NTNU – Trondheim
Norwegian University of
Science and Technology

Load-scheduling and Plug-in Hybrid Electric Vehicles in the Smart Grid

Eirik Daleng Haukedal

Master of Science in Computer Science
Submission date: June 2012
Supervisor: Pinar Öztürk, IDI
Co-supervisor: Boye Annfelt Høverstad, IDI
Axel Tidemann, IDI

Norwegian University of Science and Technology
Department of Computer and Information Science

Abstract

Peak demand is a well-known problem in the power grid. It denotes the amount of power required to supply consumers at times when demand is highest. Load peaks can have a negative impact on the stability of the power grid and maintenance costs for transmission and generation companies. Because of this, electricity prices often increase during peak demand. Currently, increasing use of plug-in hybrid electric vehicles (PHEV) further proliferates this problem because charging patterns are expected to coincide with peak demand hours, and especially the afternoon peak hours when people return home from work.

To avoid the problem that increasing PHEV demand will further aggravate peak demand hours, several different multi-agent scheduling mechanisms have been investigated, including two centralized scheduling mechanisms and two decentralized scheduling mechanisms. For both of the decentralized mechanisms, the PHEV agents choose their own charging plans without relying upon a centralized scheduler, while in the centralized scheduling mechanisms, the PHEVs agents defer control to a central agent for creating their charging plans.

From the results, we found that while both the centralized mechanisms and the decentralized mechanisms helped to reduce the average maximum peak, the performance of the centralized mechanisms proved to be highly dependent on how the day-ahead portfolio was calculated. Because of this, the overall best performer was decentralized mechanism, which gave the best results for reducing the average maximum peak without compromising the ability of the PHEVs to charge their batteries too much.

Preface

This thesis has been written as part of my Master's degree in Computer Science at the Norwegian University of Science and Technology in Trondheim. It is the conclusion of 20 weeks of work between January and June, 2012. The background for this thesis was established as part of a project report written in the fall of 2011.

Acknowledgements

I would like to thank my supervisors, Pinar Öztürk, Boye Annfelt Høverstad and Axel Tidemann. Their close support and guidance throughout the semester has been invaluable for the quality of the work and this thesis. I have enjoyed our weekly meetings, both as a source of inspiration, but also for keeping me motivated and focused. And while Pinar has been on sabbatical leave for much of this semester, she has always been accessible and actively interested in my progress, to which I will always be grateful.

June 10, 2012

Contents

List of Figures

List of Tables

1 Introduction

According to official Norwegian estimates, The Climate and Pollution Agency [2012] expects that plug-in hybrid electric vehicles (PHEVs), electric vehicles and hydrogen-based vehicles will account for 5% of the total car population in Norway by 2020, and that figure is expected to rise up to 26% by 2030. While this is arguably a positive development from a climate perspective, one of the anticipated challenges in the future smart grid is how to efficiently handle the extra load associated with charging the growing number of PHEVs. This is mainly a challenge, not because there is an insufficient overall capacity to accommodate the extra charge, but because the extra demand resulting from PHEVs recharging their batteries is expected to coincide with times at which demand is already at its highest, namely peak hours [Hadley and Tsvetkova, 2008].

One of the technologies from computer science that the next generation of power grids can benefit from is multi-agent systems. Multi-agent systems is a branch within computer science and artificial intelligence that is based upon theory from many disciplines of science. They can provide flexibility, scalability and fault tolerance to many areas in the next generation power grids [IEEE PES Multi-agent Systems Working Group, 2011] such as in diagnostics, distributed control (including managing loads) and modeling or simulation. Much of the theory is shared with economic theory and social theory. The close relationship to economic theory makes this type of technology very suitable to problem solving in domains where problems are highly interdependent. This is because the problem of distributing resources such as electricity is a multi-faceted problem. On the one hand, it is an optimization problem where there is a need for finding the most efficient way of managing and distributing electricity. On the other hand, it is also an economic problem where potentially limited resources need to be distributed fairly. These are problems that often can not be solved independently. This means that close relationship with economic theory makes multi-agent systems very applicable to this domain.

Today, energy is traded on spot markets, forward markets and futures markets. This trading is based on predictions about the future energy demand in the grid. Because electricity can be considered a zero-sum game where supply needs to balance demand; any imbalances need to be compensated for by either increasing demand or decreasing supply so that they cancel each other out. Hence, imbalances in the day-ahead portfolio from real-time electricity consumption is traded for on the spot markets at a penalty. In the current situation today, trading on the intraday market can only be limited by improving accuracy in the predictions that are used when trading electricity on the day-ahead market. However, one of the new capabilities of the coming Smart Grid is two-way communication between customer and the power companies. This opens up new possibilities for demand-side management and demand response whereby consumers can increase or decrease their electricity consumption, either by responding to price signals, or from participating in some type of demand-side management program.

Addressing the problem of peak hours, we have investigated two different approaches to load-scheduling where the real-time consumption of energy is influenced by controlling the charging of PHEVs; a centralized mechanism and a decentralized mechanism. This is enabled by the two-way communication capability of the Smart Grid, which allow real-time interaction between consumer and the electricity companies. In the current power grid, these kind of scheduling mechanisms are not possible because of the lack of communication capabilities. However, with the coming Smart Grid, PHEVs will be able to communicate their intentions to charge or to obtain information from electricity companies about future electricity demand, either of which can be used by the PHEV to schedule their charging plans socially economical.

In the first approach – the centralized scheduling mechanism – we have investigated an approach to load scheduling where the real-time consumption of energy is influenced by controlling the charging of PHEVs. If we at any time have information about 1) the current, real-time demand for energy in the grid, 2) the energy traded on the forward market (day-ahead) at that time, and 3) the state of the PHEVs in the grid, then we have sufficient information to schedule the charging of the PHEVs in a way that will be socially economical. By carefully constructing the day-ahead portfolio, the scheduling mechanism can be used to obtain a peak-shaving effect on the real-time energy consumption. To help the scheduling mechanisms we have implemented a simple learning mechanism that the PHEV agents can use in order to learn their own usage patterns. By learning the usage patterns, the PHEVs are better able to predict when they are likely to leave which will help the BRP agent to find the best schedule for the PHEV charging profiles.

In the second approach to load-scheduling, we have investigated a decentralized mechanism where no centralized schedulers are used. The idea is that that the PHEVs will generate their charging strategies on their own accord, by randomizing over a probability distribution that they create themselves. The important thing to consider with this mechanism is how this probability distribution is created. For this, we have considered two possible methods. One in which the probability distribution is uniform, and another in which the distribution is generated by the PHEV agent by interacting with a central agent that provides the PHEV agents with predictions about the future energy consumption in the grid.

1.1 Motivation

The main objective of this thesis is to investigate methods of how multi-agent systems can aid in demand-side management applications for the Smart Grid. To investigate such possibilities, a multi-agent system and a Smart Grid simulator was both developed. The first part of the thesis will focus on the simulator, which uses mathematical models of the different components of the simulated power grid. The most important model is the PHEV model, which is used by the simulator to model the behavior of consumers that own and operate PHEVs. The focus is to reflect uncertainty in the consumer usage patterns of the PHEVs. This is done to challenge the scheduling mechanisms that we have developed and experimented with.

The second part of this thesis relates to the multi-agent system that was developed. It is the MAS that implements the scheduling mechanisms that are used, while the models and the simulator discussed in the first part of the thesis provide the environment which the agents operate in. To successfully schedule the charging of the PHEVs, the different agents will cooperate and communicate with each other with the goal of finding a mutually beneficial schedule for all the agents.

1.2 Outline

The structure of this thesis is as follows: In Section 2, the background for the thesis will be discussed. This includes general information about the Smart Grid, Demand-side Management, Demand Response and Plug-in Hybrid Electric Vehicles, as well as some related work that has been done with respect to these areas. In Section 3, a set of research questions and hypotheses will be defined. Subsequently, in Sections 4 and 5, we will discuss the overall approach, how the simulator, models and multi-agent system relate to each other. Then, in Section 3, we will discuss the mathematical models used by the simulator to provide an environment for the agents in, while in Section 6, the multi-agent system will be outlined, where we will discuss the different agents, and their roles and responsibilities. In that Section we will also discuss the different multi-agent mechanisms that were developed, two centralized mechanisms and two decentralized mechanisms. In Sections 7, 8 and 9, we will outline how the experiments were performed, what the results of these experiments were, and discuss whether they proved or disproved the hypotheses made. Finally, in Sections 10 and 11, we will discuss which mechanism performed best, comparing them against each other, and present some suggestions for future work.

2 Background

Parts of this section is included from my project in the previous semester. My supervisor advised me to include these in this thesis for the sake of possible readers who are not familiar with the power domain.

2.1 Smart Grid

The current electric grid is faced with many challenges. For instance, 20% of the entire generation and transmission capacity in the current electric grid exist only to meet peak-demand, meaning that this capacity is only used about 5% of the time [Farhangi, 2010]. In countries with adequate generation and transmission capabilities, this solution works most of the time, although inefficiently. In developing-countries where generation or transmission capacity might be limited, peak-demand often means rolling blackouts over parts of the network. In countries like India, this may affect the lives of a significant amount of people.

Smart Grid has been called the next-generation electricity-grid [Farhangi, 2010], and it is difficult to provide an exact definition of what it is, because it is still evolving as a concept. However, what can be said about it is, it refers to a class of technology aimed to help bring utility electricity delivery systems into the 21st century, taking advantage of the modern advances in computing and automation [Office of Electricity Delivery and Energy Reliability, 2011]. One of the things that is expected of it, is that is should address the shortcomings of the existing grid and the challenges that were created with the deregulation of the market and the increasing interdependencies among the critical infrastructures in the modern electricity grid [Amin and Wollenberg, 2005].

It is expected that Smart Grid will bring about a revolution in how customers interact with the electric companies. Presently, there is little interaction between customer and utility companies. Consumers buy their electricity from retailers, paying a certain price for the electricity that is consumed, but it is difficult for the utility companies to know how much electricity is being consumed in real-time, and this can affect the stability of the system. However, one of the key steps towards the Smart Grid is to open up two-way communication lines between the consumer and the utility companies, giving both consumer and utility company more real-time information on what is happening in the grid. This, and some of the other expected key differences between the current grid and the next generation grid, is shown in table 2.1.

Because consumers pay for electricity consumed, and not for how much that is generated, utility companies have to ensure that there is a sufficient capacity of electricity in the grid to meet consumer demands. Today, this is done by keeping energy reserves, but this method is inefficient and expensive. It is also becoming an increasingly difficult problem for the utility companies to expand their generation capacity in line with the increasingly growing demands for electricity [Farhangi, 2010]. This means that the next-generation grid need to move in a direction of becoming less dependent on keeping energy reserves, so that as much of the generation capacity as possible can be effectively utilized.

Existing Grid	**Intelligent Grid**
One-Way Communication	Two-Way Communication
Centralized Generation	Distributed Generation
Hierarchical	Network
Few Sensors	Sensors Throughout
Blind	Self-Monitoring
Manual Restoration	Self-Healing
Failures and Blackouts	Adaptive and Islanding
Manual Check/Test	Remote Check/Test
Limited Control	Pervasive Control
Few Customer Choices	Many Customer Choices

Table 2.1: The smart grid compared with the existing grid [Farhangi, 2010]

One of the key steps towards the evolution of the Smart Grid is the introduction of Advanced Metering Infrastructure (AMI) units. For instance, critical capability such as demand response may not be possible without them [Farhangi, 2010]. In essence, an AMI unit can be thought of as a more advanced version of the Automatic Meter Reading unit that exist today. The most significant differences is the two-way communication capability and near real-time sensors that the AMI unit has. These devices will record real-time electricity consumption where they are installed, and they will be able to communicate this information directly to the utility companies in fixed intervals of time. In Norway, a recording frequency of 15 or 60 minutes is being considered [NVE, 2011].

2.2 Demand-side Management

The term demand-side management (DSM) refer to techniques for influencing consumer demand of energy through various means. Many demand-side management techniques relate to ways of reducing peak-demand in the electric grid, also called peak-shaving. Reducing these is important for many reasons. For instance, rolling blackouts are intentionally engineered blackouts that are used when demand for electricity in the grid exceeds the power supply capabilities of the network. The need for intentionally engineered blackouts arise either when there is insufficient generation capacity, or when there is insufficient transmission capacity to cope with the current demand for electricity in the grid. Because demand-spikes put extra pressure on the network, these are often the primary cause for the need for rolling blackouts. Furthermore, because of the zero-sum nature of the electric grid, if demand for electricity increase, then production must also increase. Because of this, utility companies have to keep energy reserves to be able to cope with the extra demand. If demand-spikes could be reduced or eliminated in the network, there would be less need for peaking plants and top load generators. Many ways and techniques for peak-shaving have been discussed, such as load-deferring through micro-storage units [Ramchurn et al., 2011] and district heating systems [Wernstedt et al., 2007].

An important note on load-shifting techniques is that the overall goal is to reduce the peak-to-average ratio. If load-shifting is done without some form of coordination, or structure to enforce an even distribution of consumption over time, then there is a risk that peak-demand is simply shifted from peak-hours to nonpeak-hours. This is called load-synchronization [Mohsenian-Rad et al., 2010]. In other words, the overall goal in load-shifting is to minimize the peak-to-average ratio, or $_{peak}/L_{avg}$, where $_{peak}$ is the overall load during peak hours, and $_{avg}$ is the average load during a day. An example of load-shifting can be seen in figure 2.1, where demand is shifted away from peak-hours.

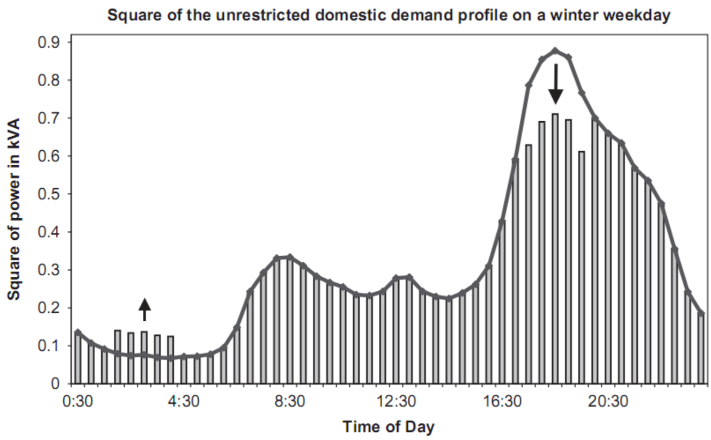

Figure 2.1: Effects of demand-side management and load-shifting techniques [Shaw et al., 2009]

2.3 Demand Response

demand response (DR) is one of the important new features of the Smart Grid. In short, it is the ability to influence consumer demand for electricity by price, monetary incentives or utility directives [Stavrogiannis, 2010]. Although DR techniques might seem similar to demand-side management (DSM) techniques, there are some subtle differences. Typically, in DR programs, end-users respond to price signals to reduce their electricity consumption in the short term. In DSM programs, however, the overall goal is to achieve energy efficiency, such as load-shifting. In load-shifting techniques the overall demand for electricity is not necessarily reduced, but is more evenly distributed over a day. While this may incidentally reduce electricity pricing for consumer, this is indirectly as a consequence of the load-shifting. In DR programs, however, any shift in energy consumption is caused directly by consumers responding to price signals.

