\sum_{j=2}^{+\infty} jP(X = j) - \sum_{j=2}^{+\infty} P(X = j)\right) \tag{5.8}$$

$$= \frac{1}{1 - \mathrm{e}^{-\lambda/k} - \lambda\mathrm{e}^{-\lambda/k}} \left((\frac{\lambda}{k} - \lambda\mathrm{e}^{-\lambda/k}) - (1 - \mathrm{e}^{-\lambda/k} - \lambda\mathrm{e}^{-\lambda/k})\right) \tag{5.9}$$

We use the fact that $\mathbb{E}\left[Pois(\frac{\lambda}{k})\right] = \frac{\lambda}{k}$. We get the desired term by multiplying the last line by $Err1$ and simplifying. □

In Theorem 5.1, we show that the discrete demand process converges to the Poisson process when refining the disretization of time. However, it does not tell us about the quality of the approximation for a given $k$. Theorem 5.2 explicitly quntifies the two kinds of errors in the approximation in terms of number if timesteps $k$ and the expected number of arrivals $\lambda$. Since the error depends on the expected number of arrivals $\lambda$, in experiments, we will use the relative error

$$\frac{Err_2(k,\lambda)}{\lambda} \tag{5.10}$$

to quantify the quality of the approximation. The interpretation of this relative error is *number of missed arrivals per expected arrival.*

## 5.3 Dynamic Pricing Algorithm Using MCTS

**Algorithm 3** Dynamic pricing MCTS algorithm for MDPs. Based on [MK16; Ego+17]

1: **procedure** MCTS(state; $c, d_{\max}, n_{\text{iter}}$)
2: $\quad n \leftarrow 0$
3: $\quad$ **while** $n < n_{\text{iter}}$ **do**
4: $\quad\quad n \leftarrow n + 1$
5: $\quad\quad s_0 \leftarrow$ state
6: $\quad\quad r_0 \leftarrow 0$
7: $\quad\quad d \leftarrow 0$
8: $\quad\quad$ **for** $i = 1$ to $d_{\max}$ **do** ▷ *Selection-Expansion loop to max. tree depth* $d_{\max}$
9: $\quad\quad\quad$ **if** $s$ not encountered yet **then** ▷ *Expansion*
10: $\quad\quad\quad\quad n_s \leftarrow 0$
11: $\quad\quad\quad\quad$ **for** $a \in A$ **do**
12: $\quad\quad\quad\quad\quad n_{s,a} \leftarrow 0$
13: $\quad\quad\quad\quad\quad q_{s,a} \leftarrow 0$
14: $\quad\quad\quad d \leftarrow i$
15: $\quad\quad\quad$ **if** $\exists a \in A$ s.t. $n_{s,a} = 0$ **then** ▷ *Avoids undefined operation on line 18*
16: $\quad\quad\quad\quad a_i \leftarrow a$
17: $\quad\quad\quad$ **else**
18: $\quad\quad\quad\quad a_i \leftarrow \arg\max_{a \in A} q_{s,a} - c\sqrt{\frac{\ln(n_s)}{n_{s,a}}}$ ▷ *Selection (using UCB1)*
19: $\quad\quad\quad s' \leftarrow \mathcal{T}(s'|a_i, s)$ ▷ *Sample execution of action a in s*
20: $\quad\quad\quad r_i \leftarrow r_{i-1} + R(s, a_i, s')$ ▷ *Cumulative reward received up to iteration i*
21: $\quad\quad\quad s_i \leftarrow s$
22: $\quad\quad\quad s \leftarrow s'$
23: $\quad\quad\quad$ **if** $s$ terminal **or** $n_{s,a_i} = 0$ **then**
24: $\quad\quad\quad\quad$ **break**
25: $\quad\quad r_d \leftarrow r_{d-1} + \text{ROLLOUT}(s)$ ▷ *Value estimation*
26: $\quad\quad$ **for** $i = 1$ to $d$ **do** ▷ *Backpropagation*
27: $\quad\quad\quad q_{s_i,a_i} \leftarrow \frac{n_{s_i,a_i} q_{s_i,a_i} + (r_d - r_{i-1})}{n_{s_i,a_i} + 1}$ ▷ *Average* $q_{s,a}$ *with reward from levels below*
28: $\quad\quad\quad n_{s_i} \leftarrow n_{s_i} + 1$
29: $\quad\quad\quad n_{s_i,a_i} \leftarrow n_{s_i,a_i} + 1$
30: $\quad$ **return** $\arg\max_{a \in A} q_{s_0,a}$

This section describes the method we use to derive the dynamic pricing policies. Our solution method of choice for large-scale problems is MCTS. Unlike tabular methods,

**Algorithm 4** Rollout algorithm used in Algorithm 3

```
procedure ROLLOUT(state = (c, t, p))
    s ← state
    r ← 0
    while s is not terminal do
        a' ← select random action from A
        s' = (c', t', p') ← T(s'|a_i, s)        ▷ sample result of action a' in s
        r ← r + R(s, a', s')
        Δt ← sample inter-arrival time distribution of demand at timestep t
        p' ← sample product request at t + Δt
        s ← (c', t + Δt, p')
    return r
```

such as Value Iteration (VI) or policy iteration, MCTS does not need to enumerate the whole state space. Instead, it looks for the best action from the current state and expands only states that the system is likely to develop into. However, unlike VI, for every state, MCTS only approximates best actions. MCTS improves its approximations of best action with the number of iterations. See Section 2.5.4 for more background information.

Nonetheless, it can be stopped anytime to provide currently the best approximation of optimal action. These properties make it a helpful methodology in dynamic pricing. With MCTS, we can quickly apply changes in the environment to the solver. Even in large systems, the price offer can be generated quickly enough for a reasonable customer response time. To the best of our knowledge, this is the first attempt to solve the EV charging dynamic pricing problem using MDP and MCTS.

In our MCTS implementation, we use the popular Upper Confidence-bound for Trees (UCT) variant of MCTS [ACF02; Cou06; MK16] that treats each node as a bandit problem and uses the upper confidence bound formula to make the exploration-exploitation trade-off. We split the presentation of the algorithm into two parts. The first part covers the first two and the last step in Figure 2.14, the selection and expansion that form the tree policy, and backpropagation. This part of the algorithm is shown in Algorithm 3. The second part of the algorithm is the rollout policy, which is shown in Algorithm 4.

### 5.3.1 Tree Policy and Backpropagation

The input of Algorithm 3 is the current state for which we seek to estimate the best action $a$. The algorithm has three parameters, the exploration constant $c$, the number of iterations $n_{\text{iter}}$, and the tree depth limit $d_{\max}$.

The input state corresponds to a tree's root from which the MCTS algorithm builds the search tree. Each tree node corresponds to a state $s$ of the MDP, and for each node, we keep track of how many times was the given node visited, $n_s$, how many times was action $a$ used in given state, $n_{s,a}$ and a running average of the q-value of each state-action pair, $q_{s,a}$. These values are iteratively updated during the run of the algorithm.

Each iteration of the MCTS algorithm works in 4 stages, as shown by Figure 2.14. The existing tree nodes are first traversed in the *Selection* step (Lines 8-24). Each tree node selects an action using the UCB formula (Line 18). The exploration constant $c$ sets the appropriate exploration-exploitation trade-off between selecting actions with high q-value averages and actions that were not yet tried often. The node to transition to is determined by sampling the MDP transition function $\mathcal{T}$ for the outcome of action $a$ in the state $s$ (Line 19)[3]. If the resulting state has already initialized tree node variable $n_s$, the same process is repeated from this node. This process is repeated for $d_{\max}$ iterations or until a node corresponding to a terminal state is reached or until a state that does not yet have a corresponding tree node is encountered.

When previously unseen state is encountered, the tree is *Expanded* with a new node corresponding to this state (Line 9). Additionally, the stat-action counters and q-values are initialized for this node, for all possible $a \in A$. However, for the freshly expanded node, the exploration term in the UCB formula on Line 18 has undefined value. Therefore, if there is an untried action in some state $s$, we first try any unused action $a$ with $n_{s,a} = 0$ (Line 16). Additionally, when we encounter an untried action, we terminate the selection-expansion stage, even if the tree depth limit has not been reached yet (Line 23). As a result, our MCTS implementation proceeds in the breadth-first manner.

During the selection phase of the algorithm, the cumulative reward collected up to $i$th iteration is recorded (Line 20). The estimated value of the newly explored state-action pair $q_{s,a}$ is then calculated in the *Value estimation* stage of the iteration using the rollout (Line 25).

The fourth stage of the MCTS iteration is the *backpropagation* (Lines 26). Here, the cumulative rewards are used to update the average q-value of the node on the $i$th level of the tree, $q_{s_i,a_i}$. The value is updated with $r_d - ri - 1$, *all* reward collected *below* the $i$th level in the tree, including the reward from the rollout. The update is done by averaging the previous update values with the new value (Line 27).

The four-stage (selection—expansion—value estimation—backpropagation) iteration visualized in Figure 2.14 is repeated as many times as allowed. Finally, when the main loop terminates, the algorithm estimates the best action based on the q-value of actions in the tree's root (Line 30).

## 5.3.2 Rollout Policy

Algorithm 4 presents the second part of the MCTS algorithm, the rollout. It is applied from state-action pairs that were used for the first time in the tree policy. The rollout approximates the reward of selected action by quickly reaching terminal state. In our

[3]The description of the algorithm may lead one to expect that during the selection phase of the algorithm, each selection leads to a new state and one level lower in the tree. This is not necessarily true for all MDPs since the transition from state $s$ in $i$th step in the selection loop can result in a state encountered in some previous step. However, this cannot happen in our finite-horizon MDP as it includes timestep as a state variable.

experiments, we use the uniformly random rollout policy that applies random actions until a terminal state in the MDP is reached.

Because we use Poisson process (Bernoulli processes after discretization) as customer arrival processes, we can speed up the rollout by sampling the time to the next arrival from the inter-arrival distribution which is geometric in the discretization (Line 8 in Algorithm 4). Therefore, we can arrive at the terminal state in fewer steps. In the rollout, we simulate actions without storing their resulting states until reaching a terminal state. When terminal state is reached, the rollout terminates immediately, returning the accumulated reward.

### Implementation

Our implementation is based on the MCTS implementation in the POMDPs.jl[Ego+17] library that uses recursion when traversing the tree, unlike the description in Algorithm 3. We provide the unrolled iterative description for clarity.

The implementation we use reuses the constructed decision tree between the steps of the "real" MDP that happen outside of the MCTS algorithm, improving the convergence speed. In our experiments, we build the tree to the maximum depth $d_{\max} = 3$ with exploration constant set to $c = 1$. The number of iterations is capped at $n_{\text{iter}} = 800$. We find that these low numbers are sufficient for sufficiently good performance in our experiments and result in a reasonable computation time.

## 5.4 Experiments and Results

This section presents the experiments carried out with the proposed MDP-based pricing model and MCTS solver described in the previous sections. We compare our solutions against multiple baselines on artificially generated problem instances modeled on a real-life charging station dataset. The instances are generated from the same dataset as the ones described in Section 4.3.1; however, they use the original data differently.

The main goal of the experiments is to demonstrate the viability of the MCTS dynamic pricing algorithm for EV charging. To show this, we run a number of simulations with problem instances created using different problem parameters and compare the average performance of the MCTS algorithm with baseline methods.

The instances are generated by sampling the EV user charging requests from the dataset and sampling the user budgets from the user budget distribution for the given charging session length. Each problem instance has a form of a charging request sequence, as shown by Equation (5.1). The requests cover one day, from 00:00 to 23:59.

Each pricing method, including the baseline methods described below, then prices the requests in the order they come in. The accumulated reward is averaged across multiple simulated runs to measure the performance of each method for comparison with other methods.