In figure 2.2 and 2.3, the different effects of hypothetical DR and DSM techniques can be seen. Figure 2.2 shows the effect of a DSM program on overall

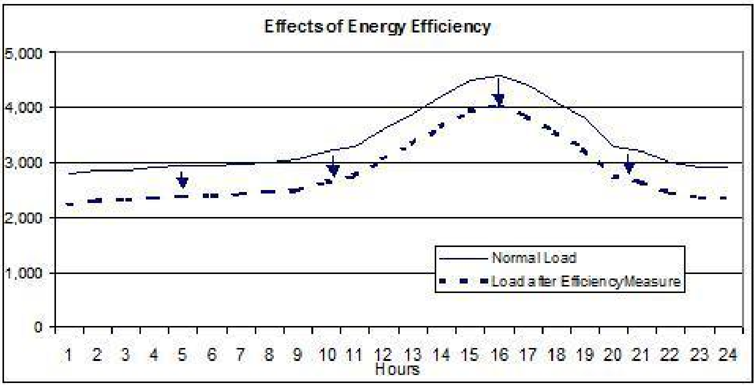

Figure 2.2: Improved energy efficiency using demand-side management techniques [Rowles, 2011]

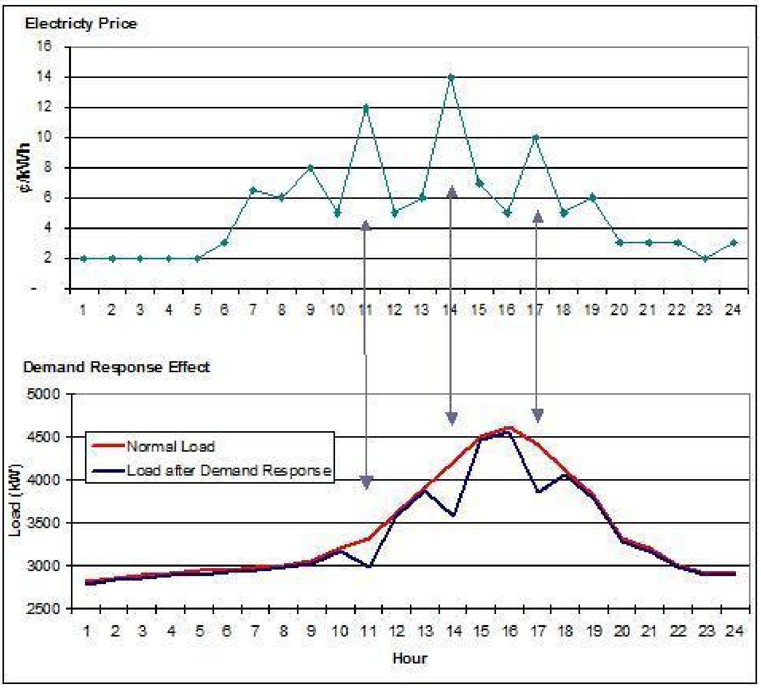

Figure 2.3: Potential effect of a demand response program [Rowles, 2011]

electricity demand, showing how the overall demand is reduced uniformly. This represents how DSM techniques can be used to improve energy efficiency. In figure 2.3, however, it can be seen how the price per kWh of electricity affects consumer demand for electricity. The upper graph shows price compared to time of day, while the lower graph shows aggregated demand in the grid, both during normal load and with the effects of DR applied. As can be seen, the effects of DR on demand is most noticeable during peak-hours, lowering the overall average demand. However, while the overall demand during peak-hours is lowered, the upper limit of peak-demand is almost untouched, meaning that peak-to-average ratio actually increases in the DR program. This is because prices goes down when the consumption goes down. But if prices go down too far, then consumers are given incentive to increase their consumption again, so the consumption goes back up. Because one leads to the other, you will inevitably get these periodic patterns of fluctuating demand. This does not necessarily mean that DSM is a better solution, however. The benefits of DR is that they are often relatively cheap to implement if automated control systems are already in place. Considering how many countries are also mandating the introduction of the AMI, this is an argument in favor of DR. Also, while DR and DSM might have different potential and limitations, they are not mutually exclusive techniques. In fact, they are more likely to be complementary techniques in the next generation power grid, each contributing with its own advantages.

So what is the difference between demand-response and demand-side management? In short, it is a difference in approaches. In demand response programs, the general idea is to influence consumer demand by price signals, so this approach directly involves and depends upon consumer interaction. In demand-side management, however, the focus is more directed towards utility-, transmission- or generator companies. This includes load-shifting, reducing peak-to-average ratio and increasing energy efficiency. While some DR programs might achieve some of the same effects as in DSM, the consequences are caused by different means. For instance, different pricing mechanisms can theoretically produce a load-shifting effect similar to load-shifting programs in DSM, but it is not guaranteed, and it ultimately depends on how the consumer reacts to the price signals. However, in load-shifting programs in DSM, the ultimate goal is the load-shifting itself, without affecting consumer awareness of it. This means that it should be transparent to the consumer whether there is a load-shifting DSM program running, or not. This would not be the case in a DR program, as it directly involves and depends on consumer interaction (either through an agent-proxy, or through the consumer itself). If a consumer lowers its consumption by responding to a price signal, then this reduction must come at an expense of something, for instance by turning off some lights or postponing some activity that depends upon electricity to a later time. However, this also means that the benefits of DSM programs may not be directly obvious to the consumer, as the benefits are primarily for the electricity industry. But, indirectly, and in a market economy, if a utility company can reduce its margins, then it is also able to offer lower prices on its services. Load-shifting through DSM would improve energy efficiency, which would save costs for the utility companies, which would allow them to lower the margins on their electricity prices.

2.4 Plug-in Hybrid Electrical Vehicles

One of the concerns about the future of the Smart Grid is how plug-in hybrid electric vehicles will be integrated into the Smart Grid. Plug-in Hybrid Electric Vehicles (PHEVs) are expected to become a significant challenge for the power grid in the future, as they are expected to considerably increase the demand for electricity (predictions expect that around 30% of the cars in Belgium will be PHEVs by 2030, and that they will account for 5% of the total electricity consumption in Belgium by 2020 [Herbruggen and Zeebroeck, 2006]). Further troubling is that people are likely to start recharging the batteries immediately after having used them (i.e. when they get home from work). Because most people use their cars in the same way (to get to and from work), we get what we call rush hours in the transportation grid. This is analogous to peak-hours in the power grid, and because people use electricity when they are home, this means that peak-hours will generally coincide with the end of rush hour. This also means that it is likely that people will start recharging the batteries coincidentally during peak-hours, adding to an already considerable problem. It is therefore necessary to consider ways of managing or scheduling consumption in a way such that the safety and stability of the power grid is preserved.

The plug-in hybrid electric vehicle (PHEV) is a close cousin to the hybrid electric vehicle that exist on the market today. The difference is mainly that it has a larger battery and that, in addition to be chargable through the gasoline engine, it can also be charged by plugging the vehicle into the electric grid [U.S. Department of Energy, 2012]. This has the advantage that the PHEV can further displace fuel usage compared to the HEV.

2.5 Multi-agent systems

Multi-agent systems (MAS) is a field of study within A.I. that takes inspiration from a number of disciplines, such as economics, software engineering and social sciences to name a few. From a computer science perspective, autonomy is the key feature of MAS [Russell and Norvig, 2010]. Agents often possess other qualities. For instance, an autonomous agent is an entity that is capable of learning through experiences on how their actions affect the environment, reasoning, and having beliefs about their environments. They could also have desires and intentions to act upon those desires, to achieve a goal that they desire. There are many different forms of agents, and the field of MAS is the study of how these agents interact with each other when there are several agents in the environment.

Because agents can have different goals, intentions and beliefs, it becomes important to understand how populations of heterogeneous or homogeneous agents will interact. For instance, in some scenarios it could be beneficial for agents to form coalitions, co-operating to achieve some common goal. And when they have different goals, they may co-operate if this is mutually beneficial to each other. In other scenarios, however, they could be primarily self-interested and competitive.

In the rest of this Section, we will discuss how agents can make rational decisions using Game Theory, when competing against other agents in environments where they are primarily self-interested. First, we will look at a simple type of game, the Prisoner's Dilemma, in which we find and describe the Nash equilibrium. Second, we will look at a special type of game, called load balancing games, which defines the problem of load scheduling computing tasks over multiple machines in computer science. Also, while Game Theory does not preclude co-operative behavior, the underlying assumption is that the Agents choices are made based on their own self-interest. However, it is possible to also describe altruistic behavior in Game Theory as well, using Evolutionary Game Theory, but this outside the scope of this thesis and has been omitted.

2.5.1 Game Theory

Game Theory is often used to analyze what is the rational decisions that an agent might make in a multi-agent system. It is originally a mathematical theory that dates back as early as the 1950s through the works of many scholars, including John von Neumann among others. It was the study of two-person zero-sum games and its associated minimax theorem [Wooldridge, 2002]. However, practical applications of Game Theory were limited until John Forbes Nash extended the theory with his concept of a Nash equilibrium. The concept of a Nash equilibrium made it possible to extend the theory beyond two-person, zero-sum games, and it has since been used in many areas of science, including economics, computer science and biology to name a few. Later, other branches of Game Theory have appeared, including Evolutionary Game Theory that extends the theory into multi-player games, although the type of game has slightly different characteristics than in original game theory (GT).

To explain Game Theory, a classic example of the Prisoner's Dilemma is often used. In it, two prisoners, p_1 and p_2 are kept apart, unable to communicate with each other, and both are given the choice to either cooperate (C) or defect (D). If both prisoners chooses C, then both agents will get sentenced to 1 month in prison. However, if one prisoner, p_1 chooses C and the other, p_2 chooses D, then prisoner p_1 gets 1 year in prison, while p_2 goes free. Conversely, the same holds true for prisoner p_2 if p_2 chooses C while p_1 chooses D. If both prisoners choose D, then they both get 3 months in prison. The problem in deciding what to do, then becomes dependent on what the other agent chooses. The choices of actions and their effects are illustrated in a pay-off matrix in Table 2.2 where the outcome of their choices are represented with payoff values.

The problem for the prisoners in the Prisoner's Dilemma example is that the reward of choosing some action becomes dependent on the choice of the other prisoner. Because the prisoners have no way of communicating with each other, they both have to make an assumption of what the other prisoner is likely to do. The best possible scenario for p_1 is if p_1 chooses D while p_2 chooses C. However, if p_1 chooses D, then it risks that p_2 might also have chosen D, and they both get 3 months. In this case, they would have been better off by choosing C instead, where they both would have gotten 1 month. But if p_1 had chosen C, then it would have had the risk that the other prisoner would have chosen D, leading to the worst possible outcome it could have had.

In Game Theory, finding a Nash equilibrium is often used to help make rational decisions about which strategies to play in games like the Prisoner's Dilemma. The idea is that there might exist some strategies (s_1, s_2) in which neither player have any incentive to deviate from their choice of strategy. More formally, a Nash equilibrium is defined for two strategies s_1 and s_2 when [Wooldridge, 2002]:

1. under the assumption that prisoner p_1 plays s_1, prisoner p_2 can do no better than play s_2, and

2. under the assumption that prisoner p_2 plays s_2, prisoner p_1 can do no better than play s_1.

In the Prisoner's Dilemma, a Nash equilibrium can in fact be found when both prisoners choose to defect, (D, D). Assuming that prisoner p_1 defects, p_2 can do no better than to defect either, otherwise it would get the worst possible outcome by changing its strategy, and conversely, the same argument holds true for the other prisoner. In fact, (D, D) is the only Nash equilibrium that can be found in this game. Even though the sum of the total payoff for both players would be highest with (C, C). However, while (C, C) would give the highest total payoff, it is not a Nash equilibrium, because either agent would benefit individually from changing their strategy and defecting. (C, D) and (D, C) are neither Nash equilibriums, because in this case, the prisoner that is playing strategy C would benefit from changing its strategy. In this sense, playing strategy D would be a rational choice for both prisoners, even though it is the third least preferable outcome.

		Player p_1	
		C	D
Player p_2	D	3 3	0 5
	C	5 0	1 1

Table 2.2: A pay-off matrix for the Prisoner's Dilemma.

2.5.2 Load balancing games

A classic problem in computer science is how to optimally allocate computing tasks to machines so that the maximum load over all machines, or *makespan*, is minimized. In Game Theory, such problems are called load balancing games. The problem is defined by considering an agent that has a task that it wishes to place on one of a set of machines, assigning the task to the machine so that the work load is most evenly balanced over all the machines. More formally, the problem is defined by a set of m identical-, or uniformly related machines, with speeds $s_1 \dots s_m$. The intentions is to distribute $[n]$ tasks with weights $w_1 \dots w_n$. If the set of pure strategies for an agent is [m], then the the goal is to find an optimal assignment $A : [n] \rightarrow [m]$ that will minimize the makespan:

$$l_j = \sum_{i \in n j = A(i)} \frac{w_i}{s_j},$$

where l_j is the load on machine j.

(a) Nash equilibrium (b) Socially optimal (c) Non-equilibrium

Figure 2.4: Load scheduling assignment. Figure 2.4a, shows an assignment which is a Nash equilibrium, while Figure 2.4b and 2.4c shows the socially optimal assignment and a non-equilibrium assignment

These load balancing games also applies to Mixed strategies in Game Theory, where probabilities are assigned to the set of the pure strategies. Formally, it is defined as a probability distribution that assigns to each available action a likelihood of being selected. For Mixed strategies this means that: If p_i^j is the probability that agent i assigns it task to j, then a strategy profile is defined by $P = (p_i^j)_{i \in [n], j \in [m]}$. This means that if the strategy profile P contains the probabilities for all the possible assignments to the machines, over all the tasks, then the expected load of machine j is defined by

$$E[l_j] = \sum_{i \in n} \frac{w_i p_i^j}{s_j},$$

and the *social cost* of *strategy* P is $cost(P) = E[cost(A)] = E[max_{j \in [m]}]$, the expected maximum load over all the machines.

Nash equilibriums in load balancing games

In load balancing games, a Nash equilibrium is defined for a strategy profile P when only assignments that minimize the makespan have positive probabilities. If the strategy profile P has only one assignment with a positive probability, then the strategy is said to be *pure*. For pure strategies the Nash equilibrium is defined when no agent has any incentive to assign its task to another machine. Also, we can further assume that *every instance of the load balancing game admits at least one pure Nash equilibrium.* Proof of this proposition can be found in the textbook

Algorithmic Game Theory by Nisan et al. [2007]. Essentially this means that for an assignment A that is not in a Nash equilibrium, there exists a *finite sequence of improvement steps* whereby an agent can improve its utility by changing its assignment, ultimately leading to a Nash equilibrium. Remember, however, that a Nash equilibrium is not guaranteed to be the socially optimal solution. This is illustrated in Figure 2.4, where a Nash equilibrium is compared to the socially optimal assignment and a non-equilibrium assignment where further improvement steps can be made. Further, since it is possible that there may be several Nash equilibriums, there is no guarantee that the sequence of improvement steps will not lead to the worst Nash equilibrium. The cost of selfish load balancing games is then calculated as the ratio between the cost of the worst possible Nash equilibrium and the socially optimal cost. This is called the *price of anarchy*.

2.6 Related work

Boucké and Holvoet [2011] discuss strategies for real-time scheduling of the charging of PHEVs. The system is a hierarchical MAS, with 3 different kind of agents, where one agent is responsible for charging the PHEV, another agent is responsible for preventing overload of the transformer, and one agent that is responsible for minimizing the cost of imbalances, which is the penalty that is payed on the intraday market for a negative or positive imbalance of electricity from the day-ahead portfolio to the real-time consumption. The last agent is called a Balancing Responsible Party agent (BRP), and is responsible for creating the load-schedules for all of the PHEVs that are connected to it. It does this by creating an intention graph of all the PHEVs in its distribution grid, a graph containing information about each of the PHEVs intentions to charge, including how much they want to charge, and until when they want it charged. This is done by each PHEV communicating their desire to charge to the transformer agent. The BRP agent aggregates the information from all its transformer agents, and then selects which agents should charge for the next time step depending on the given strategy.