The best way to evaluate dynamic pricing methods would be to deploy them in a real-world setting and compare their performance with other pricing strategies. This approach is rarely feasible in research as it requires the opportunity to experiment with real consumers using real products. Realistic pricing experiments with customers are difficult to reproduce since many external factors affect purchasing behavior.

Additionally, directly comparing the performance of our method with other dynamic pricing methods is difficult because all published, readily accessible dynamic pricing methods have been known to use restrictive assumptions on the underlying problem or incompatible models and generally are not as flexible as the MCTS-MDP-based approach. For example, although the approximate dynamic programming solution proposed in [KZZ19] can provide optimal pricing solutions in large problem instances, it only does so with restrictive assumptions on the problem, such as allowing only for linear demand models. Another issue is that there are no established benchmark datasets for comparing the performance of dynamic pricing strategies so far. That said, we can still obtain valuable insights into the performance of our MCTS heuristics-based pricing algorithm by comparing it to well-defined baseline methods.

## 5.4.1 Baseline Methods

Because of the difficulties of evaluating the dynamic pricing policies, we evaluate our proposed MCTS solution against three baseline methods: *flat rate*, *MDP-optimal VI*, and *Oracle* pricing methods. The flat rate represents the lower bound on the revenue we might expect from a dynamic pricing solution. The VI baseline returns an optimal pricing policy and represents the best possible pricing method for the MDP model. Finally, the Oracle policy represents the unachievable upper bound on dynamic pricing performance. Oracle provides the best possible allocation by assuming the CS operator has a perfect knowledge of future requests and EV users' budgets, which is unrealistic in real-world use cases.

### Flat Rate

This baseline provides a lower bound for our MCTS pricing method and a reference for showing how much improvement dynamic pricing could bring. The flat-rate is a single flat price per minute of charging that is used for all charging requests.

The flat price is determined from a set of "training" pricing problem instances that take the form of sequences shown in Equation (5.1). In each each problem instance, we evaluate every possible flat-rate price that corresponds to an action in the MDP, and measure the resulting revenue. The price that maximizes the average revenue across all training sequences is then used as the flat-rate price in the testing simulation runs. This pricing method still uses reservations and the reservations are allocated in in sequence, resulting in first-requested, first-reserved allocation.

### Value Iteration

The optimal MDP policy generated by a VI algorithm is our second baseline pricing method. VI is a simple yet accurate method for solving MDPs that converges to an optimal policy for any initialization. See Section 2.5.3 for a detailed method description. The advantage of VI is that it quickly converges to a complete near-optimal pricing policy at the cost of enumerating the whole state space in memory.

Based on the structure of our MDP state, the state-space size of our MDP model is $kc_0^n 2^n$, where $k$ is the number of timesteps, $n$ is the number of charging timeslots and $c_0$ is the initial charging capacity. If we limit the reservations to contain only the contiguous timeslots, as we do, the state-space size reduces to $kc_0^n n(n+1)/2$. This gives VI an exponential space complexity in the number of timeslots. Thus, it does not scale well to larger problem instances. Therefore, we use VI only to obtain optimal policies on smaller problem instances to validate the heuristic approach of MCTS.

Note that there are other exact solution methods for MDP problems than VI, such as policy iteration used in Section 4.3 or linear programming. All these methods can provide the same optimal pricing policy as VI. However, just like VI, all these methods require enumeration of the whole state space. Our choice of VI is, therefore, arbitrary in this sense.

Since VI gives optimal pricing policy, we use it to benchmark the performance of our MCTS approach, which, since it is heuristic, is expected to provide worse results. How much worse is the question we want to answer by comparing the performance of the two methods on small-enough instances that VI can solve.

### Oracle

Finally, we compare our MCTS-based pricing method against the *Oracle* baseline strategy. It's important to note that the Oracle strategy, while used for benchmarking, isn't practically applicable due to its nature. Unlike other pricing strategies, Oracle relies on having prior knowledge of the entire request sequence and the budgets of EV users to determine prices.

Using this knowledge, Oracle maximizes the optimization metric to provide a theoretical upper bound on the revenue and resource usage achievable by any pricing-based allocation strategy. It works for large and small problem instances; therefore, we can use it to track the performance of MCTS across a wide range of problem sizes.

The Oracle pricing solution is obtained from a linear program. For $k$th sequence of charging requests $d_k$ with requests indexed by $i$, the optimum revenue is the result of a simple binary integer program:

$$\text{maximize} \sum_{i \in \{1 \dots |d_k|\}} x_i \lfloor b_i \rfloor_A, \text{ subject to:} \tag{5.11}$$

$$\sum_{i \in \{1 \dots |d_k|\}} x_i \boldsymbol{p}_i^j \leq \boldsymbol{c}_0^j \qquad j = 1, \dots, |R| \tag{5.12}$$

$$x_i \in \{0, 1\} \qquad i = 1, \dots, |d_k|.$$

where, $x_i$ are the binary decision variables that determine which requests from $d_k$ are accepted by the CS operator. In the objective function (Equation (5.11)), the term $\lfloor b_i \rfloor_A = \max_{a \in A, a \ \leq b_i} a$ denotes the fact that the budget values in the sequence $d_k$ are mapped to the closest lower values in the action space $A$. Conditions (Equation (5.12)) mean that the accepted charging sessions have to use fewer resources than the initial supply $\boldsymbol{c}_0$.

### 5.4.2 Problem Instances

The experimental evaluation is based on the dataset described in Section 4.3.1. For our problem instances in this chapter, we use charging data to instantiate the models of charging session starts and duration. We model the session starts using a normal distribution and the charging session duration using the exponential distribution (see histogram in Figure 4.4). In the dataset, the start times and duration appear uncorrelated (Pearson's correlation coefficient = 0.008), so we assume these distributions to be independent. Furthermore, the charging sessions do not go beyond 23:59; we also make this assumption in our problem instances.

A notable difference from the previous model is that we only allow for the sale of full timeslots. In the previous model, two charging sessions could share the same capacity unit in a timeslot if the first ends before the second one starts, as illustrated in Figure 4.3. In the new model, once a unit of capacity in a timeslot is sold, no other charging session can use it. This is done mainly for simplicity; the new model could be modified to calculate capacity on the fly based on continuous start and end times instead of timeslots. However, this is more complex and makes interpretation more difficult. Additionally, the new model uses finer discretization of timeslots where we assume some buffers between charging sessions will be required in practical application, so stacking reservations back to back within short timeslots does not make practical sense.

The start time and duration distributions are essential parts of the demand model used in the MDP demand model. However, since the datasets captured only realized charging sessions, we do not have any data on the time distribution between EV user's charging request and the requested charging time (the lead time distribution). For simplicity, we assume this distribution to be uniform and independent of the other two distributions.

The distribution of charging session start time, duration, and lead time let us fully define the demand model. To generate the EV charging request sequences (as in Equation (5.1)), we only miss the customer budget distribution. Since we do not have any data on the budget distribution (budget per minute of charging), we model it using a normal distribution, a common approach in the absence of other information, using the central limit theorem as the justification [WB18]. However, our MDP model does not rely on the properties of any of the selected distributions, so any distribution would work.

The budget distribution is unitless, implying no particular currency. To connect the demand side with the supply side, we use the budget distribution to determine the price levels in the MDP action space so that the actions cover the whole range of possible budgets for charging sessions of all durations. Having described the start time, duration, and lead time distributions and the user budget distribution, we can sample the EV pricing problem instances in the form of charging request sequences from Equation (5.1).

Given the distribution parameters, we leave only one free parameter for the problem instances. This parameter sets the level of absolute demand in a given instance. However, a high demand for the CS can be formed with either a small number of long-duration charging sessions or a large number of short charging sessions. Therefore, we use the expected total duration of requested charging timeslots as the scaling parameter, which makes the problem instances comparable across different demand structures. The expected duration of requested timeslots is the demand scaling parameter used in all problem instances.

### 5.4.3 Results

In our experiments, we use multiple pricing problem instances to show the scalability and competitiveness of our MCTS approach. Dynamic pricing is usually not as useful in low demand when requests don't conflict, and resources are underutilized. Pricing as a signal of resource scarcity only makes sense when there is more demand than available resources. Therefore, in our first experiment, we analyze the performance of the pricing methods with increasing overall demand. We do this by increasing the demand scaling parameter.

The second experiment focuses on a technical aspect of our MDP model, the fineness of the discretization of the charging timeslots. Also, since the instances in the first experiment cannot be solved by the optimal VI method, we use the second experiment to compare the performance of the MCTS method with the optimal VI baseline.

#### Implementation

The experiments were implemented in the Julia programming language, using the POMDPs.jl framework [Ego+17]. The framework also provided the VI algorithm[4] we uses and the basis for the MCTS implementation[5]. The development of the MCTS solution was aided by the use of the D3Trees.jl package[6] that we have extended to interactively load and visualize subtrees in the MCTS search tree. The Oracle pricing problem was implemented in Julia and solved with the Gurobi mathematical optimization toolbox [Gur23].

The experiments were run on a single core of an Intel(R) Core(TM) i7-3930K CPU @ 3.20 GHz. The MCTS algorithm with parameters as described in Section 5.3.2 takes on average at most 9 ms to generate the estimate of the best action in any state of any problem instance discussed in this work. The flat-rate prices were trained using 25 randomly sampled sequences from the set 100 problem instances used in the evaluation.

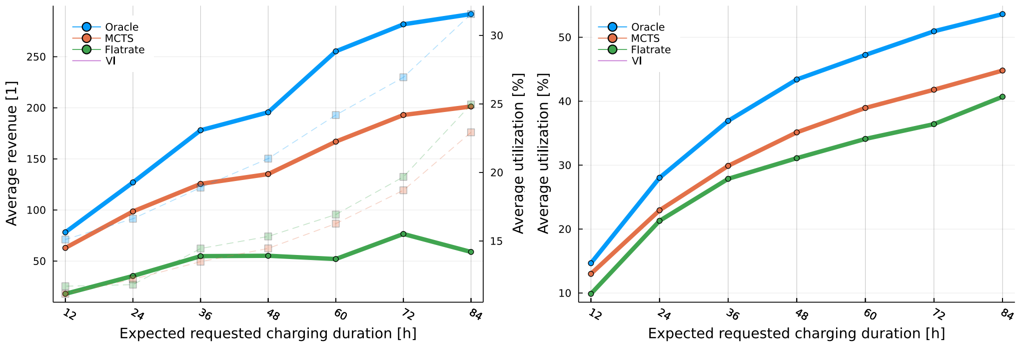


**(a)** Revenue maximizing experiment.

**(b)** Utilization maximizing experiment.

**Figure 5.5:** Experiments with increasing demand for EV charging. The right-hand y-axis and thin plot lines in figure 5.5a show the achieved utilization of each method (when optimizing revenue). Here, the number of timeslots is fixed at 48, meaning 30 min charging timeslots, while the expected demand increases from 1/6th of the available charging capacity (3 chargers with 48 charging timeslots each) to 7/6. The VI baseline could not compute pricing policy for any of the instances in this experiment.

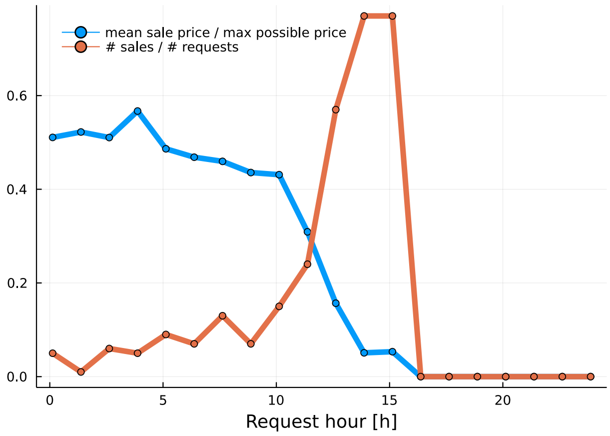


**Figure 5.6:** Pricing response of MCTS to the states $s = (\boldsymbol{c}_0, t, \boldsymbol{p})$ where the capacity $\boldsymbol{c}_0$ and request $\boldsymbol{p}$ are fixed. The timestep of the request $t$ is varied from 00:00 to 23:59. The requested product $\boldsymbol{p}$ is a charging session lasting 1 hour, starting at 16:00, and the capacity $\boldsymbol{c}_0$ is fixed to 3. The figure shows average MCTS price divided by a maximum possible price ($\max(A)$) and average number of accepted price offers from 100 randomly sampled user budgets.

### Variable-Demand Experiment

In the first experiment, we fix the number of timeslots to 48, resulting in each timeslot lasting 30 min, which corresponds to the peak (mode) in Figure 4.4b. For the instances, we vary the total expected duration of requested charging timeslots from 12 h to 84 h between the different problem instances. Since the charging station is selling 24 hours worth of charging sessions from 3 chargers, 84 hours worth of charging requests is well above the capacity of the charging station, especially considering that most charging sessions start in peaks shown in Figure 4.4a. The lower limit of 12 h corresponds to a low demand situation, with the expected duration of requested timeslots being only 17% of the total CS capacity of 72 charging hours.

We sample 100 charging request sequences from each problem instance and average the results. When optimizing for revenue, these results are shown in Figure 5.5a, while the optimization for utilization is in Figure 5.5b.

MCTS outperforms flat rate in both experiments, with increasing gains as demand increases. When optimizing for revenue, the revenue of MCTS is much greater than that of flat-rate, while the utilization remains comparable (Figure 5.5a).

Using the same problem instance with total expected demand of 48 charging hours, we also illustrate the pricing response of MCTS to different states in Figure 5.6. The figure demonstrates the price changes offered by the MCTS for a one-hour long charging session at 16:00 when the request is is priced at timesteps increasingly closer to the start of charging. At the start of the selling period at 00:00, the price is quite high. However, as the start of the charging session approaches, the price is reduced to almost 0. After 16:00, the charging session unnder consideration has already started and the zero price signifies impossible requests. The second curve in the figure shows the driver response to the changing prices, with acceptance rate increasing as the price decreases.

Overall, the MCTS improves revenue and utilization as demand increases, shadowing the performance of the Oracle baseline. Unfortunately, the VI baseline could not compute pricing policy for any of the instances in this experiment. However, we consider the VI in the second experiment, where we vary the state-space size.

### Fixed-Demand Experiment

In the second experiment, we look at the influence of fineness of the MDP discretization on the MCTS performance, given the distribution of charging session durations as represented by Figure 4.4b. We do this by fixing the demand scaling parameter to a fixed expected duration of requested charging timeslots, and we vary the number of charging timeslots from three, 8 h long timeslots to 24, 1 h timeslots.

However, varying the number of timeslots in the 24 h period has a twofold effect. First, since many sampled charging sessions are much shorter than 8 h (see Figure 4.4b), reducing

[4] Package DiscreteValueIteration.jl

[5] Package MCTS.jl

[6] Package D3Trees.jl

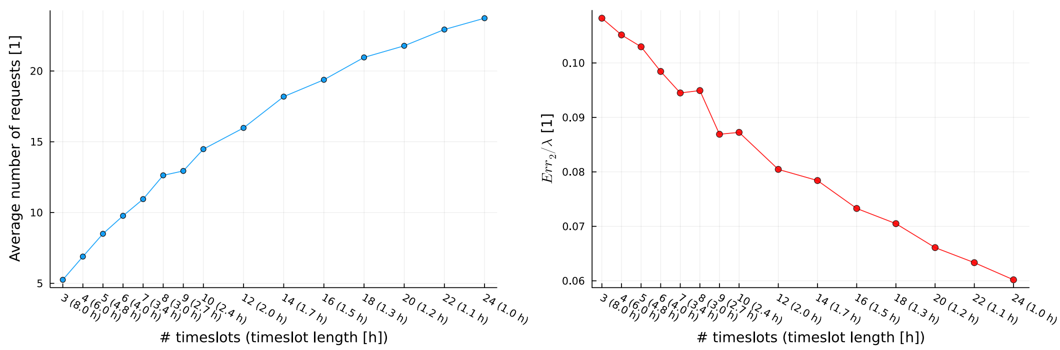


**(a)** Average number of requests. **(b)** Average relative error $Err_2/\lambda$.

**Figure 5.7:** Experiment with constant demand scaling parameter (the expected duration of all requested charging timeslots), but a varying number and length of these timeslots. Here, we use the average number of requests as $\lambda$ and number of requests as $k$ to calculate the $Err_2(\lambda, k)$ according to equation Equation (5.5).

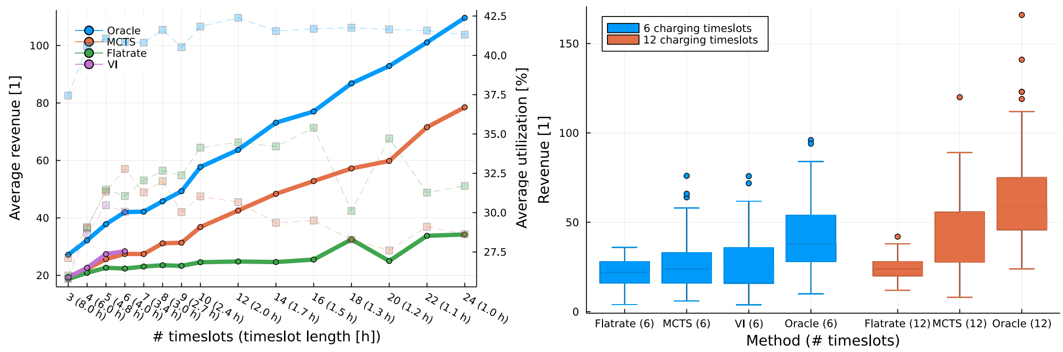


**(a)** **(b)**

**Figure 5.8:** Performance of *revenue* maximizing pricing with fixed demand duration and increasing number of timeslots and requests. The plot of mean values for all instances is in (5.8a) while the boxplot (5.8b) shows problems with 6 and 12 timeslots only. The right-hand y-axis and thin plot lines in the left figure show the utilization of each method (that was not the optimization criterion). In this experiment, we have fixed the expected number of charging hours while varying the number of timeslots in the 24 h selling period across the different problem instances.

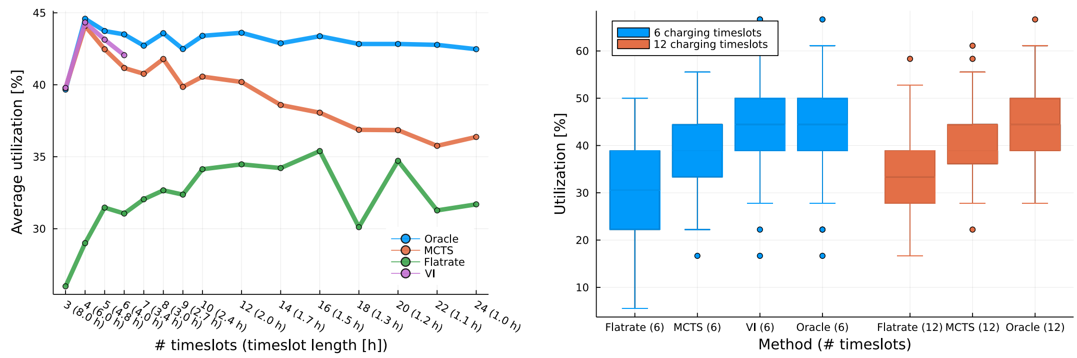


(a) (b)

**Figure 5.9:** Performance of *utilization* maximizing pricing with fixed demand duration and increasing number of timeslots and requests. The revenue is not reported. Figure (5.9a) shows mean values for all instances, while the boxplot (5.9b) shows problems with 6 and 12 timeslots only. We have fixed the expected duration of requested charging timeslots while varying the number and length of timeslots in the 24 h selling period across the different problem instances.

the timeslot length while keeping the expected charging duration of all requested timeslots constant increases the number of requests. For example, three sampled requests for 30 min charging sessions will result in three 8-hour timeslots in the crudest timeslot discretization of 3 timeslots with a total duration of 24 h. However, in the finest discretization of 24 timeslots, the same three requests will result in 3 timeslots. So, when sampling charging duration distribution based on Figure 4.4b for charging session lengths, we get more requests in the finer discretization if we force the expected duration of requested timeslots to be the same. This is shown in Figure 5.7a.

The second effect is that the finer discretization of the timeslots means that the MDP state space is larger. This means that the VI algorithm can only solve instances with long timeslots.

In the second experiment, we generate 15 parametric problem instances with different charging timeslot sizes (between 8 h and 1 h, see x-axis in Figure 5.7a) and a different number of timesteps (each instance having the number of timesteps equal to the number of timeslots times 8). We set the demand scaling parameter, the expected duration of requested timeslots so that the total requested charging duration of all charging sessions is 2/3 of the total charging capacity. Note that the charging location has three charging points capable of serving three simultaneous charging sessions.

In this scenario, we first optimize for revenue and utilization with each method. For each problem instance, we again sample 100 charging request sequences and calculate the revenue and utilization of every method. The averaged results of these experiments when optimizing revenue are given in Figure 5.8, while the utilization-optimizing results are shown in Figure 5.9. In both figures, we see that Oracle is indeed the upper bound on the performance of any pricing method due to the availability of perfect information that this method assumes.

In the case of revenue optimization (Figure 5.8a), the VI baseline, which provides realistic in-expectation optimal pricing solutions, can solve problem instances with up to 6 timeslots. Above 6, we ran out of memory for representing the state space. The MCTS solution performs well compared to the VI and manages to outperform the flat rate by a wide margin, shadowing the performance of the Oracle baseline. However, the flat rate generates higher utilization when optimizing for revenue than MCTS in this medium demand scenario.

Notably, in Figure 5.8a, the revenue of the Oracle baseline and MCTS increases with increasing timeslots. This is caused by keeping the total duration of requested charging timeslots the same in all problem instances. Increasing the number of timeslots increases the number of individual charging requests. The finer timeslot discretization allows the total expected charging duration to be split into more differently sized charging sessions. The Oracle and MCTS can leverage this larger number of requests to prefer requests from EV users with higher budgets. Since the flat rate can not respond this way, the revenue from the flat rate remains level.

When optimizing utilization (Figure 5.9a), unlike in the previous case, the dynamic pricing methods can not exploit the high budgets of some EV users to boost performance. However, while both flat-rate and Oracle give, for the most part, steady performance, the utilization of MCTS slowly falls as the number of requests and timeslots increases. However, for instances where the optimal VI baseline provides results, the MCTS closely follows it's performance. Here, the cause is again a larger number of more fine-grained requests. The smaller granularity means there is a greater chance of overlap among the requests. While Oracle can readily prevent overlap due to its nature, this is not the case for MCTS. The task becomes increasingly difficult as the number of requests (and timeslots) increases. Unlike in the case of revenue maximization, where MCTS dynamic pricing exploits differences in budgets to boost the optimization metric, the utilization optimization does not allow for such exploitation as every request improves the optimization metric only by its duration.

The higher number and shorter duration of charging requests (while the total duration of requested charging timeslots is kept constant) provide an opportunity for dynamic pricing to increase revenue through allocating resources to higher-paying EV drivers, as seen in Figure 5.8. We show that MCTS dynamic pricing can leverage this, closely shadowing the optimal VI where the comparison is available. Coincidentally, introducing long-distance, very fast EV charging stations such as the Tesla superchargers means a higher number and shorter charging sessions in high-demand locations such as highways. Such locations could be a prime candidate for the dynamic pricing scheme discussed in this chapter.