For load-scheduling, they evaluate two strategies, a reactive- and a proactive strategy. In the reactive strategy, imbalances are postponed as long as possible. That is, it will schedule PHEV charging to perfectly balance the real-time electricity consumption with the predictions that were made on the day-ahead market. By doing so, the portfolio remains balanced for as long as possible, but at a risk of large future imbalances. In the proactive strategy, the BRP agent will use its intention-graph to reduce the average distance between the real-time consumption and the predicted day-ahead consumption. This reduces the risk of large future imbalances, but is more dependent on accurate real-time predictions. In simulations, they show that the reactive strategy is able to reduce imbalance costs by 14%, and the proactive strategy by up to 44%

An extension of the system described above, is also discussed in [Vandael et al., 2011], which builds upon previous work in Vandael et al. [2010], Boucké and Holvoet [2011]. However, in this article the focus is not on load scheduling, but rather on how plug-in hybrid vehicles can be used as a primary reserve capacity in the power grid. For instance, in Europe, the constant net frequency is 50 Hz, so to maintain stability, the actors in the grid have to ensure that the frequency do not deviate too much from this value. If the frequency is higher, then there is an oversupply of electricity to the grid. If it's lower, there is an undersupply. To ensure stability, the grid should be able to react quickly to deviations in frequency, e.g. by using primary reserve capacities when the frequency falls below 50 Hz. However, primary reserve capacities are typically expensive and they generate normally little power. This means that, because of the potential technology such as District Heating Systems (DHS), PHEVs might be used as primary reserve capacities.

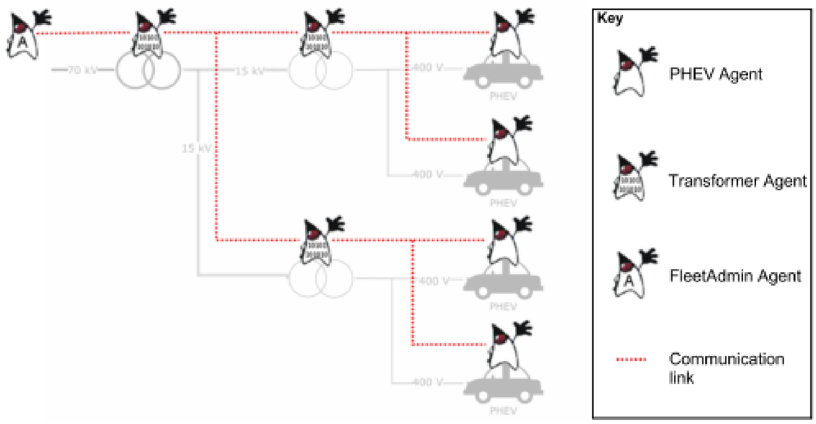

Figure 2.5: Architecture of the system

An overview of the architecture of the system can be seen in figure 5.7a. It is a hierarchical system of agents, each with their own responsibilities. There is a PHEV agent, a Transformer agent and a FleetAdmin agent. The PHEV agent is responsible for charging the battery of its PHEV in time, the Transformer agent is responsible for flattening the load of its transformer so that it doesn't overload, and the FleetManager is responsible for managing a number of PHEVs as primary reserve capacities. In essence, the system works by the PHEV agent communicating its intentions to charge to the Transformer agent. The FleetAdmin will then create an intention-graph of all the PHEVs in the grid, through its Transformer agent. It uses this intention-graph to decide which PHEV will charge when, and how much based upon one of a few predetermined strategies. This information is then communicated back to the PHEV agents. The extended version of this system, which can make PHEVs act as primary reserve capacities add two extra steps: Each PHEV calculates its potential to increase or decrease its charging power. And, this information is transmitted back to the FleetAdmin agent, which aggregates this information to calculate a percentage of charging power that the PHEVs should adjust their power in direction of.

3 Objectives

In this Section, we will first define the research questions posed by this thesis, after which we will define a few hypotheses regarding those questions.

3.1 Research Questions

In this section we will define three different research questions concerning peak-load, stability and end-user quality of service.

Research Question 3.1 *Is it possible to obtain a peak-shaving effect in the Smart Grid by using a centralized or decentralized scheduling mechanism for the scheduling of PHEV loads?*

This question will address peak-load in the system.

Research Question 3.2 *Can a mechanism be designed that will schedule the charging of PHEV loads, while also maintaining the stability of the system with respect to transformer capacity constaints?*

This question will address the stability of the system.

Research Question 3.3 *Can a scheduling mechanism be designed that will be fair with respect to the end-user?*

This question will address end-user quality of service

3.2 Hypotheses

To investigate the research questions made above, we have designed a few hypotheses that we are going to test.

Hypothesis 3.1 *We hypothesize that a carefully constructed day-ahead portfolio can be used by a centralized scheduling mechanism to obtain a peak-shaving effect, where PHEV loads are scheduled away from peak-hour.*

This hypothesis addresses research question 3.1

Hypothesis 3.2 *We hypothesize that a decentralized strategy can reduce peak-load in the system by distributing PHEV demand, either uniformly, or according to a weighted probability distribution based on predictions about the future energy consumption.*

This hypothesis addresses research question 3.1

Hypothesis 3.3 *We hypothesize that the TRF agent and the scheduling mechanisms can be designed in a way that will minimize the significance and number of transformer capacity violations.*

This hypothesis addresses research question 3.2

Hypothesis 3.4 *We hypothesize that by PHEV agents participating in the scheduling mechanisms, the scheduling mechanisms can be made so that they do not compromise the ability of the PHEV agents to maintain an acceptable quality-of-service for its end-user. This assumes that the PHEV agents can be made to learn about their own usage patterns, to help create the best possible charging strategies.*

This hypothesis addresses research question 3.3

Hypothesis 3.5 *We hypothesize that a carefully designed decentralized mechanism can perform at least as well as a centralized scheduling mechanism in obtaining a peak-shaving effect, and in charging the PHEVs and maintaining stability in the grid.*

This hypothesis addresses research question 3.1

4 Overall Approach

To test the hypotheses defined in Section 3.2, an open source simulator was developed, where PHEVs, Transformers, and other entities capable of producing or consuming energy are simulated within a grid delimited by a balancing responsible party, where each agent in the simulator is assigned a model that it is responsible for. The simulator is designed to be modular so that other agents and mechanisms can be implemented and integrated without needing to make changes to the simulator, or the mathematical models used to simulate the environment.

The multi-agent system, and the different mechanisms will be implemented in the simulator to be used in experiments where historical data is used to simulate houses and small businesses. Consumers that own and operate PHEVs will be simulated mathematically, modeled to reflect uncertainty in the usage patterns of the consumers. The results of the experiments performed on the scheduling mechanisms will be compared against results from running the simulator using no scheduling mechanisms.

4.1 Development Tools

The simulator, models and multi-agent system was developed in the .NET framework using the F# language[1]. F# is a multi-paradigm, functional and object-oriented programming language based upon OCaml. The reason behind this choice in framework and programming language was the strong support for *agent design patterns*. For instance, the support for discriminated unions and its strong typing system make F# a good fit for systems in which type-safe message passing is important. Being an officially supported .NET language, it also has access to existing libraries for numerics and statistics that can be found in the .NET framework and community. It also has good support for cross-language interoperability. Additionally, threads in F# are light-weight and use the ThreadPool concept of the .NET framework. This is an advantage over other frameworks, such as JADE and Java, where each agent has its own thread.

Further, the multi-agent framework was built from the bottom up, and no other frameworks, like JADE or Zeus were used. The custom-designed messaging protocol will be further discussed in Section 6.5. For sending and receiving messages, the simulator uses the MailboxProcessor available through the .NET framework.

4.2 Architecture

The architecture of the system can be categorized into three main groups, the simulator itself, the mathematical models that provide the environment for the agents, and the multi-agent system. For the models, we have three main types that represent physical entities, the PHEV model, the Transformer model and the PowerNode model. The PowerNode model represent the individual houses, small businesses or distributed energy facilities in the grid. In the simulator they are

[1]Source code for the open source simulator is available through github [Haukedal, 2012]

modeled as a node in the grid, which either consumes or produces electricity. For the Transformer model, this is modeled as a node in the grid where electricity passes through. For the electricity throughput, they have a maximum capacity that they can handle at any given time. In addition, they are either connected to a parent Transformer or a Balancing Responsible Party, and they can have one or more PowerNodes or PHEVs assigned to it as child nodes in the grid. The last model, the PHEV, represents a physical PHEV entity. It is modeled by an average charging- and discharge rate, a maximum battery capacity and a current throughput. They also have a set of probability distributions that determine the frequency and duration of when they are used. This is generated from pre-defined PHEV profiles when creating the PHEV models, and are meant to represent individual consumers.

For each of the models described above, there is a corresponding agent associated with it. In total, there are 4 types of agents: A PHEV agent, a Transformer agent, a PowerNode agent and a BRP agent. The PHEV agent has two targeted tasks: First, its main task is to ensure that the PHEV is adequately charged when it needs to be. Second, it should try to charge the PHEV batteries so that the social welfare of the power grid is considered.

For the other agents, the PowerNode agent has no particular targeted task. They observe their current demand as time passes in the simulator, and can provide information about this demand to the other agents if requested. The BRP can have up to two targeted tasks, depending on which mechanism is used. Common for all mechanism, it has a targeted task to reduce the average maximum peak in the power grid, while in the centralized mechanism, it has an additional targeted task to minimize the imbalances between the day-ahead portfolio and the real-time consumption. The last agent, the Transformer agent, has one targeted task, to prevent the electricity throughput to exceed the maximum capacity of its Transformer model.

In Figure 4.1, an overview of the system can be seen, where the relationship between the models, the simulator and the agents is distinguished.

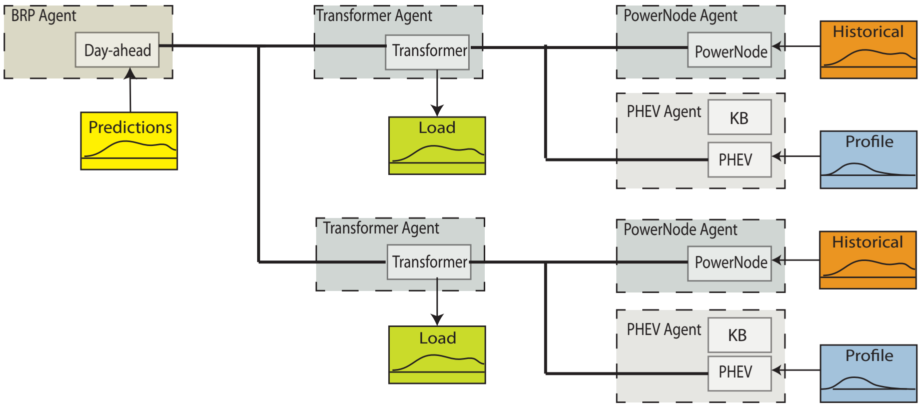

Figure 4.1: The relationship between the different components in the system, showing the different models and agents. The figure shows how the PHEV and PowerNodes are connected to its parent Transformer. Note, also the input that is used to generate the PowerNode and PHEV, illustrated by the orange and blue boxes. The Historical data that is used as input to the PowerNode is a function f(t) that determines its electricity consumption, while the Profile that is used as input when creating the PHEV model is a function f(t) that returns a set of vectors, where each vector contains a probability that the PHEV leaves on a trip, the duration of the trip, and also how much of the trip is spent driving. In addition, each PHEV agent has a knowledge base (KB) that it uses for a simple learning mechanism that will be discussed further in Section 6.2.1. The green boxes represent the output from the Transformer model to the simulator, representing the accumulated consumption over all its child nodes. The yellow box represent the input when creating the day-ahead portfolio. How the day-ahead portfolio is calculated will be further discussed in Section 5.4

5 Models

The architecture of the simulator is based on 4 main models: The Transformer model, the PHEV model, the PowerNode model and the Day-ahead model. The relationship and hierarchy between these models can be seen in Section 4, Figure 4.1. In addition to these, there is also a model for the day-ahead profile.

5.1 PowerNode model

The PowerNode model is the simplest model in the simulation. Essentially, it represents anything in the grid that can consume or produce power, and which is not a PHEV. For instance, it can represent a household, consuming power, in which case it will contribute a negative flow of electricity to the power grid. Or it can represent a distributed energy facility, contributing a positive flow of electricity.

In terms of the complexity of the model, the main thing to consider is how the model decides what the flow of electricity will be. In this simulator, the flow of electricity is determined by historical data collected from smart meters installed in real houses. In essence, this means that each PowerNode model in the simulator will be represented by one household, or equivalent, from the dataset used. While the PowerNode model could also represent a power producing unit in the grid, the historical data that we used to generate the PowerNode models included only recordings made on power consuming units. The flow of electricity at time t in the simulator is then determined by the recorded consumption at corresponding time t in the dataset for that household.

Considering that the resolution of the simulator, with respect to time, is in 15 minute intervals, and the historical data used was recorded in hourly intervals, the dataset was processed to match the resolution used in the simulator. To process the dataset, power levels in between the hourly readings were interpolated using splines. More specifically, the Akima spline interpolation algorithm [Alglib, 2012] was used to interpolate the readings to match a 15 minute resolution. The benefit of this particular spline interpolation algorithm, its stability with respect to outliers.

5.2 Transformer model

In the simulation, the transformers are the nodes in the power grid that will connect the PHEV agents. They are relevant in the simulation in the sense that they connect the different parts of the grid, and they have a maximum capacity of electricity throughput that should not be exceeded. In other words, BRP agents *should not* schedule the charging of PHEVs so that it violates the capacity constraints of the transformer models. This means that the Transformer model will need to have a known value for the maximum capacity of energy that it can handle at any given time, as well as to be able to hold information about the amount of electricity that it is currently handling. For maximum capacity, we have used standardized values for power grids in Norway, with capacities in the range of 220-240V for low-voltage transformers, 11kV for high-voltage transformers and 66kV for regional voltages.

5.3 PHEV model

To model PHEVs in the simulation, the simulator tries to reflect the behavior of PHEVs and consumers realistically. It does this by assuming that the behavior of the PHEVs can be modeled using statistical distributions such as, for instance, distributions around when people go to work in the morning and around the time at which they are expected to return home. Additionally, each PHEV has a value for how long it takes for the PHEV to reach its destination, including information about the average battery consumption while driving and average recharge rate when connected to the grid. During the simulation, the simulator will decide whether the PHEV departs on a drive according to a profile determined by a composition of probability distributions.

For instance, assuming a profile composed of a single Gaussian distribution, the probability at which the PHEV leaves for work at any given time, would be given by:

$$f(x, \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

where σ^2 is the variance, and μ is the mean time when a PHEV will leave for work. The mean value for each PHEV should be given, respectively by the time at which the PHEV needs to start its journey so that it arrives at work in time, as well as the time at which the PHEV finishes work. An illustration of this can be seen in figure 5.1, where a time-line corresponding to 24 hours is illustrated with 2 Gaussian distributions centered around peak hours $\mu_1 = 8, \mu_2 = 16$, and with standard deviations of $\sigma = 0.25$. Since time is modeled discretely in units of 15 minutes in the simulator, the probability that a PHEV departs between 07.45 and 08.00 would, for instance, be $\sim 34.13\%$.

Figure 5.1: Probability distribution of time of departure for a PHEV over a 24 hour period, showing two Gaussian distributions with mean values 07:00 and 18:00

Additionally, because of the nature of the Gaussian normal probability distribution, the probability that a PHEV might leave for work will be evenly distributed around the mean. However, this might not accurately reflect the behavior of a population of PHEV users. For instance, some individuals might be consistently earlier for work than others, as well as vice versa. To account for more varied behavior amongst individuals in the simulator, the PHEV models are based upon fixed user profiles which may contain different number of probability distributions. For instance, some individuals might be earlier for work than others, and as such, their probability distribution during early peak hours. To model this, we might consider a distribution that is more like a LogNormal distribution or a Weibull distribution with an early peak and a long tail, illustrated in Figures 5.2 and 5.3. Also, the choice of distribution is likely to depend on the task to be performed as well. For instance, an individual might be consistently early for work, which he or she might consider to be a more important task than, for instance, meeting a friend for watching a game of football. To account for many different types of behavior, the simulator and model is designed to support many different types of distributions, but this thesis will mainly concentrate on the Gaussian, LogNormal and Weibull distributions.