## 5.5 Summary

In this chapter, we make number of impromevements to our initial dynamic pricing model presented in Chapter 4 that make it much more viable. The proposed model of dynamic pricing together with the MCTS solver maximize either (1) the revenue of the CS operator or (2) the overall utilization of the available capacity.

We have presented an MDP model that considers demand comprehensively, throughout the day, removing the splitting of MDPs done in Chapter 4. Additionally, we have generalized the demand model as a collection of discretized Poisson processes. Not only being more realistic on its own, this model allows us to easily incorporate demand varying in time. We describe these improvements using domain-independent language to give clear description of our model.

Since the model uses discretization of the Poisson processes and relies on their properties to generalize to time-dependent demand, we formally analyze the relationship between the continous and discretized demand processes. In the process, we develop a notion of error relevant to our problem and determine appropriate parameters for the discretization.

Since the new model is more complex, we propose a heuristic solution method based on MCTS that uses the properties of the model. This method scales much better than the optimal solution through value iteration with only small loss of optimality. The successful application of MCTS to the dynamic pricing of EV charging problem is a novel contribution of this chapter. The performance of this solution is then demostrated on the EV charging dataset.

To validate the performance of the method, we carried out the experiments using a real-world dataset from a German CS operator and compared flat-rate, optimal pricing throug value itration, and Oracle baselines with our proposed MCTS-based pricing. The results of our second experiment show that the proposed MCTS method achieved over 90% of revenue of the optimal pricing policy provided by the VI, and it did so without significantly increasing the variance of the results. Additionally, we have shown that the heuristic MCTS solution scaled to up to ten orders of magnitude larger problem instances than VI (in terms of the state space size).

CHAPTER 6

# Conclusion

In this thesis, we had three objectives: 1. to find an appropriate approach to reservations of EV charging and thus determine the product for our dynamic pricing solutions, 2. to develop a MDP model of dynamic pricing around the chosen reservation product that could be solved optimally, and 3. to develop a practical MDP model that can handle realistically sized instances.

**Objective 1:** We achieved the first objective in Chapter 3 through work published in workshop publication [Bas+19], where we conducted multi-agent simulation study comparing different options for an EV charging reservation system. The results of the simulations show that the limited variant of reservations that is feasible with existing systems is in fact not very useful to any of the stakeholders and we should concentrate our efforts on full-fledged reservations for EV charging.

The main contribution of the chapter is a thorough analysis of different reservation types and their benefits for electromobility stakeholders, a topic that has not received significant attention in the literature before. The results also identify the fully-reserved EV charging sessions as the basis for the dynamic pricing models developed as part of Objective 2 and 3.

**Objective 2:** The second objective was achieved in Chapter 4 through two conference publications [MKJ18a; MKJ18b]. With the developed concept of product of EV charging, we focused on the problem of dynamically pricing these products so that the EV charging capacity is allocated optimally. The developed Split-MDP model is compact (each MDP segment having at most $10^4$ states) and can be solved optimally with basic MDP methods. In simulations it shows promising results when compared to the baseline pricing strategies. This is the main contribution of the chapter.

In the results, the Split-MDP pricing works well across all price-demand functions and various demand levels. The higher the demand and price elasticity of demand, the better the revenue. When maximizing revenue, the Split-MDP model also increases the utilization of the charging station by more efficient allocation of the capacity.

However, this Split-MDP model is limited in considering the price in different timeslots as independent, which is not realistic. Additionally, the cyclical state space structure in the MDP transition function results in a simplified user demand model that cannot capture real-world demand processes.

**Objective 3:** Therefore, in Chapter 5, we developed a more general Full-MDP model of dynamic pricing that addressed the third objective of this thesis. This approach will be the content of a journal paper titled "Dynamic pricing of EV charging reservations with MDPs" that we are preparing for publication. Some preliminary results are also published in [MB22]. The model is built around the Poisson demand process, a popular and theoretically sound choice for modeling user arrivals in the literature. However, since we use finite MDP model, we have to discretize the time. To do this in a sound way, we derive two types of error based on the theoretical properties of the discretization that we can use to estimate the error of the model. Additionally, instead of considering the pricing of different intervals independently as is done by the Split-MDP model, this Full-MDP model considers the temporal relationships between existing and future bookings in pricing as well. This Full-MDP is the first contribution of this chapter.

As a result of these changes, the state space of the Full-MDP model is much larger than that of the Split-MDP model. This makes the model intractable for exact solution methods in realistically large instances. Therefore, we adapt a heuristic solution method based on MCTS to solve the Full-MDP model. This heuristic solution and its application to the Full-MDP model is the second contribution of this chapter. In our experiments, we show that this heuristic solver is able to find near-optimal solutions in small instances (state-space in order of $10^7$) where optimal solutions can be found. In practical instances, with state-space sizes in order of $10^{28}$ and thousands of timesteps, we compare the performance to several baselines, with results suggesting that the heuristic solver generalizes well to larger instances. Additionally, the solver can provide pricing actions in a reasonable time frame (in order of milliseconds) with consumer-grade hardware.

The developed model is general enough to be used in other domains and relies on small number of assumptions. Furthermore, it does not rely on specific properties of today's EV charging infrastructure and can easily incorporate signals from different electromobility stakeholders.

**Overall,** the main contribution of the thesis is the comprehensive handling of the problem of dynamic pricing of EV charging in three steps: 1. identification of an appropriate product, the fully-reserved EV charging paid as a service instead of hourly rate, 2. validation of the MDPs as a solution methodology to the pricing constructed around the said product, and 3. development of a practical MDP model and heuristic solution method that can be applied in practice. This approach to dynamic pricing of EV charging has not been a major focus of the literature before and the thesis is the first to address it in such a comprehensive way.

Together with the multicriterial route and charging planner [CJM24], the comprehensive approach to dynamic pricing of EV charging with full reservations brings several advantages to both EV drivers and charging station operators over other approaches common in the field. For the drivers using the planner, the driver has multiple cost and travel time options to choose from and the charging is predictable and reliable. For the charging station operator, the ability to plan charging in advance and dynamically price the new incoming requests allows for better utilization of the charging capacity and increased revenue which can be used to support further growth of the charging infrastructure.

## 6.1 Future Work

Over the course of writing this thesis, several possible directions for future work have become apparent.

Starting with the simulation study in Chapter 3 that we used to compare different reservation types, we could extend the study as an evaluation of the dynamic-pricing models developed in this thesis. This would involve the Full-MDP dynamic pricing model applied to the CS side of the simulation and the multicriterial route and charging planner [CJM24] applied to the driver side. Here, we could evaluate the performance of the whole system compared to the flat-rate baseline.

Regarding the MDP models described in Chapters 4 and 5, there are a number of ways they could be improved. First, some of the simplifications made in the Split-MDP model that significantly reduce its complexity could be applied to the Full-MDP model as well. Specifically, the self-loop in the transition function in Figure 4.2 that significantly reduces the state space and the action space could potentially be applied to the Full-MDP model as well with careful choice of discretization parameters. While the Poisson distribution of demand in the Full-MDP model can be only poorly approximated by the geometric distribution induced by the self-loop, the resulting model could still be useful in practice as the maximization is done in expectation.

The second change to the model would involve replacing the finite horizon MDP model with an infinite horizon. Superficially, this is a cosmetic change, as the complexity of the problem would not be directly affected; the states would still have to encode the cyclical daily patterns involved in the problem. However, the infinite horizon model would allow for a more natural representation of the problem that could enable investigation of the fixed points of the model.

Finally, the MCTS solution method leaves plenty of space for improvement. For example, incorporating other domain-specific heuristics into the rollout policy would likely improve solver performance. Moreover, we have not significantly investigated the stability of the solver, which could theoretically pose regulatory as well as practical problems in practice. The solver could also be significantly sped up with parallelization, which we have not yet seriously attempted.

We also have not yet put a significant effort into considering other MDP solvers and approaches, such as a large body of literature in reinforcement learning. While there are good reasons to avoid it, it would be an interesting comparison to make.

Finally, an area of research we have only lightly touched upon is the topic of fitting the various distributions and random processes involved in the model to real-world data. There are statistical methods such as [Gug+20] that we could use and such evaluation on a testing dataset would be a good way to further validate the model and provide a more realistic evaluation of the performance of the solver. At the same time, such extended evaluation would provide a guideline for the practical application of the model in the real world.

APPENDIX A

# Publications

The list below is a summary of the publications of the author. Publications in individual categories are sorted by citation counts (excluding self-citations) from Web of Science (WoS).

The contributions use the CRediT author statement [Els24] based on [Bra+15]. The author statement categories are reproduced in Table A.1.

| | | | |
|---|---|---|---|
| 1 | Conceptualization | 8 | Data curation |
| 2 | Methodology | 9 | Writing – original draft |
| 3 | Software | 10 | Writing – review and editing |
| 4 | Validation | 11 | Visualization |
| 5 | Formal analysis | 12 | Supervision |
| 6 | Investigation | 13 | Project administration |
| 7 | Resources | 14 | Funding acquisition |

**Table A.1:** CRediT contribution cathegories [Bra+15].

## A.1 Publications related to the thesis

### Journal publications (with IF)

[CJM24] Marek Cuchý, Michal Jakob, and Jan Mrkos. "Route and Charging Planning for Electric Vehicles: A Multi-Objective Approach". In: *Transportation Letters — Accepted and waiting for publication* (2024)
(Contributions: `JM: 10, MC: 1-3,5,6,8,9,11, MJ: 1,2,10,12`)

[MB22] Jan Mrkos and Robert Basmadjian. "Dynamic Pricing for Charging of EVs with Monte Carlo Tree Search". In: *Smart Cities* 5.1 (1 Mar. 2022)
(Contributions: `JM: 1-3,5,6,8-11, RB: 1,4,10,14`)

### Other WoS publications

[MKJ18b] Jan Mrkos, Antonín Komenda, and Michal Jakob. "Revenue Maximization for Electric Vehicle Charging Service Providers Using Sequential Dynamic Pricing".

In: *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018).* Stockholm, Sweden: IFAAMAS, July 10–15, 2018
(**6** citations, Contributions: `JM: 1-3,5,6,8-11, AK: 1,2,9,10, MJ: 1,10,12,14`)

[Bas+19] Robert Basmadjian, Benedikt Kirpes, Jan Mrkos, Marek Cuchý, and Sahar Rastegar. "An Interoperable Reservation System for Public Electric Vehicle Charging Stations: A Case Study in Germany". In: *Proceedings of the 1st ACM International Workshop on Technology Enablers and Innovative Applications for Smart Cities and Communities.* TESCA'19. New York, NY, USA: Association for Computing Machinery, Nov. 13, 2019
(**5** citations, Contributions: `JM: 2,3,5,6,8-11, MC: 2,3,6,8, RB: 1,2,9,10,11,14, BK: 1,2,9,10,11`)

[MKJ18a] Jan Mrkos, Antonín Komenda, and Michal Jakob. "Dynamic Pricing Strategy for Electromobility Using Markov Decision Processes". In: *10th International Conference on Agents and Artificial Intelligence (ICAART).* Apr. 6, 2018
(Contributions: `JM: 1-3,5,6,8-11, AK: 1,2,9,10, MJ: 1,10,12,14`)

## A.2 Other publications

**Journal publications (with IF)**

[Bas+20] Robert Basmadjian, Benedikt Kirpes, Jan Mrkos, and Marek Cuchý. "A Reference Architecture for Interoperable Reservation Systems in Electric Vehicle Charging". In: *Smart Cities* 3.4 (4 Dec. 2020)
(**6** citations, Contributions: `JM: 2,3,5,6,8-11, MC: 2,3,6,10, RB: 1,2,9-11,14, BK: 1,2,9-11, SR: 2,9,10`)

[Ega+18] Malcolm Egan, Jan Drchal, Jan Mrkos, and Michal Jakob. "Towards Data-Driven On-Demand Transport". In: *EAI Endorsed Transactions on Industrial Networks and Intelligent Systems* "5".14 (June 27, 2018)
(Contributions: `JM: 3,5,6,8-11, MJ: 9,10,12, ME: 1,2,9-11, JD: 2,9,10,12`)

**Other WoS publications**

[Sch+19] Martin Schaefer, Michal Čáp, Jan Mrkos, and Jiří Vokřínek. "Routing a Fleet of Automated Vehicles in a Capacitated Transportation Network". In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).* Nov. 2019
(**2** citations, Contributions: `JM: 2,10,11, MS: 1,2,3,5,6,8,9,10,11, MČ: 1,2,9,10, JV: 1,10,12`)

[FM23] David Fiedler and Jan Mrkos. "Large-Scale Ridesharing DARP Instances Based on Real Travel Demand". In: *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC).* Bilbao, Spain: IEEE, Sept. 24, 2023
(Contributions: `JM: 2-6,10,11, DF: 1-6,9,10,12,13`)

**Patents**

[MGK16] Jan Mrkos, Martin Grill, and Jan Kohout. “Tracking Users over Network Hosts Based on User Behavior”. U.S. pat. 20160352760A1. Cisco Technology Inc. Dec. 1, 2016

# Bibliography


[12] *Smart Grid Reference Architecture*. CEN-CENELEC-ETSI Smart Grid Co-ordination Group, Nov. 2012, p. 231. URL: https://www.cencenelec.eu/media/CEN-CENELEC/AreasOfWork/CEN-CENELEC_Topics/Smart%20Grids%20and%20Meters/Smart%20Grids/reference_architecture_smartgrids.pdf.

[17] *D2.2 – Initial Description of Scenarios, Business Requirements and Use Cases*. Project deliverable D2.2. UNIMA, May 31, 2017. URL: https://electrific-project.eu/wp-content/uploads/2020/05/D2.2_Initial-Description-of-Scenarios-Business-Requirements-and-Use-Cases.pdf (visited on 10/16/2023).

[18] *Global EV Outlook 2018*. Paris: OECD/IEA, 2018. URL: https://www.iea.org/reports/global-ev-outlook-2018 (visited on 02/02/2024).

[20a] *Continous Pricing*. Lufthansa Group, Oct. 2020. URL: https://www.lufthansaexperts.com/shared/files/lufthansa/public/mcms/folder_102/folder_3212/folder_4994/file_148593.pdf.

[20b] *D9.6 - Final Exploitation Framework: Market Analysis, Project Impact and Sustainability Plan*. Project deliverable D9.6. BCNEco, May 3, 2020. URL: https://electrific-project.eu/wp-content/uploads/2020/05/D9.6_Final-exploitation-framework_Market-analysis-project-impact-and-sustainability-plan.pdf (visited on 10/16/2023).

[22a] *Electric Vehicle Outlook 2022*. BloombergNEF, 2022. URL: https://bnef.turtl.co/story/evo-2022 (visited on 10/16/2023).

[22b] *Enabling Seamless Electromobility through Smart Vehicle-Grid Integration | ELECTRIFIC Project | Fact Sheet | H2020*. CORDIS | European Commission. Aug. 2022. URL: https://cordis.europa.eu/project/id/713864 (visited on 02/18/2024).

[22c] *Here's How to Accelerate the Electric Vehicle Revolution*. World Economic Forum. Jan. 31, 2022. URL: https://www.weforum.org/agenda/2022/01/the-ev-revolution-obstacles-solutions/ (visited on 02/02/2024).

[23a] *Global EV Outlook 2023*. Paris: IEA, 2023, p. 142. URL: https://www.iea.org/reports/global-ev-outlook-2023 (visited on 10/16/2023).

[23b] *Interconnected Network of Continental Europe 2023*. European Network of Transmission System Operators for Electricity, Sept. 13, 2023. URL: https://www.entsoe.eu/Documents/Publications/maps/2023/230922/Map_Continental-Europe-2.500.000.pdf (visited on 02/20/2024).

[23c] *Připojovací podmínky pro výrobny elektřiny*. 23_28-0023r00. ČEZ Energetické služby, s. r. o., Jan. 1, 2023. URL: https://www.cez.cz/webpublic/file/edee/2023/01/23_28-0023r00_pripojovaci_podminky_vyrobny_komplet_final.pdf (visited on 02/24/2024).

[23d] *Solar Industry Research Data | SEIA*. 2023. URL: https://www.seia.org/solar-industry-research-data (visited on 02/04/2024).

[ACF02] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. "Finite-Time Analysis of the Multiarmed Bandit Problem". In: *Machine Learning* 47.2 (May 1, 2002), pp. 235–256. ISSN: 1573-0565. DOI: 10.1023/A:1013689704352. URL: https://doi.org/10.1023/A:1013689704352 (visited on 12/05/2022).

[AE08] M. H. Albadi and E. F. El-Saadany. "A Summary of Demand Response in Electricity Markets". In: *Electric Power Systems Research* 78.11 (Nov. 1, 2008), pp. 1989–1996. ISSN: 0378-7796. DOI: 10.1016/j.epsr.2008.04.002. URL: http://www.sciencedirect.com/science/article/pii/S0378779608001272.

[AKR14] O. Ardakanian, S. Keshav, and C. Rosenberg. "Real-Time Distributed Control for Smart Electric Vehicle Chargers: From a Static to a Dynamic Study". In: *IEEE Transactions on Smart Grid* 5.5 (Sept. 2014), pp. 2295–2305. ISSN: 1949-3053. DOI: 10.1109/TSG.2014.2327203.

[ARK12] O. Ardakanian, C. Rosenberg, and S. Keshav. "RealTime Distributed Congestion Control for Electrical Vehicle Charging". In: *SIGMETRICS Perform. Eval. Rev.* 40.3 (Jan. 2012), pp. 38–42. ISSN: 0163-5999. DOI: 10.1145/2425248.2425257. URL: http://doi.acm.org/10.1145/2425248.2425257 (visited on 10/18/2016).

[Bas+19] Robert Basmadjian, Benedikt Kirpes, Jan Mrkos, Marek Cuchý, and Sahar Rastegar. "An Interoperable Reservation System for Public Electric Vehicle Charging Stations: A Case Study in Germany". In: *Proceedings of the 1st ACM International Workshop on Technology Enablers and Innovative Applications for Smart Cities and Communities*. TESCA'19. New York, NY, USA: Association for Computing Machinery, Nov. 13, 2019, pp. 22–29. ISBN: 978-1-4503-7015-8. DOI: 10.1145/3364544.3364825. URL: https://doi.org/10.1145/3364544.3364825 (visited on 02/10/2020).

[Bas+20] Robert Basmadjian, Benedikt Kirpes, Jan Mrkos, and Marek Cuchý. "A Reference Architecture for Interoperable Reservation Systems in Electric Vehicle Charging". In: *Smart Cities* 3.4 (4 Dec. 2020), pp. 1405–1428. DOI: 10.3390/smartcities3040067. URL: https://www.mdpi.com/2624-6511/3/4/67 (visited on 11/22/2020).

[Bas20] Robert Basmadjian. "Optimized Charging of PV-Batteries for Households Using Real-Time Pricing Scheme: A Model and Heuristics-Based Implementation". In: *Electronics* 9.1 (1 Jan. 2020), p. 113. ISSN: 2079-9292. DOI: 10.3390/electronics9010113. URL: https://www.mdpi.com/2079-9292/9/1/113 (visited on 02/17/2024).

[Bau+13] Moritz Baum, Julian Dibbelt, Thomas Pajor, and Dorothea Wagner. "Energy-Optimal Routes for Electric Vehicles". In: *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2013, pp. 54–63.

[Bay+13] I. S. Bayram, G. Michailidis, M. Devetsikiotis, and F. Granelli. "Electric Power Allocation in a Network of Fast Charging Stations". In: *IEEE Journal on Selected Areas in Communications* 31.7 (July 2013), pp. 1235–1246. ISSN: 0733-8716. DOI: 10.1109/JSAC.2013.130707.

[BD18] Robert Basmadjian and Hermann De Meer. "A Heuristics-Based Policy to Reduce the Curtailment of Solar-Power Generation Empowered by Energy-Storage Systems". In: *Electronicsweek* 7.349 (12 2018). ISSN: 2079-9292. DOI: 10.3390/electronics7120349. URL: https://www.mdpi.com/2079-9292/7/12/349.

[Bel57] Richard Bellman. "A Markovian Decision Process". In: *Indiana University Mathematics Journal* 6.4 (1957), pp. 679–684. ISSN: 0022-2518. DOI: 10.1512/iumj.1957.6.56038. URL: http://www.iumj.indiana.edu/IUMJ/fulltext.php?artid=56038&year=1957&volume=6 (visited on 08/29/2018).

[BG12] Sandford Bessler and Jesper Grønbæk. "Routing EV Users towards an Optimal Charging Plan". In: *International Battery, Hybrid and Fuel Cell Electric Vehicle Symposium*. Vol. 56. 2012. URL: http://www.ea.tuwien.ac.at/fileadmin/t/ea/projekte/KOFLA/EVS26.pdf (visited on 11/04/2016).

[Bha+16] Saptarshi Bhattacharya, Koushik Kar, Joe H. Chow, and Aparna Gupta. "Extended Second Price Auctions with Elastic Supply for PEV Charging in the Smart Grid". In: *IEEE Transactions on Smart Grid* 7.4 (2016), pp. 2082–2093. URL: https://ieeexplore.ieee.org/abstract/document/7440866/ (visited on 02/17/2024).

[BMD12] Daehyun Ban, George Michailidis, and Michael Devetsikiotis. "Demand Response Control for PHEV Charging Stations by Dynamic Price Adjustments". In: *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*. IEEE, 2012, pp. 1–8. URL: https://ieeexplore.ieee.org/abstract/document/6175601/ (visited on 02/17/2024).

[Bra+15] Amy Brand, Liz Allen, Micah Altman, Marjorie Hlava, and Jo Scott. "Beyond Authorship: Attribution, Contribution, Collaboration, and Credit". In: *Learned Publishing* 28.2 (2015), pp. 151–155. ISSN: 1741-4857. DOI:

10.1087/20150211. URL: https://onlinelibrary.wiley.com/doi/abs/10.1087/20150211 (visited on 02/28/2024).

[Bro+12] Cameron B. Browne et al. “A Survey of Monte Carlo Tree Search Methods”. In: *IEEE Transactions on Computational Intelligence and AI in Games* 4.1 (Mar. 2012), pp. 1–43. ISSN: 1943-0698. DOI: 10.1109/TCIAIG.2012.2186810.

[CCX06] Wen-Chyuan Chiang, Jason C.H. Chen, and Xiaojing Xu. “An Overview of Research on Revenue Management: Current Issues and Future Research”. In: *International Journal of Revenue Management* 1.1 (Oct. 30, 2006), pp. 97–128. ISSN: 1474-7332. DOI: 10.1504/IJRM.2007.011196. URL: http://www.inderscienceonline.com/doi/abs/10.1504/IJRM.2007.011196.

[Cha+10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. “Multi-Parameter Mechanism Design and Sequential Posted Pricing”. In: *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*. STOC ’10. New York, NY, USA: Association for Computing Machinery, June 5, 2010, pp. 311–320. ISBN: 978-1-4503-0050-6. DOI: 10.1145/1806689.1806733. URL: https://dl.acm.org/doi/10.1145/1806689.1806733 (visited on 02/17/2024).

[Cha23] Robert N. Charette. *The EV Transition Explained*. IEEE Spectrum, 2023. 49 pp. URL: https://spectrum.ieee.org/the-ev-transition-explained-2659602311 (visited on 02/02/2024).