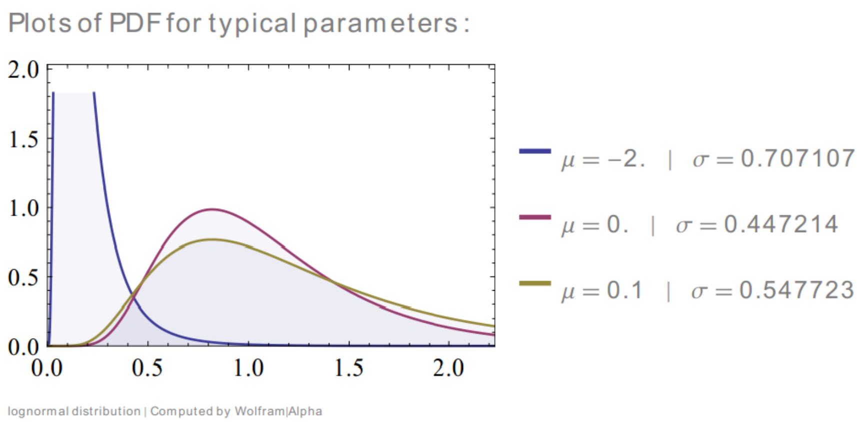

Figure 5.2: An illustration of a few different LogNormal distributions generated using Wolfram Alpha [Wolfram Alpha, 2012]

5.3.1 PHEV profiles

While it is reasonable to believe that the probability distributions will vary from person to person, it also seems reasonable to believe that the probability distributions will vary depending on the time of the day, within the daily patterns for an individual. For instance, a person should be more likely to arrive early at work, or at least somewhat on time, rather than arriving very late, very often. However, while this may be true for work, the probability distributions may look very different for other regular activities. To account for this, the probability distributions for the PHEV model in the simulator may be composed of several different types of probability distributions over a single day. These combinations of probability distributions for a single PHEV make up what will be called a *PHEV profile*, and

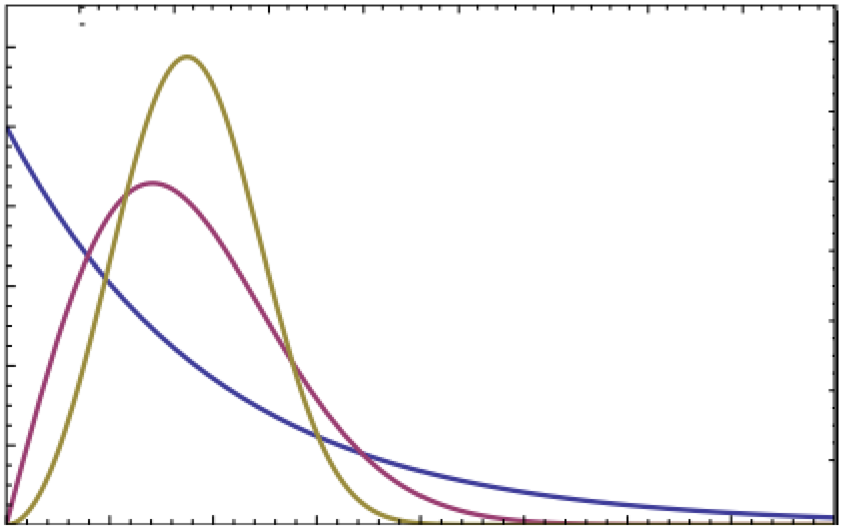

Figure 5.3: An illustration of a few different Weibull distributions generated using Wolfram Alpha [Wolfram Alpha, 2012]

each PHEV model in the simulations are generated from such a profile. Which profile that is used to generate the PHEV model is pre-defined in a configuration file that contains information about the configuration of the simulated power grid. Pseudo-code for how the simulator decides when a PHEV leaves, based upon the generated probability distributions, can be found in Algorithm 4.

5.4 Day-ahead model

When determining the day-ahead profile, it is preferable that the day-ahead profile will be as accurate as possible with respect to the real-time energy consumption. In real-life, historical data is often used to determine the day-ahead profile. In the simulator, a two-step process is used in the simulator. First, the baseline is determined as the expected contribution from all of the non-schedulable loads. For the experiments, the non-schedulable loads will be defined by the PowerNode models based on the historical data that was sampled from residential houses with smart meters.

After establishing the baseline, the expected PHEV demand need to be added to the day-ahead profile. To do this, it is necessary to know how much needs to be added. However, since no historical data is available for the PHEV demand, and since the PHEVs are modeled by the simulator, there are two different options that can be used. One option estimates the PHEV demand by running the simulation in advance, recording the demand by the PHEVs as the simulation progresses and without using any scheduling mechanisms. This option will be further discussed in Section 5.4.1. The other option calculates the expected demand from the PHEVs

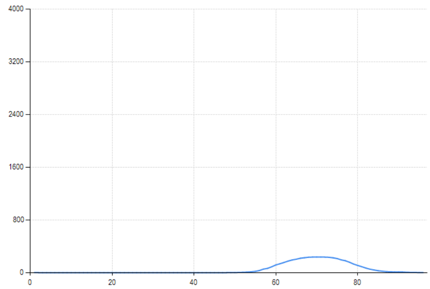

(a) Simulated values

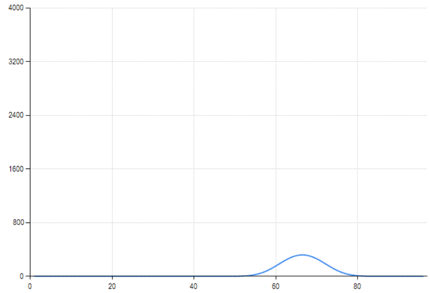

(b) Expected values

Figure 5.4: Comparison between simulated and expected values for 616 PHEVs, for a single, Gaussian probability distribution with mean 12:15, and a standard deviation of 60 minutes.

statistically by using known values for the probabilities that a PHEV would leave at a given time, plus the duration of the travel during which it would be discharging its batteries, the average rate with which it would be discharging, and the rates with which it would recharge upon return. Further discussion of this option can be found in Section 5.4.2.

Again, after the baseline and the expected demand from the PHEVs is determined there are different options available. One option is to add them together, superimposing them. This would create an accurate prediction of the demand profile when no scheduling mechanisms are used. However, considering that we also desire to influence the reduction of peak-loads, there are two other options available in the simulator for adding the PHEV predictions and the baseline together. In one, the baseline and the PHEV predictions are added together, after which one of the two algorithms is used. One algorithm, Algorithm 2, is a peak-shaving algorithm, loosely based on the Kohonen algorithm for self-organizing maps, and will be discussed further in Section 5.4.3. The other is an algorithm, Algorithm 3, where the PHEV demand is distributed onto the day-ahead portfolio based upon a distance rule. This algorithm will be discussed in Section 5.4.4. Common to both algorithms is that they shift loads away from peak-hours, while the overall volume of the portfolio is preserved. Both algorithms result in a portfolio that can be used to schedule PHEV loads to help reduce peak-load.

A comparison between simulated and expected values can be seen in Figure 5.4 showing the demand from running the simulator with 616 PHEV agents for a selected day in Figure 5.4a, as well as the expected demand calculated using the expected values based on the PHEV profiles in Figure 5.4b. For simplicity, all the PHEV profiles used for the comparison contained a single Guassian distribution. The Simulated values showed a consumption of 4734.62 kWh by the PHEV agents for that day, while the Expected values predicted a demand of 4235.0 kWh, a difference of 11,7%. However, as will be discussed later in Section 7, it is preferable that the expected values predict a slightly higher consumption than the actual consumption.

5.4.1 Simulated values

In using the simulated values, the method for calculating the day-ahead profile is the result of running the simulations for the desired number of days, without scheduling. This means that the contribution of the PHEVs will depend on the composition of profiles among the PHEVs modeled. Or in other words, the combination of probabilities that the PHEV might depart at any given time, as well as the duration of departure. These are the only influencing factors in determining the amount of power that a PHEV will need to charge when returning.

The apparent advantage in using this method for calculating the day-ahead profile, would be to obtain a fairly accurate day-ahead profile, for any day in the simulation. Using simulated values, the day-ahead portfolio will very closely approximate a load profile in which no scheduler is employed. The most obvious disadvantage in using simulated values, is the overhead of running the simulations twice, once for obtaining a day-ahead portfolio, and twice for running the simulations with the day-ahead portfolio obtained.

5.4.2 Expected values

Another approach to calculating the PHEV demand, is simply: If the load profiles of the power nodes are given, then an expected aggregated load profile can be calculated using the probability distributions, profiles, and average recharging/discharging rate of the PHEV models. In other words: Each PHEV profile has a set of probability distributions, and each probability distribution has an associated duration. This means that the expected load of a PHEV at a time t depends only on the probability that a PHEV is arriving at that time, the amount of power discharged since it left, as well as the recharging rate of the PHEV. Assuming that the PHEV will *fully recharge its batteries upon return*, then the aggregated expected load contribution from a singular probability distribution, is the probability that it left some time d ago, multiplied by the product of a duration vector and a scalar recharging rate. For instance, if we use a normal Gaussian distribution, and assuming that the PHEV has fully discharged its batteries, then the aggregated expected demand of that distribution to the overall expected load profile is 1 (the sum of the area under the normal distribution) multiplied by the capacity of the battery.

For determining the expected demand at some specific time t, we need to consider the contribution of the expected demand from any of the probability distributions in the profile. This means that the expected load for a single distribution, at some *specific* time t is the product of the probability that it was discharging at a time d earlier, and its (constant) recharging rate, assuming it has not recharged its batteries yet. The aggregated expected load contribution for the entire profile then becomes the aggregated contribution of the demand caused by all the distributions in the profile. Furthermore, the aggregated expected load demand from all the

PHEV models then, similarly, is the aggregated sum of all the profiles for all the PHEV models.

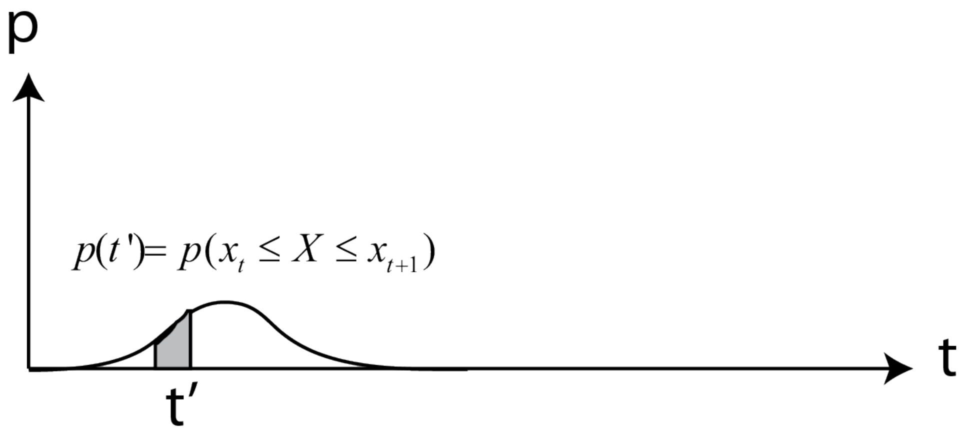

Figure 5.5: A sample probability distribution for a profile i. The marked area shows the probability that a PHEV leaves during a given 15 minute interval.

The process for calculating the expected PHEV demand is also illustrated in Algorithm 1. $profiles$ is the set of *probability profiles*, where a probability profile, Φ is represented as a vector containing the parameters for a probability distribution, the duration of the trip, how much of that trip is spent driving, the average discharge rate for that trip, and also the average charging rate for the PHEV. The average charging rate is constant over all probability distributions for a single PHEV.

First, the total amount of discharged power is calculated, iterating over the total time of the duration, accumulating the average discharge rate when PHEV is driving, given by the function $\kappa(\Phi, t')$. Second, for each time step t, we calculate the expected demand for time t plus the duration of the trip forward in time, until the accumulated discharged power is depleted. λ is a step.function that takes as parameters the probability profile, Φ and a time t, and returns the probability that the PHEV departs at that time. ξ is the expected demand from a particular *probability profile*, while the expected demand caused by the entire *PHEV profile* is given by the set of expected demand from all the probability profiles, Ξ.

After having calculated the PHEV demand, either by using simulated- or expected values, there are a couple of different options as to how this is added to the day-ahead portfollio. For this, we have considered further two approaches. One, where the PHEV demand is distributed using a distance rule algorithm, and one where the PHEV demand is superimposed on the PowerNode predictions, after which a peak-shaving algorithm is applied.

Algorithm 1 The algorithm for calculating the expected PHEV demand

for each $\Phi \in profiles$ **do**
  $\chi \leftarrow 0$
  for each $t' \in 0 \ldots \Phi.duration$ **do**
    if $\kappa(\Phi, t')$ **then**
      $\chi \leftarrow \chi + \Phi.drate$
    end if
  end for
  for each $t \in T$ **do**
    for each $t' \in t \ldots \Phi.duration$ **do**
      if $\chi > 0$ **then**
        $\tau \leftarrow \Phi.crate * \lambda(\Phi, t)$
        $\xi[t + t'] \leftarrow \xi[t + t'] + \tau$
        $\chi \leftarrow \chi - \tau$
      end if
    end for
  end for
  $\Xi \leftarrow append(\Xi, sort(\xi))$
end for

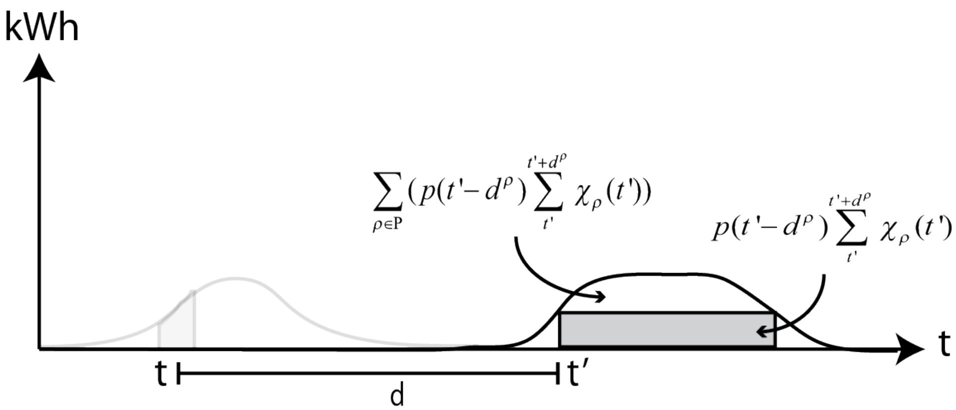

Figure 5.6: Expected load distribution caused by profile i. This figure shows how the expected load is calculated from the probability distribution in Figure 5.5. The expected load at time t' is the product of the weighted probability that it left some time d ago times the recharging rate of the PHEV, added to the expected load that is caused by earlier probable events.

5.4.3 Peak-shaving algorithm

After having calculated the predicted demand for all the PHEVs, the next step in the process is to determine how to distribute this across the load profile for all the power nodes. By superimposing the PHEV demand onto the predicted PowerNode demand, the peak-shaving algorithm can be used. The idea is that by shifting the PHEV demand away from peak-hours, the day-ahead imbalances can be best capitalized by the PHEVs early on, helping the PHEVs ability to charge as early as possible.

Algorithm 2 Processing the day-ahead profile using the peak-shaving algorithm. D is the dayahead profile, which contains the energy traded for on the dayahead market for each 15 min interval for the specific day. $\vec{w_i}$ is the maximum energy during the day, or the peak energy. S is the energy pool. Initially, the energy pool contains the energy shaved from peak hours. After peak-shaving, the energy in S is distributed to nearby hours until S is empty.

$\vec{w_i} \leftarrow \arg\max D, k \leftarrow i+1$
$\vec{x_i} \leftarrow \arg avg(D)$

let $\delta(\vec{w_j}) = \alpha\theta(\vec{x_i} - \vec{w_j})$

function $update(\vec{w_j})$
  $\psi \leftarrow \delta(\vec{w_j})$
  $\vec{w_i} \leftarrow \vec{w_i} + \psi$
  $S \leftarrow S + \psi$
end function

while $(\vec{w_{i\pm(k+1)}} + \delta(\vec{w_{i\pm(k+1)}})) > \vec{w_{i\pm k}}$ **do**
  $update(\vec{w_{i\pm k}})$

  $k \leftarrow k+1$
end while

while $S < 0$ **do**
  $update(\vec{w_{i+k}})$

  $k \leftarrow k+1$
end while

$\phi \leftarrow \frac{S}{2k}$

for $j \in i-k \ldots i+k$ **do**
  $\vec{w_j} \leftarrow \vec{w_j} - \frac{S}{2k+1}$
end for

In algorithm 2, the peak-shaving algorithm is applied to the simulated day-ahead profile. It starts out by finding the maximum peak in the day-ahead profile, D, $\vec{w_i}$. From there, the algorithm explore the energy profile for nearby times, moving the predictions closer to average and keeping track of how much. This

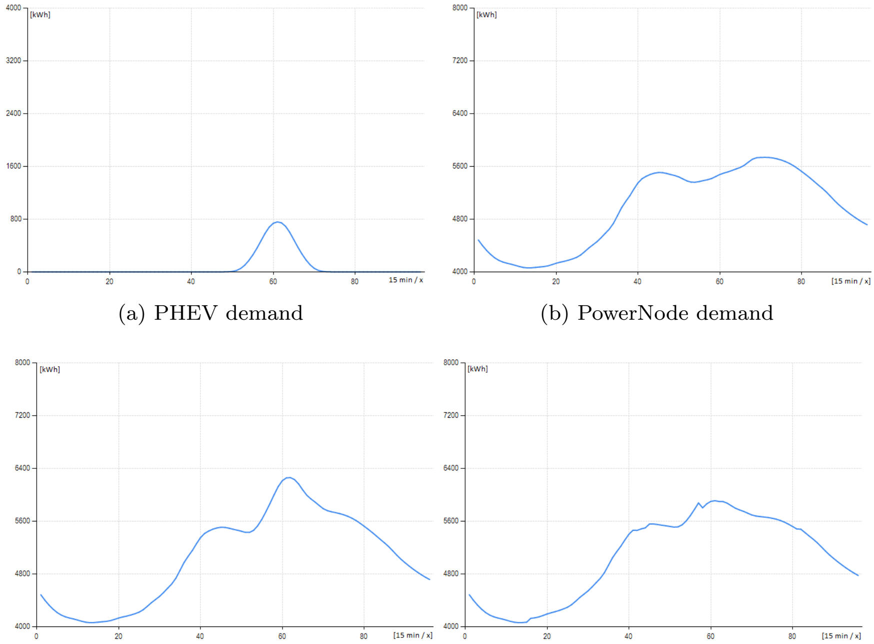

(a) PHEV demand

(b) PowerNode demand

(c) PHEV + PowerNode

(d) PHEV + PowerNode – after peak-shaving

Figure 5.7: Sample images from the simulator showing the different steps for calculating the day-ahead profile using the peak-shaving algorithm. Figure 5.7a shows the calculated expected PHEV demand, while the Figure 5.7b shows the PowerNode demand predictions. The aggregated demand of the PHEV and PowerNode is shown in Figure 5.7c. Finally, Figure 5.7d shows how the peak-shaving algorithm has been applied to the aggregated demand from Figure 5.7c.

process is repeated until the adjusted demand of the next time-step would have been greater than the adjusted demand of the current iteration. This process is illustrated in Figure 5.8. So far, the blue colored area has been shaved from the peak hours, accumulating the area in S meanwhile. This is the first part of the algorithm. In the second part, the accumulated energy pool S is distributed forward in time. It is important to note this difference, as it is the only difference between the first and the second part of the algorithm, except for the conditions with which they terminate. The same process is then repeated until the accumulated energy pool, S, is empty. Examples of the algorithm in action can also be seen in Figure 5.7, where the different steps of the entire process for calculating the day-ahead portfolio using the peak-shaving algorithm is discussed.

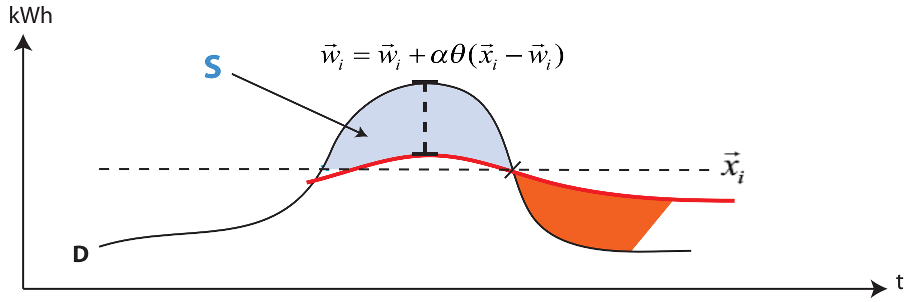

Figure 5.8: A graphical illustration of the effects of Algorithm 2. The initial update is shown by adjusting the peak by the update rule $\vec{w_i} \leftarrow \vec{w_i} + \alpha\theta(\vec{x_i} - \vec{w_i})$. The blue colored area represent the total area that will be shaved from the peak, accumulated in the energy pool S. The first part of the algorithm will terminate in the intersection between the blue and red colored areas, from where the shaved area will be subsequently distributed forward in time, until the energy pool S is depleted.

5.4.4 Distance-rule algorithm

The first method for distributing the expected contribution is simply to place the contribution at a time t which maximizes the function

$$f(x) = \frac{\theta^{|x-x'|}}{dayahead_{x'}}$$

where $\theta^{|x-x'|}$ is the distance rule away from the original position x', and $dayahead_{x'}$ is the value of the load profile of the power nodes at some time t. An illustration of the algorithm in action can be seen in Figure 5.9, where the initial expected load from all of the PHEVs are shown in the top left corner and the expected load of all the power nodes in the bottom left corner. In the top right corner, you can see the remaining load from the expected load of the PHEVs, where the blue area represents the amount of energy that has been distributed at the current time t, while in the bottom right corner you can see where the distributed energy from the top right graph has been placed. Pseudocode for the algorithm can also be found in Algorithm 3, where $vv(x,' t)$ is the distance rule function. The algorithm iterates over the expected PHEV demand, adding the demand to the time in the dayahead portfolio that maximizes the value of the distance rule.

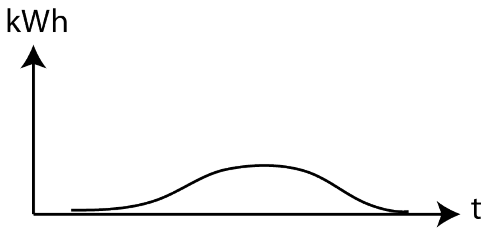

(a) Expected PHEV

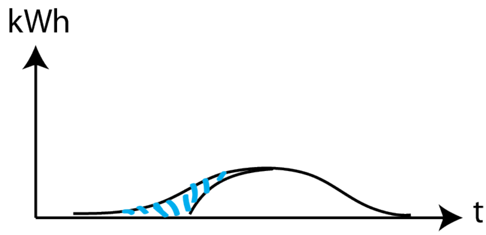

(b) Expected PHEV demand remaining

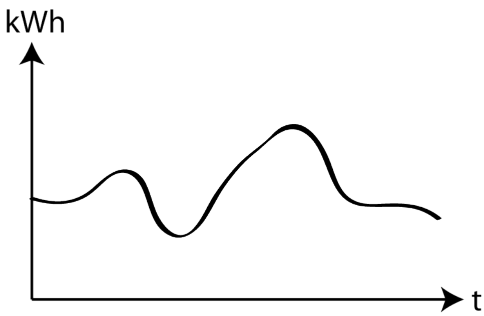

(c) PowerNode

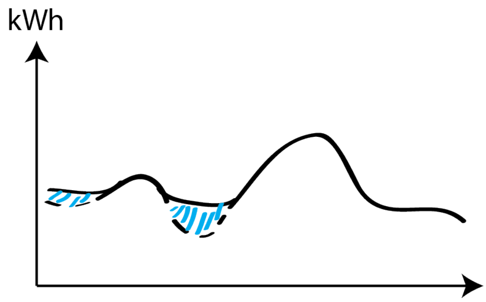

(d) PowerNode + PHEV demand distributed so far

Figure 5.9: A graphical illustration of Algorithm 3, showing how the expected PHEV demand is distributed onto the PowerNode predictions based on the distance-rule algorithm. Figure 5.9a shows the expected PHEV demand that is to be distributed, while Figure 5.9c shows the PowerNode predictions. In Figure 5.9c, some of the demand from Figure 5.9a has been distributed. This is illustrated by the blue colored area, and Figure 5.9d illustrates how this blue colored area has been distributed onto the PowerNode predictions based on the distance-rule algorithm.

Algorithm 3 Distribute expected load from PHEVs onto to expected load contribution of power nodes

let $\upsilon(x, x') = \frac{\theta^{|x-x'|}}{dayahead_{x'}}$

for each $i \in 0 \ldots len(expected)$ **do**

  let $demand = expected_i$

  while $demand > 0$ **do**

    let $\upsilon'(x) = \upsilon(x, i)$

    $i \leftarrow \arg\max_{\upsilon'(x)} dayahead$

    $dayahead_i \leftarrow dayahead_i - rate_{min}$

    $demand \leftarrow demand - rate_{min}$

  end while

end for

6 Multi-agent system

In the previous Section we discussed the different underlying models that the simulator and agents depend on. We discussed how the day-ahead portfolio was calculated, which is used by the BRP agent, and how this could be calculated by using a combination of simulated or expected values based on statistical tools. Further we discussed how the PHEV models were created. The PHEV models represents the mathematical model of the consumer, and determines when and how long the PHEV is used. And we also discussed how the constraints of the Transformer model are determined, and how the PowerNode models are generated based on historical data that is used as input to the simulator.

All the different models that we discussed in the previous Section is the foundation for the environment in which we intend to implement the multi-agent system and the different mechanisms that are used for scheduling the PHEV charging plans. In the rest of this Section, we will discuss each of the individual multi-agent scheduling mechanisms, as well as each of the individual agents and their responsibilities. Of the different mechanisms, two of them are based on a centralized load scheduling mechanism, and further two are based on a decentralized mechanism where the agents either choose their charging plans based on a uniform distribution, or a mixed strategy distribution based on a weighted distribution on the future energy demand in the grid. We will first discuss the individual agents and the message protocol they use, after which the four different mechanisms will be discussed, explaining the load scheduling algorithms for the centralized mechanisms, and how the PHEV agents choose their own charging plans in the decentralized mechanisms.

6.1 PowerNode agent

The PowerNode agent in the simulator have only one main responsibility, which is to inform its parent Transformer agent about its demand for the next 15 minutes. The demand is given by the PowerNode model that was generated from the historical data that was discussed in Section 5.1. The PowerNode agent knows its demand by observing changes to its PowerNode model as time moves forward in the simulator.

6.2 PHEV agent

In both mechanisms, the PHEV agent have two main responsibilities. Its primary responsibility is to maximize the battery levels within time of departure. Its secondary responsibility is to achieve its primary responsibility in the most socially economical way possible. How it attempts to achieve these goals, however, depends upon which mechanism is used by agents. In the decentralized mechanisms, the PHEV agent chooses its charging strategies by itself, while in the centralized mechanism it defers control of creating the charging plan to a centralized scheduler. A sample interaction between a PHEV agent, Transformer agent and BRP agent in the centralized mechanism can be seen in Figure 6.2, where the PHEV agent is negotiating with the other agents in creating a charging plan that will maximize the social welfare of all the agents. Sample interactions of the decentralized mechanisms can also be seen in Figure 6.5 for the Uniform strategy scenario, and 6.7 for the Mixed strategy scenario. In the Random strategy mechanism, the PHEV agent will choose its charging plan based upon a uniform distribution from the time the plan is created and until its expected time of departure, while similarly, in the Mixed strategy mechanism, the charging plan will be generated based upon a weighted distribution. The different mechanisms will be discussed in more detail in Sections 6.6 and 6.7.

6.2.1 PHEV learning

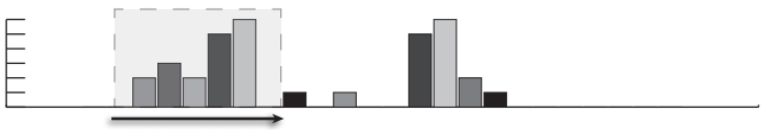

Figure 6.1: A sliding window over a histogram for a PHEV leaving

To estimate expected time to depart, the PHEV needs a way to collect information about its usage patterns. With this in mind, each time the PHEV is used, the starting time of its journey will be recorded, as well as the state of its battery at the time of departure. While they are related, this learning mechanism must not be confused with the PHEV *profiles* that determines when a PHEV model leaves on a drive. The PHEV agent has no direct knowledge of the PHEV profiles that determines the behavior of its model. This means that it will have to learn from experience the usage patterns of the model it is responsible for. Ultimately, this should mean that the probability distributions that are defined in the PHEV *profiles* will be reflected in the PHEV *knowledge base* (KB). The PHEV will use this KB to estimate an expected time to depart, which it will communicate to the BRP agent each time the PHEV connects to the grid. The BRP agent will use this information to find a good scheduling plan for all its agents. Additionally, the PHEV agent also records information about the state of its battery at time of departure. This can be used to measure the effectiveness of participating in the scheduling mechanism, from the point of view of the PHEV agent.

While using the pre-defined mean values from the probability distributions in the PHEV profiles could have been used to determine the expected time of departure, it is preferable that the simulator should be as realistic as possible. In reality, people are generally unpredictable, and PHEV units are unlikely to have direct knowledge about the usage patterns of a potential owner. To this end, the PHEV learning mechanism is used to reflect this uncertainty.

To determine the most likely time to departure from the KB of the agent, the PHEV agent applies a sliding window, of size k, to determine the mode of the data within that window of time. Where the window is defined from the present time,

Algorithm 4 Updating the state of PHEVs. This is the algorithm that decides whether a PHEV leaves home, thereby disconnecting from the grid.

```
let norm = normal(μ = 60, σ² = 15)

for each phev in phevs do
    if phev.status is connected then
        for each probdist in phev.profile do
            let prob = (probdist.cdf(now + 15 min) - probdist.cdf(now))
            let rand = (rand(0.0, 1.0))

            if rand < prob then
                leave(phev, tick, phev.battery)
            end if
        end for
    else
        if phev.status is arrived at destination then
            set(phev.status, connected)
        end if
    end if
end for
```

and until k units of time later. An illustration of this principle is shown in figure 6.1. By using a sliding window, we limit the probability of the pitfall that the PHEV agent converges on an expected time to depart that may, or may not be far off into the future. However, it also has the disadvantage that the performance of the PHEV agent in determining its time of departure will depend on the size of the window. For instance, let us assume a window of 5 hours is used. If, using this window size, the PHEV agent applies the window to its KB during night hours, then it is less likely to find any significant patterns than if the window size had been 8 hours which would likely have presented more valuable information. And disregarding the time of the day the window is applied to, it is always a risk that the just beyond the window, there might be more valuable information than what is currently present. However, for the purposes of the experiments that will be discussed in later Sections, this mechanism will perform adequately well that it can be used. Since we know the PHEV profiles that determines the usage patterns of the PHEVs, we can select a window size which will perform adequately well.

6.3 Transformer agent

The transformer agent has one main responsibility, to ensure that the capacity of the physical transformer it is assigned to is never exceeded. To enforce this constraint, it has the capability to filter demand messages made by PHEV agents. After the PHEV agent have negotiated a charging plan with the BRP agent, or if the PHEV agent have already done so in a previous timeframe, the PHEV agent will send a Demand message containing their energy demand for the current time-frame to the Transformer agent. The Transformer agent will collect these Demand messages from all of its connected PHEV agent and PowerNode agents. If the sum of the demand of all the Demand messages exceeds its capacity, then it will filter away Demand messages that are coming from PHEV agents until its capacity constraint is satisified. While it could also, theoretically, filter Demand messages from agents other than PHEV agents also, it is not safe to assume that those messages represent deferable loads.