[Che+16] Lixing Chen, Zhong Chen, Xueliang Huang, and Long Jin. “A Study on Price-Based Charging Strategy for Electric Vehicles on Expressways”. In: *Energies* 9.5 (2016), p. 385. URL: http://www.mdpi.com/1996-1073/9/5/385/htm (visited on 10/06/2017).

[CJM24] Marek Cuchý, Michal Jakob, and Jan Mrkos. “Route and Charging Planning for Electric Vehicles: A Multi-Objective Approach”. In: *Transportation Letters — Accepted and waiting for publication* (2024). ISSN: 1942-7867. DOI: 10.1080/19427867.2024.2315359. URL: https://doi.org/10.1080/19427867.2024.2315359.

[Cla+20] Christopher T. M. Clack, Aditya Choukulkar, Brianna Coté, and Sarah A. McKee. *Why Local Solar For All Costs Less: A New Roadmap for the Lowest Cost Grid*. Technical report. Boulder, Colorado: Vibrant Clean Energy, LLC, Dec. 2020, p. 97. URL: https://www.esig.energy/coordinated-deployments-of-transmission-and-distribution-scale-resources-provide-the-lowest-cost-electricity/.

[Cou06] Rémi Coulom. “Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search”. In: *International Conference on Computers and Games*. Springer, 2006, pp. 72–83.

[Cro08] David Crossley. *Assessment and Development of Network-driven Demand-side Management Measures*. Research Report Second Revised Edition. Australia: International Energy Agency, Oct. 10, 2008, p. 72. URL: https://userstcp.org/task/task-15-network-driven-dsm/ (visited on 11/28/2023).

[CŠJ18] M. Cuchý, M. Štolba, and M. Jakob. “Benefits of Multi-Destination Travel Planning for Electric Vehicles”. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. 2018 21st International Conference on Intelligent Transportation Systems (ITSC). Nov. 2018, pp. 327–332. DOI: 10.1109/ITSC.2018.8569385.

[Cui+23] Li Cui et al. “Dynamic Pricing for Fast Charging Stations with Deep Reinforcement Learning”. In: *Applied Energy* 346 (Sept. 15, 2023), p. 121334. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2023.121334. URL: https://www.sciencedirect.com/science/article/pii/S0306261923006980 (visited on 06/15/2023).

[De +13] Mathijs Michiel De Weerdt, Enrico H. Gerding, Sebastian Stein, Valentin Robu, and Nicholas R. Jennings. “Intention-Aware Routing to Minimise Delays at Electric Vehicle Charging Stations”. In: (2013). URL: http://repository.tudelft.nl/view/ir/uuid:492b1dbd-e735-4ac8-a025-bae4a3a42d5e/ (visited on 10/18/2016).

[Dei+11] Sara Deilami, Amir S. Masoum, Paul S. Moses, and Mohammad A. S. Masoum. “Real-Time Coordination of Plug-In Electric Vehicle Charging in Smart Grids to Minimize Power Losses and Improve Voltage Profile”. In: *IEEE Transactions on Smart Grid* 2.3 (Sept. 2011), pp. 456–467. ISSN: 1949-3061. DOI: 10.1109/TSG.2011.2159816. URL: https://ieeexplore.ieee.org/abstract/document/5986769 (visited on 02/17/2024).

[DWH19] Bingqian Du, Chuan Wu, and Zhiyi Huang. “Learning Resource Allocation and Pricing for Cloud Profit Maximization”. In: The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19). Honolulu, Hawai, USA, Jan. 27, 2019.

[Ega+18] Malcolm Egan, Jan Drchal, Jan Mrkos, and Michal Jakob. “Towards Data-Driven On-Demand Transport”. In: *EAI Endorsed Transactions on Industrial Networks and Intelligent Systems* "5".14 (June 27, 2018). ISSN: 2410-0218. URL: http://eudl.eu/doi/10.4108/eai.27-6-2018.154835 (visited on 09/05/2018).

[Ego+17] Maxim Egorov et al. “POMDPs.Jl: A Framework for Sequential Decision Making under Uncertainty”. In: *Journal of Machine Learning Research* 18.26 (2017), pp. 1–5. URL: http://jmlr.org/papers/v18/16-300.html.

[Eid+17] Markus Eider et al. “Seamless Electromobility”. In: *Proceedings of the Eighth International Conference on Future Energy Systems*. New York,NY, USA: ACM, 2017, pp. 316–321.

[Els24] Elsevier. *CRediT Author Statement*. www.elsevier.com. 2024. URL: https://www.elsevier.com/researcher/author/policies-and-guidelines/credit-author-statement (visited on 02/28/2024).

[Ete+17] S. Rasoul Etesami, Walid Saad, Narayan Mandayam, and H. Vincent Poor. "Smart Routing in Smart Grids". In: *Decision and Control (CDC), 2017 IEEE 56th Annual Conference On*. IEEE, 2017, pp. 2599–2604.

[Fai09] Peter Fairley. "Speed Bumps Ahead for Electric-Vehicle Charging". In: *IEEE Spectrum* (Dec. 31, 2009). URL: https://spectrum.ieee.org/speed-bumps-ahead-for-electricvehicle-charging (visited on 02/28/2024).

[Fan+21] Cheng Fang, Haibing Lu, Yuan Hong, Shan Liu, and Jasmine Chang. "Dynamic Pricing for Electric Vehicle Extreme Fast Charging". In: *IEEE Transactions on Intelligent Transportation Systems* 22.1 (Jan. 2021), pp. 531–541. ISSN: 1558-0016. DOI: 10.1109/TITS.2020.2983385. URL: https://ieeexplore.ieee.org/abstract/document/9057557?casa_token=XJ6JCoJBM5AAAAAA%3AqAUAoOHbZMok5ju8pHviT4XW2o94D5NA0AIZL4V3c7xRAiV7mtfWwhWv (visited on 02/10/2024).

[FM23] David Fiedler and Jan Mrkos. "Large-Scale Ridesharing DARP Instances Based on Real Travel Demand". In: *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*. 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC). Bilbao, Spain: IEEE, Sept. 24, 2023, pp. 2750–2757. ISBN: 9798350399462. DOI: 10.1109/ITSC57777.2023.10422146. URL: https://ieeexplore.ieee.org/document/10422146/ (visited on 02/26/2024).

[Fra+11] Inês Frade, Anabela Ribeiro, Gonçalo Gonçalves, and António Pais Antunes. "Optimal Location of Charging Stations for Electric Vehicles in a Neighborhood in Lisbon, Portugal". In: *Transportation Research Record* 2252.1 (Jan. 1, 2011), pp. 91–98. ISSN: 0361-1981. DOI: 10.3141/2252-12. URL: https://doi.org/10.3141/2252-12 (visited on 02/17/2024).

[GA16] Arnob Ghosh and Vaneet Aggarwal. "Menu-Based Pricing for Charging of Electric Vehicles with Vehicle-to-Grid Service". Nov. 30, 2016. arXiv: 1612.00106 [cs, math]. URL: http://arxiv.org/abs/1612.00106 (visited on 10/06/2017).

[Gal13] Robert G. Gallager. *Stochastic Processes: Theory for Applications*. Cambridge University Press, 2013.

[Gal15] Amy Gallo. *A Refresher on Price Elasticity*. Harvard Business Review. Aug. 21, 2015. URL: https://hbr.org/2015/08/a-refresher-on-price-elasticity (visited on 08/28/2017).

[GAM15] Jiarui Gan, Bo An, and Chunyan Miao. "Optimizing Efficiency of Taxi Systems: Scaling-up and Handling Arbitrary Constraints." In: *AAMAS*. 2015, pp. 523–531. URL: https://www.ifaamas.org/Proceedings/aamas2015/aamas/p523.pdf (visited on 02/17/2024).

[Gan+13] Jiarui Gan, Bo An, Haizhong Wang, Xiaoming Sun, and Zhongzhi Shi. “Optimal Pricing for Improving Efficiency of Taxi Systems”. In: *Twenty-Third International Joint Conference on Artificial Intelligence*. 2013. URL: https://www.ijcai.org/Proceedings/13/Papers/414.pdf (visited on 02/17/2024).

[Ger+13] Enrico H. Gerding, Sebastian Stein, Valentin Robu, Dengji Zhao, and Nicholas R. Jennings. “Two-Sided Online Markets for Electric Vehicle Charging”. In: *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 989–996. URL: http://dl.acm.org/citation.cfm?id=2485076 (visited on 12/21/2016).

[Ger+91] Anne Geraci et al. *IEEE Standard Computer Dictionary: Compilation of IEEE Standard Computer Glossaries*. IEEE Press, 1991. ISBN: 1-55937-079-3.

[GH09] Xianping Guo and Onésimo Hernández-Lerma. “Continuous-Time Markov Decision Processes”. In: *Continuous-Time Markov Decision Processes: Theory and Applications*. Ed. by Xianping Guo and Onésimo Hernández-Lerma. Stochastic Modelling and Applied Probability. Berlin, Heidelberg: Springer, 2009, pp. 9–18. ISBN: 978-3-642-02547-1. DOI: 10.1007/978-3-642-02547-1_2. URL: https://doi.org/10.1007/978-3-642-02547-1_2 (visited on 03/09/2023).

[Gu97] Zheng Gu. “Proposing a Room Pricing Model for Optimizing Profitability.” In: *International Journal of Hospitality Management* 16.3 (1997), pp. 273–277. URL: https://www.cabdirect.org/cabdirect/abstract/19981803115 (visited on 02/17/2024).

[Gug+20] Shota Gugushvili, Frank van der Meulen, Moritz Schauer, and Peter Spreij. *Fast and Scalable Non-Parametric Bayesian Inference for Poisson Point Processes*. Mar. 29, 2020. DOI: 10.48550/arXiv.1804.03616. arXiv: 1804.03616 [stat]. URL: http://arxiv.org/abs/1804.03616 (visited on 02/16/2023). preprint.

[Gur23] LLC Gurobi Optimization. *Gurobi Optimizer Reference Manual*. 2023. URL: https://www.gurobi.com/ (visited on 10/07/2020).

[Ham+20] Jamie Hamilton et al. *Electric Vehicles: Setting a Course for 2030*. Deloitte, 2020.

[Hay+15] Keiichiro Hayakawa, Enrico Gerding, Sebastian Stein, and Takahiro Shiga. “Online Mechanisms for Charging Electric Vehicles in Settings with Varying Marginal Electricity Costs”. In: (2015). URL: https://eprints.soton.ac.uk/377236/ (visited on 02/17/2024).

[He+13] Fang He, Di Wu, Yafeng Yin, and Yongpei Guan. "Optimal Deployment of Public Charging Stations for Plug-in Hybrid Electric Vehicles". In: *Transportation Research Part B: Methodological* 47 (Jan. 1, 2013), pp. 87–101. ISSN: 0191-2615. DOI: 10.1016/j.trb.2012.09.007. URL: https://www.sciencedirect.com/science/article/pii/S0191261512001336 (visited on 02/17/2024).

[How60] Ronald A. Howard. *Dynamic Programming and Markov Processes*. Dynamic Programming and Markov Processes. Oxford, England: John Wiley, 1960, pp. viii, 136. viii, 136.

[Jak+12] Michal Jakob et al. "AgentPolis: Towards a Platform for Fully Agent-based Modeling of Multi-modal Transportation (Demonstration)". In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*. AAMAS '12. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2012, pp. 1501–1502. ISBN: 978-0-9817381-3-0. URL: http://dl.acm.org/citation.cfm?id=2343896.2344081 (visited on 09/17/2018).

[JSG14] C. Jin, X. Sheng, and P. Ghosh. "Optimized Electric Vehicle Charging With Intermittent Renewable Energy Sources". In: *IEEE Journal of Selected Topics in Signal Processing* 8.6 (Dec. 2014), pp. 1063–1072. ISSN: 1932-4553. DOI: 10.1109/JSTSP.2014.2336624.