6.4 Balancing responsible party

Depending on which mechanism is used, the BRP agent has one of two responsibilities. In the centralized mechanisms, the BRP agent is responsible for scheduling the PHEV charging plans for all the PHEV agents in the grid. More detail about how the charging plans are scheduled in the centralized mechanisms can be read about in Section 6.6.1. If any of the two decentralized mechanisms is used instead, then the BRP agent is responsible for providing its predictions about the future energy demand to the PHEV agents.

6.5 Message protocol

The message protocol used for this thesis is a custom-designed protocol built upon discriminated unions, which are supported in the F# language [2]. A discriminated union is a type of data structure that can hold values which may be of different, fixed types, depending upon the tag. This means that the *Message* type used by the agents in the simulator, can take different values depending upon its tag. For instance, a message can be tagged as being a *Charge*-message, where the value is an n-ary tuple that takes two floating points and an integer, or it could be tagged as a *Demand*-message, taking on a value of a single floating point. By using discriminated unions as messages, this allows for type-safe message passing in the simulator.

In this case, the message protocol is a double-layered message structure, where the first layer includes information about recipient, sender, contents of the message, and optionally, the reply channel if the sender desires a response to the message. A list of the most important messages can be seen in Table 6.1. The full list of messages available to the agents can be found in Appendix C.1.

6.6 Centralized mechanisms

In the centralized mechanisms, the PHEV agent coordinates its charging plan by interacting with the BRP agent and the Transformer agent. The BRP agent communicates with its Transformer agent, which it can announce its intention to charge to. In turn, the Transformer agent which is responsible for keeping peak load within the constraint of its transformer, forwards the aggregated intentions of all its PHEV

[2] Also called tagged unions in Computer Science

Type	Contents	Description
Charge	U_χ - uncharged capacity *ttl* - time to departure *rate* - charging rate	Used by the PHEV agents to announce their charge intentions
Demand	E_d - current demand by the agent	Used by the PHEV agents and PowerNode agents to announce their current demand to the Transformer agent
Consume	E_c - accepted demand to agent	Used by the Transformer agents to inform the PHEV agents and PowerNode agents how much they can consume at this time
Request Predictions	-	Used by the PHEV agents in the Mixed strategy mechanism to request a window of predictions about the future energy demand of all the PowerNodes
Request Dayahead	-	Used by the simulator to request the dayahead model from the BRP agent
Predictions	$\vec{p}_t$ - predictions up to time t	Used by the BRP agent when responding to RequestPredictions from the PHEV agents
Strategy	$\vec{\chi}_t$ - charging strategy until expected time to departure	Used by the PHEV agents in the Mixed strategy mechanism when informing the BRP agent about their final charging plan

Table 6.1: A selection of the most important messages available to the agents in the simulator

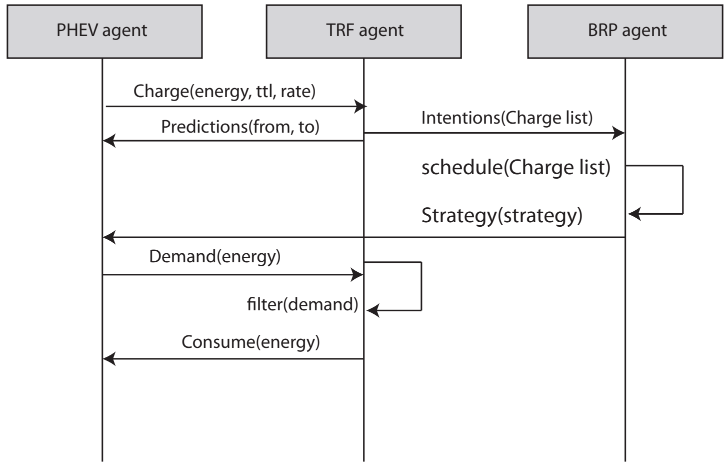

Figure 6.2: An example of multi-agent interaction in the centralized scenario

agents to the BRP agent which performs load scheduling on the intentions. After having scheduled the load profiles, the BRP agent announces its charging plans back to the Transformer agent. Depending on whether the constraints of the transformer is satisfied, the Transformer agent will either request a new plan from the BRP agent or forward the message back to the intended PHEV agents. An example of such an interaction is shown in figure 6.2.

6.6.1 Scheduling algorithms

How the PHEV agents charging plans are scheduled in the centralized mechanisms depend upon which scheduling algorithm that is used. If the centralized scheduling mechanism is used, the BRP agent can choose between two scheduling algorithms: A reactive scheduling algorithm and a proactive scheduling algorithm Boucké and Holvoet [2011]. The goal of both algorithms is to schedule the charging plans so that the overall difference between the dayahead predictions and the realtime energy demand is as low as possible.

Reactive scheduling

The principal idea behind the reactive scheduling algorithm is to schedule the PHEV demand so that the imbalances between the day-ahead portfolio and the real-time consumption is minimized as early as possible. In figure 5, pseudo-code for the proactive scheduling algorithm is shown, showing how energy is reserved for as long as the prediction of a given time is less than the dayahead consumption for the same timeframe, while there is still PHEV demand left to assign. Notice that any remaining demand is not scheduled if the dayahead portfolio imbalances does not admit it. While this will likely reduce the performance of the algorithm in charging the PHEV batteries, one of the primary intentions with the centralized scheduling mechanism is to use it to achieve a peak-shaving effect. Therefore, we want the algorithm to adhere to the constraints of the dayahead portfolio imbalances as closely as possible. The effects of this scheduling profile is shown in figure 6.3.

Algorithm 5 Pseudocode for a reactive scheduling strategy, adapted from the algorithm described by Boucké and Holvoet [2011]

function CREATE-PLAN
  $energyLeft \leftarrow sum(intentions)$

  for each $t \in now \dots ttd$ **do**
    if $sum(prediction) < sum(dayahead) \& energyLeft > 0$ **then**
      $plan_t \leftarrow \chi.rate$
      $prediction_t \leftarrow prediction_t + \chi.rate$
      $energyLeft \leftarrow energyLeft - \chi.rate$
    end if
  end for

  return plan
end function

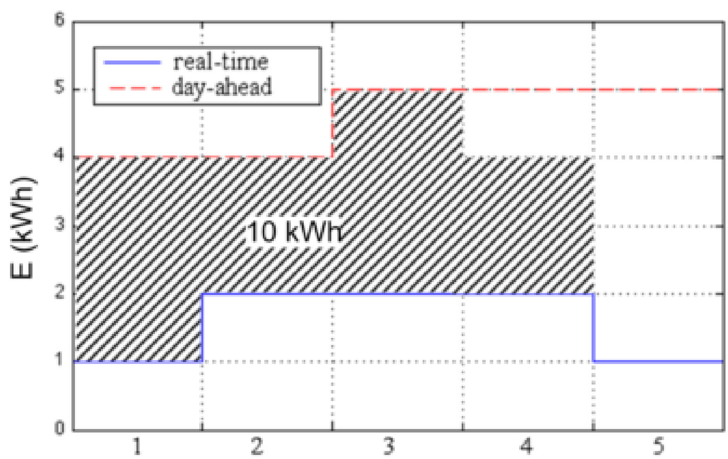

Figure 6.3: An illustration of the effects of a reactive scheduling profile [Boucké and Holvoet, 2011]

Proactive scheduling

Algorithm 6 Pseudocode for a proactive scheduling strategy [Boucké and Holvoet, 2011], adapted from the algorithm described by Boucké and Holvoet [2011]

function CREATE-PLAN
    $energyLeft \leftarrow sum(intentions)$

    while $energyLeft > 0$ **do**
        if $sum(dayahead) - sum(prediction) > 0$ **then**
            $t \leftarrow \arg\max_t(prediction(t) - dayahead(t))$
            $plan_t \leftarrow \chi.rate$
            $prediction_t \leftarrow prediction_t + \chi.rate$
            $energyLeft \leftarrow energyLeft - \chi.rate$
        end if
    end while

    return plan
end function

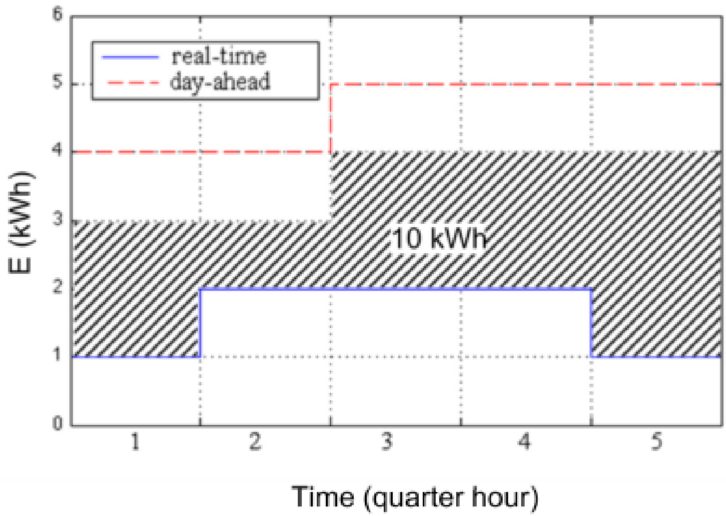

Figure 6.4: An illustration of the effects of a proactive scheduling profile [Boucké and Holvoet, 2011]

In this profile, energy is distributed in a way such that demand is assigned to times when the imbalance between the day-ahead portfolio and real-time consumption is the greatest. This is done while there is a positive difference in the day-ahead portfolio compared to the predicted real-time consumption. The predictions in this case, refers to the energy demand predictions from the PowerNodes, and does not include predictions about the demand of the PHEVs, initially, but the BRP agent will update its predictions with the PHEV demand as it generates the charging plans. Otherwise, energy is assigned to the times at which the imbalance is smallest. This is intuitively because it is desirable to minimize the average distance between the day-ahead quantity and the real-time predictions. By assigning energy to the time of largest imbalance while the difference is positive, actual real-time consumption is brought closer to the day-ahead quantity. If the difference is negative this means that the overall consumption has exceeded the net day-ahead quantity. This means that wherever the charge is placed, it will have a negative impact, but it will do the least harm at the time at which the imbalance is the least. This is illustrated in the pseudo-code shown in figure 6, and the effects are illustrated in 6.4. Also in this algorithm will the scheduler strictly adhere to the constraints of the day-ahead portfolio imbalances, similar to the Reactive scheduling mechanism.

6.7 Decentralized mechanisms

In the decentralized mechanisms, the idea is that a peak-shaving effect can be obtained by letting PHEVs choose charging strategy on their own. They do this by randomizing over a probability distribution, where the probability at a given time determines the likelihood that a PHEV chooses to charge during that time. This process is illustrated by pseudo-code in Algorithm 7. The important thing to consider with this mechanism is how the strategy for generating the charging plans is decided. Each agent is, in principle, free to decide whichever strategy they believe will benefit them the most. Assuming that the agents are interested in maintaining the stability of the grid, reducing maximum peak and PAR, we have considered two different strategies in this thesis: A Uniform strategy and a Mixed strategy which are discussed in Sections 6.7.1 and 6.7.2.

Algorithm 7 Randomizing over a distribution

let remaining = battery.max - battery.current

let $plan$ = [..]

for $t \in$ now .. ttd **do**
    if rand(0,1) $<$ $strategy_t$ **then**
        $plan_t \leftarrow$ rate
        remaining $\leftarrow$ remaining - rate
    end if
end for

6.7.1 Uniform strategy

With the uniform strategy, the PHEV agents will generate a charging strategy, assigning a uniform distribution with a probability of $\frac{1}{2}$ that it will charge in each of the 15 minutes leading up to the expected time of departure determined from the PHEV learning mechanism that was discussed in Section 6.2.1. If the KB of the PHEV is insufficient to determine an expected time of departure, it will default to the size of its learning window. For instance, if a learning window of 40 times 15 minutes is used, but it has no previously recorded departures for that window, then it will default to generating a charging plan for the next 10 hours. Ideally, this should lead to a uniform distribution of the PHEV demand over the duration of a day. While the assumption is that this will aggravate demand during peak-hours somewhat, the idea is to investigate by how much the problem is aggravated and compare it to see whether this strategy is better than a baseline scenario where no scheduling is used at all. Because of the simplicity of this mechanism, this strategy has an advantage over centralized mechanisms concerning fault-tolerance.

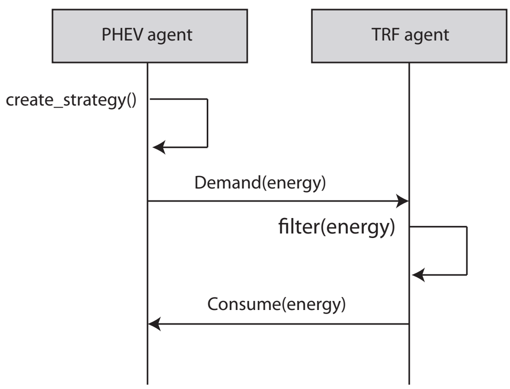

Figure 6.5: An example of multi-agent interaction in the uniform strategy scenario

6.7.2 Mixed strategy

In the Mixed strategy scenario, the distribution will be based upon predictions about the future energy demand, weighting the probabilities so that likelihood of charging during times of low demand increases compared to times of high demand. The predictions that is used to generate the weighted probability distribution come

Algorithm 8 Generating a mixed strategy charging plan

$ttd \leftarrow \arg\max_t(window \in KB)$

let predictions = ask(**BRP**, Predictions(ttd))

let max = $\arg\max predictions$
let min = $\arg\min predictions$

for each $t \in now \ldots now + window$ **do**
    let $x = predictions_t$
    $strategy_t \leftarrow (x - min)/(max - min)$
end for

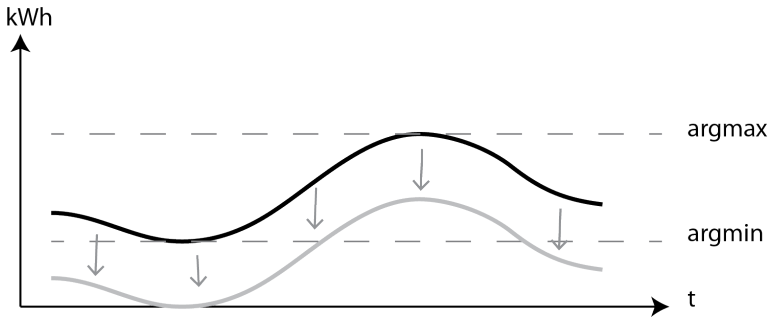

(a) Scaling the predictions by $x/\arg\max$

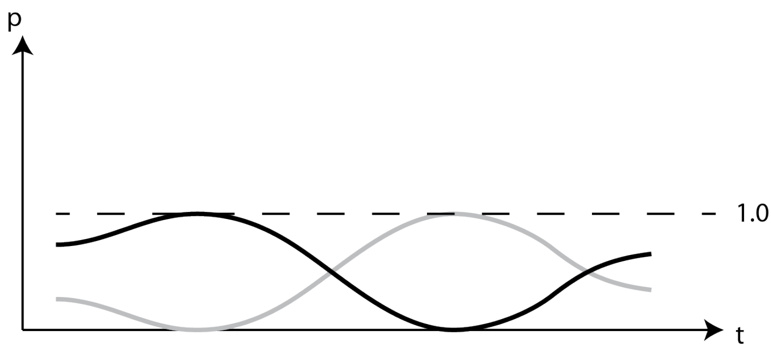

(b) Inverting the scaled predictions

Figure 6.6: An illustration of how the mixed strategy is generated. Figure 6.6a shows how the demand graph is scaled by $x - \arg\max$, and then normalized by $\frac{x}{\arg\max}$. The scaling is done to aggravate the differences between minimum demand and peak demand, so that the full range of probabilities from 0% to 100% is used. Subsequently, Figure 6.6b shows how the scaled and normalized predictions from Figure 6.6a are inverted by $1 - p$

from the BRP agent. Each time the PHEV agent recognizes that it is connected to the grid and that its current battery capacity is below maximum capacity, it will ask the BRP agent for its predictions about the future energy consumption based on a window from the present time and until the PHEV agent believes it will disconnect from the grid again. And similar to the Uniform strategy mechanism, it will default to the size of its learning window if it is unable to determine an expected time of departure from its KB. It is important to remember, however, that in the Uniform strategy mechanism, the expected time of departure is used to generate the charging plan directly, while in the Mixed strategy mechanism, the information is used to request a window of predictions about the future energy demand from the BRP.