[Kam+22] Philipp Kampshoff, Adi Kumar, Shannon Peloquin, and Shivika Sahdev. *Building the Electric-Vehicle Charging Infrastructure America Needs*. McKinsey & Company, Apr. 2022, p. 10. URL: https://www.mckinsey.com/industries/public-sector/our-insights/building-the-electric-vehicle-charging-infrastructure-america-needs#/ (visited on 02/20/2024).

[Kha+18] Saad Ullah Khan et al. "Energy Management Scheme for an EV Smart Charger V2G/G2V Application with an EV Power Allocation Technique and Voltage Regulation". In: *Applied Sciences* 8.648 (4 2018). ISSN: 2076-3417. DOI: 10.3390/app8040648. URL: https://www.mdpi.com/2076-3417/8/4/648.

[Kim+10] Hye Jin Kim, Junghoon Lee, Gyung Leen Park, Min Jae Kang, and Mikyung Kang. "An Efficient Scheduling Scheme on Charging Stations for Smart Transportation". In: *International Conference on Security-Enriched Urban Computing and Smart Grid*. Hualien, Taiwan: Springer, 2010, pp. 274–278. ISBN: 3-642-16443-9. DOI: 10.1007/978-3-642-16444-6{_}35.

[Kir+19] Benedikt Kirpes, Philipp Danner, Robert Basmadjian, Hermann De Meer, and Christian Becker. "E-Mobility Systems Architecture: A Framework for Managing Complexity and Interoperability". In: *Submitted to Energy Informatics (Preprint)* 0.0 (2019), pp. 1–30.

[KJ17] Shripad Kulkarni and Pushkar H. Joshi. “Passenger Airline Revenue Management: Research Overview and Emerging Literature”. In: *International Journal of Engineering and Management Research (IJEMR)* 7.1 (2017), pp. 387–389.

[KKC16] Yeongjin Kim, Jeongho Kwak, and Song Chong. “Dynamic Pricing, Scheduling, and Energy Management for Profit Maximization in PHEV Charging Stations”. In: *IEEE Transactions on Vehicular Technology* 66.2 (2016), pp. 1011–1026. URL: https://ieeexplore.ieee.org/abstract/document/7468568/ (visited on 02/17/2024).

[KS06] Levente Kocsis and Csaba Szepesvári. “Bandit Based Monte-Carlo Planning”. In: *Machine Learning: ECML 2006*. Ed. by Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou. Red. by David Hutchison et al. Vol. 4212. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 282–293. ISBN: 978-3-540-45375-8 978-3-540-46056-5. DOI: 10.1007/11871842_29. URL: http://link.springer.com/10.1007/11871842_29 (visited on 12/05/2022).

[KZZ19] Jiannan Ke, Dan Zhang, and Huan Zheng. “An Approximate Dynamic Programming Approach to Dynamic Pricing for Network Revenue Management”. In: *Production and Operations Management* 28.11 (2019), pp. 2719–2737. ISSN: 1937-5956. DOI: 10.1111/poms.13075. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/poms.13075 (visited on 03/11/2020).

[LDL11] Ryan Liu, Luther Dow, and Edwin Liu. “A Survey of PEV Impacts on Electric Utilities”. In: *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*. IEEE, 2011, pp. 1–8. URL: http://ieeexplore.ieee.org/abstract/document/5759171/ (visited on 05/23/2017).

[Lee+20] Jongmin Lee, Wonseok Jeon, Geon-Hyeong Kim, and Kee-Eung Kim. “Monte-Carlo Tree Search in Continuous Action Spaces with Value Gradients”. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 34.04 (04 Apr. 3, 2020), pp. 4561–4568. ISSN: 2374-3468. DOI: 10.1609/aaai.v34i04.5885. URL: https://ojs.aaai.org/index.php/AAAI/article/view/5885 (visited on 03/09/2023).

[LHG18] Chao Luo, Yih-Fang Huang, and Vijay Gupta. “Stochastic Dynamic Pricing for EV Charging Stations With Renewable Integration and Energy Storage”. In: *IEEE Transactions on Smart Grid* 9.2 (Mar. 2018), pp. 1494–1505. ISSN: 1949-3061. DOI: 10.1109/TSG.2017.2696493. URL: https://ieeexplore.ieee.org/document/7906472 (visited on 02/12/2024).

[LHZ18] Renzhi Lu, Seung Ho Hong, and Xiongfeng Zhang. “A Dynamic Pricing Demand Response Algorithm for Smart Grid: Reinforcement Learning Approach”. In: *Applied Energy* 220 (June 15, 2018), pp. 220–230. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2018.03.072. URL: https://www.sciencedirect.com/science/article/pii/S0306261918304112 (visited on 02/19/2024).

[Liu+18] Haoming Liu, Wenqian Yin, Xiaoling Yuan, and Man Niu. "Reserving Charging Decision-Making Model and Route Plan for Electric Vehicles Considering Information of Traffic and Charging Station". In: *Sustainability (Switzerland)* 10.5 (2018), pp. 1–20. ISSN: 20711050. DOI: 10.3390/su10051324.

[LLC14] A. Y. S. Lam, Y. Leung, and X. Chu. "Electric Vehicle Charging Station Placement: Formulation, Complexity, and Solutions". In: *IEEE Transactions on Smart Grid* 5.6 (Nov. 2014), pp. 2846–2856. ISSN: 1949-3053. DOI: 10.1109/TSG.2014.2344684.

[LM14] Linda Lønmo and Gerrit Muller. "7.1.2 Concept Selection - Applying Pugh Matrices in the Subsea Processing Domain". In: *INCOSE International Symposium* 24.1 (2014), pp. 583–598. DOI: 10.1002/j.2334-5837.2014.tb03169.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.2334-5837.2014.tb03169.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/j.2334-5837.2014.tb03169.x.

[LNG15] Jie Liu, M. Negrete-Pincetic, and V. Gupta. "Optimal Charging Profiles and Pricing Strategies for Electric Vehicle Charging Stations". In: *2015 IEEE Eindhoven PowerTech*. 2015 IEEE Eindhoven PowerTech. June 2015, pp. 1–6. DOI: 10.1109/PTC.2015.7232527.

[LP13] C. -Y. Cynthia Lin and Lea Prince. "Gasoline Price Volatility and the Elasticity of Demand for Gasoline". In: *Energy Economics* 38 (July 1, 2013), pp. 111–117. ISSN: 0140-9883. DOI: 10.1016/j.eneco.2013.03.001. URL: https://www.sciencedirect.com/science/article/pii/S0140988313000418 (visited on 11/21/2023).

[LPK11] Junghoon Lee, Gyung-Leen Park, and Hye-Jin Kim. "Reservation-Based Charging Service for Electric Vehicles". In: *Algorithms and Architectures for Parallel Processing*. Melbourne, Australia: Springer Berlin Heidelberg, 2011, pp. 186–195.

[LR17] Steffen Limmer and Tobias Rodemann. "Multi-Objective Optimization of Plug-in Electric Vehicle Charging Prices". In: *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2017, pp. 1–8. URL: https://ieeexplore.ieee.org/abstract/document/8285280/ (visited on 02/17/2024).

[LSP17] C. Latinopoulos, A. Sivakumar, and J. W. Polak. "Response of Electric Vehicle Drivers to Dynamic Pricing of Parking and Charging Services: Risky Choice in Early Reservations". In: *Transportation Research Part C: Emerging Technologies* 80 (July 1, 2017), pp. 175–189. ISSN: 0968-090X. DOI: 10.1016/j.trc.2017.04.008. URL: https://www.sciencedirect.com/science/article/pii/S0968090X17301134 (visited on 02/09/2024).

[LSS14] Oriol Lordan, Jose M. Sallan, and Pep Simo. “Study of the Topology and Robustness of Airline Route Networks from the Complex Network Approach: A Survey and Research Agenda”. In: *Journal of Transport Geography* 37 (May 1, 2014), pp. 112–120. ISSN: 0966-6923. DOI: 10.1016/j.jtrangeo.2014.04.015. URL: http://www.sciencedirect.com/science/article/pii/S0966692314000763 (visited on 09/17/2018).

[Lut+14] Benny Lutati, Vadim Levit, Tal Grinshpoun, and Amnon Meisels. “Congestion Games for V2G-Enabled EV Charging.” In: *AAAI*. 2014, pp. 1440–1446.

[MB22] Jan Mrkos and Robert Basmadjian. “Dynamic Pricing for Charging of EVs with Monte Carlo Tree Search”. In: *Smart Cities* 5.1 (1 Mar. 2022), pp. 223–240. ISSN: 2624-6511. DOI: 10.3390/smartcities5010014. URL: https://www.mdpi.com/2624-6511/5/1/14 (visited on 03/03/2022).

[McC01] P.S. McCarthy. *Transportation Economics: Theory and Practice*. Blackwell Publishers, 2001. URL: https://books.google.com.au/books?id=86PAtgAACAAJ.

[Mej+16] Naourez Mejri, Mouna Ayari, Rami Langar, and Leila Saidane. “Reservation-Based Multi-Objective Smart Parking Approach for Smart Cities”. In: *IEEE 2nd International Smart Cities Conference: Improving the Citizens Quality of Life, ISC2 2016 - Proceedings*. Trento, Italy: IEEE, 2016, pp. 1–6. ISBN: 978-1-5090-1845-1. DOI: 10.1109/ISC2.2016.7580840.

[MG15] Joy Chandra Mukherjee and Arobinda Gupta. “A Review of Charge Scheduling of Electric Vehicles in Smart Grid”. In: *IEEE Systems Journal* 9.4 (2015), pp. 1541–1553.

[MGK16] Jan Mrkos, Martin Grill, and Jan Kohout. “Tracking Users over Network Hosts Based on User Behavior”. U.S. pat. 20160352760A1. Cisco Technology Inc. Dec. 1, 2016. URL: https://patents.google.com/patent/US20160352760A1/en (visited on 09/05/2018).

[MK16] Mausam and Andrey Kolobov. *Planning with Markov Decision Processes: An AI Perspective*. Dec. 18, 2016. URL: https://www.microsoft.com/en-us/research/publication/planning-markov-decision-processes-ai-perspective/ (visited on 11/06/2019).

[MKJ18a] Jan Mrkos, Antonín Komenda, and Michal Jakob. “Dynamic Pricing Strategy for Electromobility Using Markov Decision Processes”. In: *10th International Conference on Agents and Artificial Intelligence (ICAART)*. 10th International Conference on Agents and Artificial Intelligence. Apr. 6, 2018, pp. 507–514. ISBN: 978-989-758-275-2. URL: http://www.scitepress.org/PublicationsDetail.aspx?ID=6c40zuVAzZE=&t=1 (visited on 04/06/2018).

[MKJ18b] Jan Mrkos, Antonín Komenda, and Michal Jakob. "Revenue Maximization for Electric Vehicle Charging Service Providers Using Sequential Dynamic Pricing". In: *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*. International Conference on Autonomous Agents and Multiagent Systems. Stockholm, Sweden: IFAAMAS, July 10–15, 2018, pp. 832–840. ISBN: 978-1-4503-5649-7.

[MV99] Jeffrey I. McGill and Garrett J. Van Ryzin. "Revenue Management: Research Overview and Prospects". In: *Transportation science* 33.2 (1999), pp. 233–256. URL: http://pubsonline.informs.org/doi/abs/10.1287/trsc.33.2.233.

[Nou+15] Mehdi Nourinejad, Sirui Zhu, Sina Bahrami, and Matthew J. Roorda. "Vehicle Relocation and Staff Rebalancing in One-Way Carsharing Systems". In: *Transportation Research Part E: Logistics and Transportation Review* 81 (Sept. 1, 2015), pp. 98–113. ISSN: 1366-5545. DOI: 10.1016/j.tre.2015.06.012. URL: http://www.sciencedirect.com/science/article/pii/S1366554515001349 (visited on 09/17/2018).