The demand predictions that are generated by the BRP is *initially* determined from the historical data that was used for input to the PowerNode models. This means that the PHEV demand is initially not included in the predictions that the PHEV agents receive from the BRP. However, once the PHEV has generated its charging strategy, randomizing over the distribution, it reports the charging plan that it created back to the BRP agent. The BRP agent will then update its own predictions with the charging plan that it got from the PHEV agent. This prevents subsequent PHEV agents from generating their charging plans based on the same distribution as the previous PHEV agents.

The algorithm for generating the mixed strategy can be seen in algorithm 8. and Figure 6.6 shows a graphic illustration of the process. The end-result is a distribution that is an inverted distribution based on the normalized predictions about the future energy consumption. A typical interaction between the agents can be seen in figure 6.7

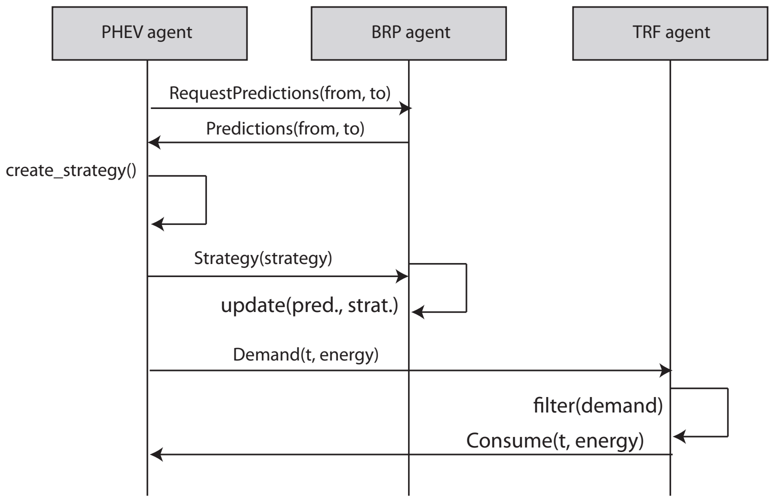

Figure 6.7: An example of multi-agent interaction in the mixed strategy scenario

7 Experiments

The experiments are designed to measure the performance of the different mechanisms with respect to the hypotheses presented in Section 3, targeting the fairness, stability and peak-shaving capabilities of the mechanisms.

Each experiment consists of running the simulations for the length of 100 days, where each combination of mechanism and day-ahead algorithm were used. This gives a total number of 18 different experiments, with 6 combinations of day-ahead algorithm and mechanisms, where each was tested with respect to stability, fairness and peak-shaving. The experiments will be performed on a grid composed of 1756 power nodes and 616 PHEVs. The number of power nodes come from the historical data that was used as a basis for the simulations, while the number of PHEVs is reflected from the predictions presented in Section 2, which state that by 2030 the ratio of PHEVs in Norway will be 26%. For good measure, this was rounded up to 30%. This assumes an average number of 1 vehicle per house. In addition, the time period for the sampled data that was used was from April 1st until July 10th, 2006.

7.1 Power Grid

For the purpose of this simulations, it was preferable to model a grid in which the capacity of the grid would be challenged under times of high demand. However, since the historical data used as a baseline for the power nodes included no information on the structure of the grid, it was decided to generate one computationally, in which the grid was deliberately designed so that it would challenge the maximum capacity during peak.

The grid was generated by finding the peak load of each of the power nodes from the historical data, and then assigning power nodes to each transformer until the sum of the peak loads exceeded the capacity of the transformers. After a low voltage transformer had been assigned enough power nodes, the low voltage transformer would, in turn, be assigned to a high voltage transformer by following the same principles. This process was iterated until the set of power nodes from the historical data were empty.

7.2 PHEV profiles

In figure 7.1, the different PHEV profiles used for the experiments can be seen. Figure 7.1a is based on the Commuter profile, with a Weibull- and LogNormal distribution centered around 07:15 and and 20:15. The different parameters for the distributions can be found in Table 7.1. Figure 7.1b has one Gaussian distribution centered around 07:15 and 20:15, with standard deviations of 30 minutes and 120 minutes. Figure 7.1c shows three Gaussians, centered around 12:15, 18:15 and 22:15, with standard deviations of 60, 120 and 120 minutes. The last figure, Figure 7.1d, shows a Weibull distribution at 07:15 with a shape parameter of 5.0 and a scale parameter of 14.0, and two Gaussian distributions centered around 18:15 and 22:15, with standard deviations of 30 and 60 minutes.

The profiles were designed so that the PHEV demand would roughly coincide with peak hours, putting extra pressure on the power grid, but also to be challenging for the individual PHEVs in being fully charged at time of departure. The Commuter-profile in Figure 7.1a is designed to have an earlier time of departure during morning hours than the other profiles, considering that a commuter likely has a longer distance to work than most others. In the afternoon hours, the Commuter-profile also has a LogNormal distribution, designed to simulate a regular behavior pattern for a typical Commuter.

For the City-profile in Figure 7.1b, the profile is designed to have a relatively short distance to work and can leave later for work than the other profiles. Since it also spends less time traveling, the overall duration that it is disconnected from the Grid is also shorter. This profile also has a LogNormal distribution in the afternoon hours, but with a longer tail than the Commuter-profile. The afternoon distribution was mainly added to challenge the battery scheduling of the mechanisms.

In the Homeworker-profile there are three Gaussian distributions. The idea is that since the Homeworker works from home there are no time-pressing considerations for a highly regular behavior. All the Gaussian distributions are designed to be somewhat overlapping with long tails. This is done to simulate a highly irregular pattern for this type of profile.

The last profile, Suburban, is designed to be a middle ground between the City-profile and the Commuter-profile with a shorter distance to work than the Commuter-profile, but longer than the City-profile. This profile also have a LogNormal distribution in the afternoon hours to put pressure on the recharging of the batteries. In addition, assuming that a Suburban household is likely to use a car more than a City household, there is an additional Gaussian distribution in the evening hours.

Also, the battery capacity and charging rate used for the PHEVs in the simulator were 16kWh, and 0.625kWh/tick respectively. This gives a total charging time for an empty battery of 6,4 hours. In addition, each PHEV has a discharge rate of 2.0 kWh/tick, which gives a total driving time of 2 hours.

7.3 Baseline and minimum extreme

To have a reference point to compare the results of the different experiments outlined in this section, we will also do experiments to use as point of reference when comparing the results: A baseline experiment which is established by running the simulator with no scheduling mechanisms or transformer filtering used, and a minimum extreme experiment which is established by running the simulator with no PHEVs present in the grid. The two experiments will represent two polar extremes, where the baseline is intended to be used as a worst-case scenario that the other mechanisms will attempt to outperform, and the minimum extreme experiment which will represent the best case scenario when comparing average maximum peak. The idea is that if that the mechanisms do worse than the baseline experiments, then it seems reasonable to assume that the mechanisms failed in achieving their goals and it would be better to not use scheduling at all. Conversely, if any of the mechanisms were to outperform the results of the minimum extreme exper-

Profiles							
type	γ	μ	σ	λ	k	dur.	driv.
City							
Weibull	31 (07:45)	-	-	5.0	14.0	36	4
LogNormal	69 (20:15)	0.0	1.467	-	-	6	4
Suburban							
Weibull	29 (07:15)	-	-	5.0	14.0	40	10
LogNormal	74 (19:30)	0.0	2.0	-	-	5	2
Gauss	81 (21:45)	81.0	4.0	-	-	6	4
Homeworker							
Gauss	49 (12:15)	49	4.0	-	-	11	6
Gauss	69 (18:15)	69	4.0	-	-	10	6
Gauss	85 (22:15)	85	4.0	-	-	10	6
Commuter							
Weibull	29 (07:15)	-	-	5.0	14.0	44	16
LogNormal	31 (20:15)	0.0	1.0	-	-	12	4

Table 7.1: Numerical values for the different profiles and probability distributions. γ is the location parameter, μ and σ is the mean and standard deviation, λ and k is the scale- and shape parameter, *dur.* is the duration that the PHEV is disconnected from the grid and *driv.* is the time spent driving.

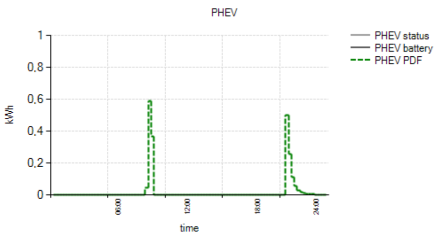

(a) Commuter

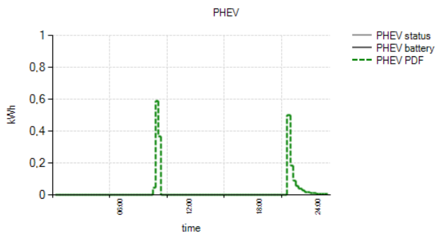

(b) City

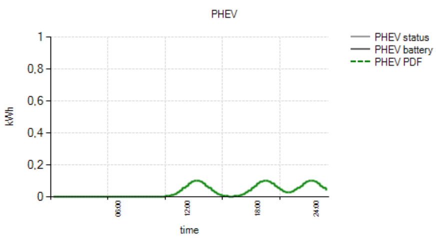

(c) Homeworker

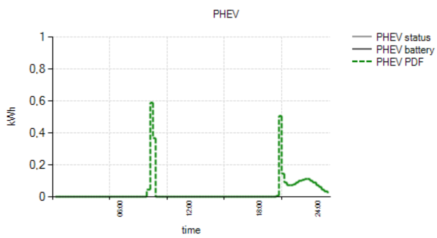

(d) Suburban

Figure 7.1: The different, discretized PHEV profiles used in the simulator. A graphical representation of the profiles and probability distributions found in Table 7.1

iment with respect to average maximum peak, then this would be an indication of a logical flaw with one or more of the mechanisms or models.

The results of these experiments can be found in Section 8.1, Experiment 8.1a.

7.4 Day-ahead profile

In this thesis we have examined 2 different options to calculating the PHEV contribution. Either by using simulated values, or calculating expected values using a statistical approach. Subsequently, there are 3 different options for the next step. Either the expected or simulated values can be added to the PowerNode contribution, after which peak-shaving is performed, or the PHEV contribution can be distributed onto the PowerNode contribution by distributing using a distance rule. One of the ideas behind calculating the PHEV demand by using expected values, was to save the overhead of using simulated values. In addition, and perhaps more importantly, it is also preferable that there is some error in the approximations, as this will make it easier to compare the performance of the different centralized mechanisms in later experiments. If the approximation of the PHEV demand is too accurate, or underestimated, then the difference between the proactive and the reactive mechanism will not be apparent. In general, the optimal result for comparing the centralized mechanisms will be if the PHEV demand is slightly overestimated.

To analyze the performance of the experiments, the simulator was run for 100 days with no scheduling mechanisms used, where the dayahead was calculated using both simulated and expected values. The experiment is designed to see how much the different dayahead profiles deviates from the total energy consumption. The results of the experiments can be found in the Section 8.2, experiments 8.2a and 8.2b.

The intention with these experiments is to verify that a day-ahead profile can be constructed by using expected values for calculating the PHEV demand. This will relate to 3.1 in the sense that it will verify whether this method of calculating the PHEV demand, using expected values, is a good approximation for when calculating the PHEV demand in later experiments.

7.5 Scheduling mechanism

Perhaps the most interesting experiments in this thesis will be the different scheduling mechanisms. In this thesis we have investigated four different approaches to load-scheduling PHEVs: Two algorithms where a centralized scheduler would schedule the charging plans of the PHEVs, and two algorithms where the PHEVs would decide upon their charging plans by themselves.

In these experiments, we will investigate hypotheses 3.1 and 3.2, of whether a centralized scheduling mechanism or decentralized scheduling strategies can help to reduce peak-load in the system. The simulator will first be run with no scheduling mechanism, where PHEVs charge whenever possible, after which experiments will be performed against the centralized scheduling algorithms and the decentralized scheduling strategies. We will observe how the scheduling mechanisms perform with respect to the average maximum peak and the peak-to-average ratio (PAR),

	Peak-shaving	Distance-rule	
Centralized (Proactive)	8.3a	8.3b	-
Centralized (Reactive)	8.3c	8.3d	-
Decentralized (Uniform)	-	-	8.3e
Decentralized (Mixed)	-	-	8.3f

Table 7.2: The different combinations of scheduling strategies. The cells refer to the experiments in Section 8.3, where the results of the experiments can be found.

	Peak-shaving	Distance-rule	
Centralized (Proactive)	8.4a	8.4b	-
Centralized (Reactive)	8.4c	8.4d	-
Decentralized (Uniform)	-	-	8.4e
Decentralized (Mixed)	-	-	8.4f

Table 7.3: The cells refer to the experiments in Section 8.4, where the results of the transformer experiments can be found.

defined by $_{peak}/L_{avg}$, where $_{peak}$ is the overall load during peak hours, and $_{avg}$ is the average load during a day

7.6 Transformer agent

These experiments will investigate hypothesis 3.3 that a centralized scheduling mechanism or decentralized scheduling strategy can help to reduce the problem of transformer capacity being exceeded. The experiments will be performed in the same manner as in Section 7.5, but the focus will be on number of times the peak-load will exceed the maximum capacity of any of the transformers in the grid. We will observe how the scheduling mechanisms perform with respect to average transformer capacity exceeded, and average PHEV charging demand filtered by the Transformer agent.

7.7 PHEV fairness

These experiments will investigate hypothesis 3.4 that a centralized scheduling mechanism or a decentralized scheduling strategy can be scheduled fairly. The

	Peak-shaving	Distance-rule	
Centralized (Proactive)	8.5a	8.5b	-
Centralized (Reactive)	8.5c	8.5d	-
Decentralized (Uniform)	-	-	8.5e
Decentralized (Mixed)	-	-	8.5f

Table 7.4: The cells refer to the experiments in Section 8.5, where the results of the PHEV fairness experiments can be found.

experiments will be performed in the same manner as in Sections 7.5 and 7.6. We will observe how the scheduling mechanisms perform with respect to the average battery capacity at time of departure for the PHEV agents.

8 Results

In this section, the results of the experiments described in Section 7 will be presented. Each experiment will be presented individually, highlighting some of the key figures. The results will be presented as is. For further discussion and comparison of the results, see Section 9.

8.1 Baseline and minimal extreme experiment

Experiment 8.1a *Baseline and minimum extreme*

The baseline and minimum extreme was established by running the simulator for 100 days, using no scheduling mechanisms or transformer filtering. Additionally, the minimal extreme experiment was established by running the simulator for 100 days in the same grid, but without PHEVs.

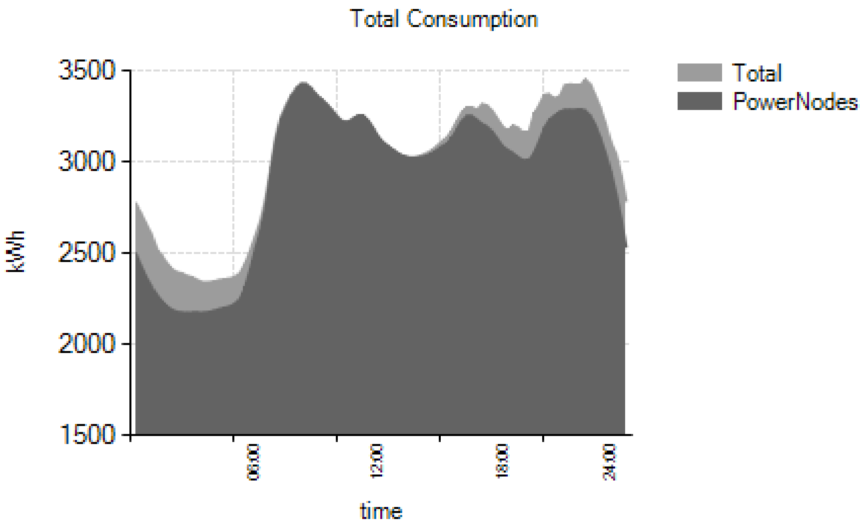

Figure 8.1: Sample image of the baseline from a selected day in the simulator. The baseline represents the scenario where no scheduling mechanisms are used. The dark gray area represent the consumption of the PowerNodes, while the light gray area represent the combined consumption of both the PowerNodes and the PHEVs.