[NR15] Maleknaz Nayebi and Guenther Ruhe. "Chapter 19 - Analytical Product Release Planning". In: *The Art and Science of Analyzing Software Data*. Ed. by Christian Bird, Tim Menzies, and Thomas Zimmermann. Boston: Morgan Kaufmann, 2015, pp. 555–589. ISBN: 978-0-12-411519-4. DOI: 10.1016/B978-0-12-411519-4.00019-7. URL: http://www.sciencedirect.com/science/article/pii/B9780124115194000197.

[Orc+18] Simone Orcioni, Luca Buccolini, Adrianna Ricci, and Massimo Conti. "Electric Vehicles Charging Reservation Based on OCPP". In: *IEEE International Conference on Environment and Electrical Engineering and 2018 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*. Palermo, Italy: IEEE, 2018.

[Pis14] Hossein Pishro-Nik. *Introduction to Probability, Statistics, and Random Processes*. Kappa Research LLC, 2014. URL: https://www.probabilitycourse.com/ (visited on 01/17/2024).

[Pol14] David Pollard. "Yale Statistics 241/541". Lecture Notes. 2014. URL: http://www.stat.yale.edu/~pollard/Courses/241.fall2014/ (visited on 10/03/2020).

[PPM08] Michael Parkin, Melanie Powell, and Kent Matthews. *Economics*. Pearson Education, 2008. 884 pp. ISBN: 978-0-13-204122-5. Google Books: Kp6Ls7j3tVIC.

[RFH14] Pooya Rezaei, Jeff Frolik, and Paul DH Hines. "Packetized Plug-in Electric Vehicle Charge Management". In: *IEEE Transactions on Smart Grid* 5.2 (2014), pp. 642–650. URL: http://ieeexplore.ieee.org/abstract/document/6733308/ (visited on 05/23/2017).

[RFK15] Nadine Rauh, Thomas Franke, and Josef F. Krems. "Understanding the Impact of Electric Vehicle Driving Experience on Range Anxiety". In: *Human Factors* 57.1 (Feb. 1, 2015), pp. 177–187. ISSN: 0018-7208. DOI: 10.1177/0018720814546372. URL: https://doi.org/10.1177/0018720814546372 (visited on 11/01/2023).

[Rig+13] Emmanouil S. Rigas, Sarvapali D. Ramchurn, Nick Bassiliades, and George Koutitas. "Congestion Management for Urban EV Charging Systems". In: *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference On*. IEEE, 2013, pp. 121–126. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6687944 (visited on 11/04/2016).

[RRB15] E. S. Rigas, S. D. Ramchurn, and N. Bassiliades. "Managing Electric Vehicles in the Smart Grid Using Artificial Intelligence: A Survey". In: *IEEE Transactions on Intelligent Transportation Systems* 16.4 (Aug. 2015), pp. 1619–1635. ISSN: 1524-9050. DOI: 10.1109/TITS.2014.2376873.

[Sac+11] Martin Sachenbacher, Martin Leucker, Andreas Artmeier, and Julian Haselmayr. "Efficient Energy-Optimal Routing for Electric Vehicles." In: *AAAI*. 2011, pp. 1402–1407.

[SBK20] Sandeep Saharan, Seema Bawa, and Neeraj Kumar. "Dynamic Pricing Techniques for Intelligent Transportation System in Smart Cities: A Systematic Review". In: *Computer Communications* 150 (Jan. 15, 2020), pp. 603–625. ISSN: 0140-3664. DOI: 10.1016/j.comcom.2019.12.003. URL: https://www.sciencedirect.com/science/article/pii/S0140366419310990 (visited on 02/10/2024).

[Sch+19] Martin Schaefer, Michal Čáp, Jan Mrkos, and Jiří Vokřínek. "Routing a Fleet of Automated Vehicles in a Capacitated Transportation Network". In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). Nov. 2019, pp. 8223–8229. DOI: 10.1109/IROS40897.2019.8967723.

[SF18] Roberto Serrano and Allan M. Feldman. *A Short Course in Intermediate Microeconomics with Calculus*. Cambridge University Press, Sept. 13, 2018. 427 pp. ISBN: 978-1-108-42396-0. Google Books: 5rBoDwAAQBAJ.

[Shi+19] Syed Arbab Mohd Shihab, Caleb Logemann, Deepak-George Thomas, and Peng Wei. "Autonomous Airline Revenue Management: A Deep Reinforcement Learning Approach to Seat Inventory Control and Overbooking". In: (May 5, 2019). URL: https://openreview.net/forum?id=rJxRICuns4 (visited on 11/11/2019).

[Sil+16] David Silver et al. "Mastering the Game of Go with Deep Neural Networks and Tree Search". In: *Nature* 529.7587 (7587 Jan. 2016), pp. 484–489. ISSN: 1476-4687. DOI: 10.1038/nature16961. URL: https://www.nature.com/articles/nature16961 (visited on 12/06/2022).

[SKS18] Arne K. Strauss, Robert Klein, and Claudius Steinhardt. “A Review of Choice-Based Revenue Management: Theory and Methods”. In: *European Journal of Operational Research* 271.2 (Dec. 1, 2018), pp. 375–387. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2018.01.011. URL: https://www.sciencedirect.com/science/article/pii/S0377221718300110 (visited on 02/08/2024).

[Sol24] Matt Solomon. *The Rise of Batteries in Six Charts and Not Too Many Numbers*. RMI. Jan. 25, 2024. URL: https://rmi.org/the-rise-of-batteries-in-six-charts-and-not-too-many-numbers/ (visited on 02/02/2024).

[SSL99] Janakiram Subramanian, Shaler Stidham Jr, and Conrad J. Lautenbacher. “Airline Yield Management with Overbooking, Cancellations, and No-Shows”. In: *Transportation Science* 33.2 (1999), pp. 147–167. URL: http://pubsonline.informs.org/doi/abs/10.1287/trsc.33.2.147 (visited on 03/16/2017).

[Ste+16] Sebastian Stein, Enrico H. Gerding, Adrian Nedea, Avi Rosenfeld, and Nicholas R. Jennings. “Bid2Charge: Market User Interface Design for Electric Vehicle Charging”. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. AAMAS ’16. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 882–890. ISBN: 978-1-4503-4239-1. URL: http://dl.acm.org/citation.cfm?id=2937029.2937053 (visited on 10/09/2017).

[THL10] Jacopo Torriti, Mohamed G. Hassan, and Matthew Leach. “Demand Response Experience in Europe: Policies, Programmes and Implementation”. In: *Energy*. Demand Response Resources: The US and International Experience 35.4 (Apr. 1, 2010), pp. 1575–1583. ISSN: 0360-5442. DOI: 10.1016/j.energy.2009.05.021. URL: https://www.sciencedirect.com/science/article/pii/S0360544209002060 (visited on 11/28/2023).

[Tho22] Hannah Thompson. *Tens of Thousands in France Sign up to EDF Energy Saving Tempo Plan*. https://www.connexionfrance.com. Dec. 13, 2022. URL: https://www.connexionfrance.com/article/French-news/Tens-of-thousands-in-France-sign-up-to-EDF-energy-saving-Tempo-plan (visited on 11/29/2023).

[VA16] Theoni Versi and Mark Allington. *Overview of the Electric Vehicle Market and the Potential of Charge Points for Demand Response*. ICF Consulting Services, 2016.

[Wal+23] Daan Walter et al. *X-Change: Batteries*. Dec. 2023. URL: https://rmi.org/insight/x-change-batteries/ (visited on 02/02/2024).

[Wan+10] Hengsong Wang, Qi Huang, Changhua Zhang, and Aihua Xia. “A Novel Approach for the Layout of Electric Vehicle Charging Station”. In: *The 2010 International Conference on Apperceiving Computing and Intelligence*

*Analysis Proceeding*. IEEE, 2010, pp. 64–70. URL: https://ieeexplore.ieee.org/abstract/document/5709852/ (visited on 02/17/2024).

[Wan+14] Ran Wang, Ping Wang, Gaoxi Xiao, and Shimin Gong. “Power Demand and Supply Management in Microgrids with Uncertainties of Renewable Energies”. In: *International Journal of Electrical Power & Energy Systems* 63.0 (2014), pp. 260–269.

[WB18] Michael D. Wittman and Peter P. Belobaba. “Customized Dynamic Pricing of Airline Fare Products”. In: *Journal of Revenue and Pricing Management* 17.2 (Apr. 2018), pp. 78–90. ISSN: 1476-6930, 1477-657X. DOI: 10.1057/s41272-017-0119-8. URL: http://link.springer.com/10.1057/s41272-017-0119-8 (visited on 01/31/2022).

[WWX15] Ran Wang, Ping Wang, and Gaoxi Xiao. “A Robust Optimization Approach for Energy Generation Scheduling in Microgrids”. In: *Energy Conversion and Management* 106.0 (2015), pp. 597–607.

[Xia+16] Qiao Xiang, Linghe Kong, Xue Liu, Jingdong Xu, and Wei Wang. “Auc2Reserve: A Differentially Private Auction for Electric Vehicle Fast Charging Reservation”. In: *Embedded and Real-Time Computing Systems and Applications (RTCSA), 2016 IEEE 22nd International Conference On*. IEEE, 2016, pp. 85–94. URL: http://ieeexplore.ieee.org/abstract/document/7579930/ (visited on 11/03/2016).

[Xio+16] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Yeng Chai Soh. “Optimal Pricing for Efficient Electric Vehicle Charging Station Management”. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 749–757.

[Xio+18] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Ana L. C. Bazzan. “Optimal Electric Vehicle Fast Charging Station Placement Based on Game Theoretical Framework”. In: *IEEE Transactions on Intelligent Transportation Systems* 19.8 (Aug. 2018), pp. 2493–2504. ISSN: 1558-0016. DOI: 10.1109/TITS.2017.2754382. URL: https://ieeexplore.ieee.org/abstract/document/8064175 (visited on 02/17/2024).

[XL13] Hong Xu and Baochun Li. “Dynamic Cloud Pricing for Revenue Maximization”. In: *IEEE Transactions on Cloud Computing* 1.2 (2013), pp. 158–171. URL: https://ieeexplore.ieee.org/abstract/document/6671562/ (visited on 02/19/2024).

[Yan+13] Shun-Neng Yangab, Wei-Sheng Chengb, Yu-Ching Hsua, Chai-Hien Gana, and Yi-Bing Lin. “Charge Scheduling of Electric Vehicles in Highways”. In: *Mathematical and Computer Modelling* 57.11-12 (2013), pp. 2873–2882.

[YU16] Bunyamin Yagcitekina and Mehmet Uzunoglu. “A Double-Layer Smart Charging Strategy of Electric Vehicles Taking Routing and Charge Scheduling into Account”. In: *Applied Energy* 167.0 (2016), pp. 407–419.

[Yue+17] Cao Yue, Ning Wang, George Kamel, and oung Jin Kim. "An Electric Vehicle Charging Management Scheme Based on Publish/Subscribe Communication Framework". In: *IEEE Systems Journal* 11.3 (2017), pp. 1822–1835.

[Zou+16] Suli Zou, Zhongjing Ma, Xiangdong Liu, and Ian Hiskens. "An Efficient Game for Coordinating Electric Vehicle Charging". In: *IEEE Transactions on Automatic Control* (2016), pp. 1–1. ISSN: 0018-9286, 1558-2523. DOI: 10.1109/TAC.2016.2614106. URL: http://ieeexplore.ieee.org/document/7577848/ (visited on 11/04/2016).

[Zuo+19] Song Zuo, Mengjing Chen, Weiran Shen, and Pingzhong Tang. "Dispatching Through Pricing: Modeling Ride-Sharing and Designing Dynamic Prices". In: (2019), pp. 165–171. URL: https://www.ijcai.org/proceedings/2019/24 (visited on 09/03/2019).