The results of establishing the baseline can be seen in Table 8.1, where the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 15,2%, and a maximum of 22,6%. Average maximum peak during the simulation was 3.31 MWh, while the maximum daily peak was 5.38 MWh. Transformer

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.310*	5.388*	681.9	1.02
PAR	1.152*	1.226	0.026	0.69
		PHEV fairness		
PHEV(Ux)	3.762*	4.278*	183.5	0.10
PHEV(Avg)	0.776	0.341	1.0	6.15
		Transformer stability		
Trf(Exc)	939.4	24204.2	3430.0	5.57
Trf(Flt)	0.0	0.0	0.0	0.00
		Day-ahead imbalance		
Dayh.	0.000*	0.000*	0.000*	0.00
Dayh.(Dx)	276.238*	460.222*	58.312*	1.07
DDx/Tot.	1.000*	1.000*	0.00	0.00
Daily.Avg	2.877*	4.794*	607.4	1.07
Total	276.238*	460.222*	58.312*	1.07

Table 8.1: Average results for the baseline experiment after 100 days of simulation. Description of the experiment can be found in Section 8.1. μ is the mean, Max is the average maximum for each day, σ is the standard deviation and γ is the skewness of the results.

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.269*	5.388*	685.2	1.02
PAR	1.180	1.248	0.026	0.00
		PHEV fairness		
PHEV(Ux)	0.000*	0.000*	0.0	0.00
PHEV(Avg)	0.00	0.00	0.00	0.00
		Transformer stability		
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	0.0	0.0	0.0	0.00
		Day-ahead imbalance		
Dayh.	0.000*	0.000*	0.000*	0.00
Dayh.(Dx)	266.414*	450.160*	58.473*	1.09
DDx/Tot.	1.000*	1.000*	0.00	0.00
Daily.Avg	2.775*	4.689*	609.1	1.09
Total	266.414*	450.160*	58.473*	1.09

Table 8.2: Average values after 100 days of simulation, for the minimum extreme experiment with no PHEVs present. Description of the experiment is found in Section 8.1

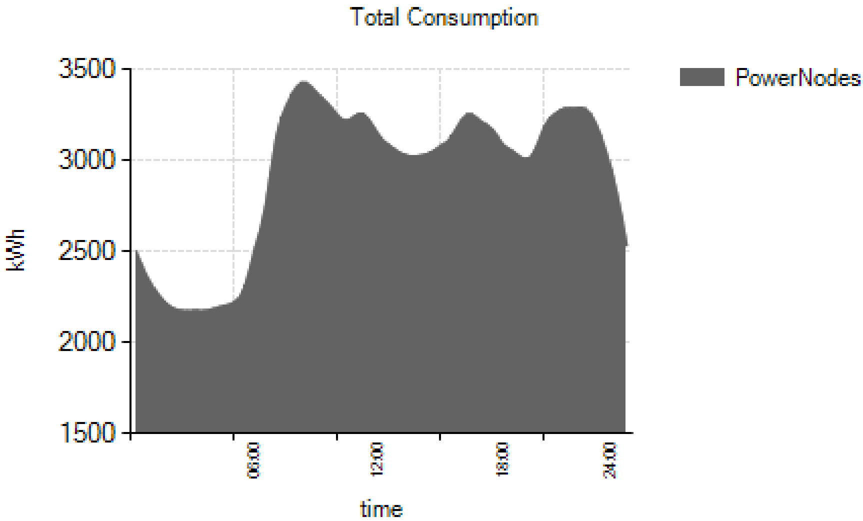

Figure 8.2: Sample image from the minimum extreme experiment for a selected day in the simulator. The minimum extreme experiment represent a scenario where no PHEVs are present in the grid.

capacity was also exceeded by an average of 939,4kW each day, with a maximum of 24204,2kW. The average total uncharged capacity of the PHEVs at time of departure was 3.62 MWh per day, which gives an average of $\frac{3{,}62e3MWh}{616} = 5.87kWh$ per PHEV.

In Table 8.2, the results of the minimal extreme experiment can be seen. The results show a peak-to-average ratio (PAR) with an average of 18,0%, with a maximum of 24,8%. Transformer capacity was exceeded by an average of 502,6kW each day, with a maximum of 13654,9kW. Average maximum peak during the simulation was 3.27 MWh, while the maximum daily peak was 5.39 MWh.

8.2 Day-ahead profile

Experiment 8.2a *Expected dayahead*

The results of running the simulator for 100 days with no scheduling mechanisms and calculating the dayahead profile using expected values can be seen in Table 8.3. The results show an average imbalance in volume between predictions and the dayahead profile of 3.77 MWh per day. This gives an average difference between dayahead and predictions of $\frac{3.77MW}{96} = 39.27kW$ per 15 min, which gives an average error of $\frac{0.03951}{2.87} = 1,36\%$. A sample day can also be seen in figure 8.3a.

Experiment 8.2b *Simulated dayahead*

Desc.	μ	Max	σ	γ
	Peak-shaving			
Peak	3.310*	5.388*	682.2	1.02
PAR	1.152	1.233	0.026	0.76
	Transformer stability			
Trf(Exc)	935.0	24164.0	3414.9	5.53
Trf(Flt)	0.0	0.0	0.0	0.00
	Day-ahead imbalance			
Dayh.	278.441*	462.187*	58.473*	1.09
Dayh.(Dx)	3.777*	7.089*	0.344*	9.16
DDx/Tot.	0.014*	0.018*	0.00	-0.39
Daily.Avg	2.877*	4.791*	607.6	1.07
Total	276.227*	459.981*	58.331*	1.07

Table 8.3: Baseline with day-ahead profile calculated using expected values, after 100 days of simulation

Desc.	μ	Max	σ	γ
	Peak-shaving			
Peak	3.310*	5.388*	683.3	1.01
PAR	1.152	1.230	0.026	0.75
	Transformer stability			
Trf(Exc)	931.1	23864.5	3383.0	5.54
Trf(Flt)	0.0	0.0	0.0	0.00
	Day-ahead imbalance			
Dayh.	276.214*	459.986*	58.350*	1.07
Dayh.(Dx)	0.446*	0.791*	0.108*	0.43
DDx/Tot.	0.002*	0.003*	0.00	0.39
Daily.Avg	2.877*	4.791*	607.4	1.07
Total	276.204*	459.941*	58.309*	1.07

Table 8.4: Baseline with day-ahead profile calculated using simulated values, after 100 days of simulation

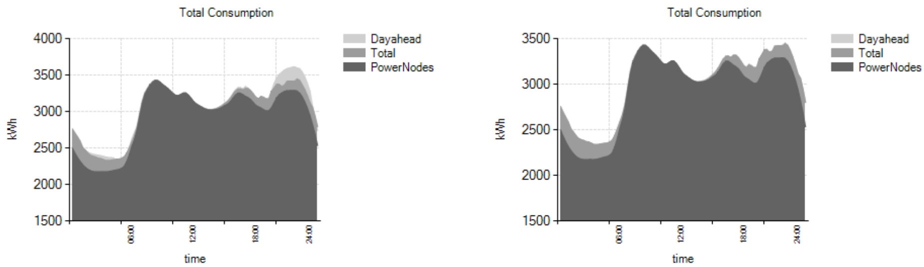

(a) Experiment 8.2a: Using expected values

(b) Experiment 8.2b: Using simulated values

Figure 8.3: Sample images from calculating the dayahead profiles for Experiment 8.2a and 8.2b. The legend is the same as described in Figure 8.1, with the addition of the day-ahead profile. The day-ahead profile is represented by the brightest area, and is barely visible in Figure 8.3a during peak-hours. In Figure 8.3b, the day-ahead profile approximation to the total consumption is too accurate to be clearly visible in the graph. Average results for all 100 days of simulation can be found in Tables 8.3

The results of running the simulator for 100 days with no scheduling mechanisms and calculating the dayahead profile using simulated values can be seen in Table 8.4. The results show an average difference in volume between predictions and the dayahead profile of 0.47MW per day. This gives an average difference between dayahead and predictions of $\frac{0.47MW}{96} = 4,89kW$ per 15 min, which gives an average error of $\frac{4,895e-3}{2.87} = 0,17\%$. A sample day can also be seen in figure 8.3b.

8.3 Scheduling mechanism

Experiment 8.3a *Proactive scheduling with peak-shaving*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.5. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 16%, and a maximum of 22,3%. Average maximum peak during the simulation was 3.272 MWh, while the maximum daily peak was 5.388 MWh.

Average dayahead volume per day was 278.46 MWh and the average total imbalance between the dayahead profile and the dayahead predictions were 7.629MWh.

Experiment 8.3b *Proactive scheduling with distance-rule*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the distance-rule algorithm for calculating the day-ahead profile,

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.272*	5.388*	683.0	1.03
PAR	1.160	1.223	0.024	0.17
		PHEV fairness		
PHEV(Ux)	10.989*	12.904*	870.8	-1.63
PHEV(Avg)	0.346	1.0	0.378	-4.06
		Transformer stability		
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	9.9	501.9	58.7	7.15
		Day-ahead imbalance		
Dayh.	278.460*	462.217*	58.476*	1.09
Dayh.(Dx)	7.629*	12.485*	0.814*	3.18
DDx/Tot.	0.029*	0.046*	0.01	0.06
Daily.Avg	2.825*	4.747*	608.8	1.08
Total	271.164*	455.665*	58.449*	1.08

Table 8.5: After 100 days of simulation using proactive scheduling with peak-shaving on expected values for PHEV demand. (*theta*=0.950, *alpha*=0.500 and PHEV learning window of 40)

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.269*	5.388*	685.2	1.02
PAR	1.150	1.209	0.024	0.09
		PHEV fairness		
PHEV(Ux)	8.678*	10.925*	516.9	-0.33
PHEV(Avg)	0.482	1.0	0.407	0.02
		Transformer stability		
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	0.0	0.0	0.0	0.00
		Day-ahead imbalance		
Dayh.	278.441*	462.187*	58.473*	1.09
Dayh.(Dx)	5.334*	12.000*	0.726*	8.10
DDx/Tot.	0.020*	0.028*	0.00	-0.09
Daily.Avg	2.845*	4.752*	605.9	1.07
Total	273.107*	456.219*	58.170*	1.07

Table 8.6: After 100 days of simulation using proactive scheduling with distance-rule on expected values for PHEV demand (*theta*=1.000 and PHEV learning window of 40)

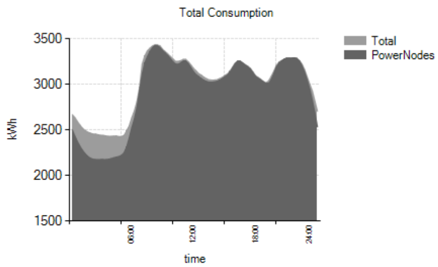

(a) Experiment 8.3f: Mixed strategy

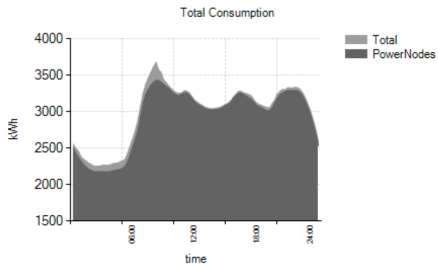

(b) Experiment 8.3e: Uniform strategy

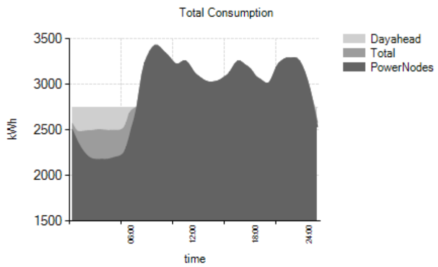

(c) Experiment 8.3d: Reactive scheduling, distance-rule

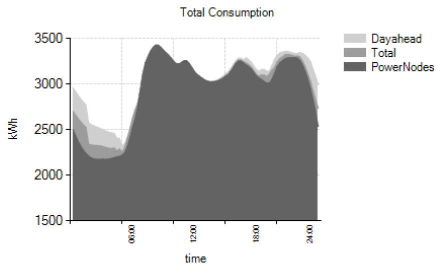

(d) Experiment 8.3c: Reactive scheduling, peak-shaving

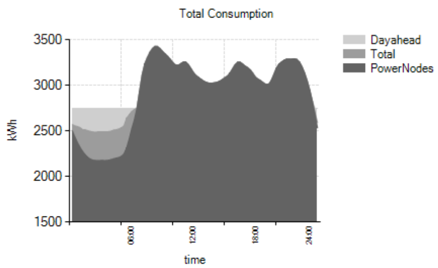

(e) Experiment 8.3b: Proactive scheduling, distance-rule

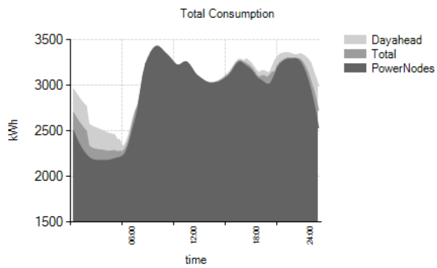

(f) Experiment 8.3a: Proactive scheduling, peak-shaving

Figure 8.4: Sample powergraphs from the scheduling experiments. The images show the consumption of the PowerNodes over a day, the combined consumption of the PowerNodes and PHEVs, and the day-ahead profile estimated for that day.

can be seen in Table 8.6. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 15%, and a maximum of 21%. Average maximum peak during the simulation was 3.27 MWh, while the maximum daily peak was 5.39 MWh.

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.272*	5.388*	682.6	1.03
PAR	1.157	1.219	0.023	0.20
		PHEV fairness		
PHEV(Ux)	10.126*	12.389*	955.0	-0.77
PHEV(Avg)	0.396	1.0	0.377	-2.61
		Transformer stability		
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	10.0	465.4	58.5	6.74
		Day-ahead imbalance		
Dayh.	278.460*	462.217*	58.476*	1.09
Dayh.(Dx)	6.894*	12.297*	0.919*	2.58
DDx/Tot.	0.026*	0.043*	0.01	0.15
Daily.Avg	2.832*	4.752*	607.5	1.08
Total	271.898*	456.195*	58.317*	1.08

Table 8.7: After 100 days of simulation using reactive scheduling with peak-shaving on expected values for PHEV demand. (*theta*=0.950, *alpha*=0.500 and PHEV learning window of 40)

Average dayahead volume per day was 278.46 MWh and the average distance between the dayahead profile and the dayahead predictions were 5,33 MWh.

Experiment 8.3c *Reactive scheduling with peak-shaving*

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.7. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 15,7%, and a maximum of 21,9%. Average maximum peak during the simulation was 3.27 MWh, while the maximum daily peak was 5.38 MWh.

Average dayahead volume per day was 278.46 MWh and the average distance between the dayahead profile and the dayahead predictions were 6,90 MWh.

Experiment 8.3d *Reactive scheduling with distance-rule*

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the distance-rule for calculating the day-ahead profile, can be seen in Table 8.8. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and

Desc.	μ	Max	σ	γ
	Peak-shaving			
Peak	3.269*	5.388*	685.2	1.02
PAR	1.152	1.214	0.024	0.08
	PHEV fairness			
PHEV(Ux)	9.308*	10.722*	485.9	-1.73
PHEV(Avg)	0.445	1.0	0.380	-0.20
	Transformer stability			
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	0.0	0.0	0.0	0.00
	Day-ahead imbalance			
Dayh.	278.441*	462.187*	58.473*	1.09
Dayh.(Dx)	5.717*	11.999*	0.667*	8.62
DDx/Tot.	0.022*	0.027*	0.00	-0.35
Daily.Avg	2.841*	4.750*	606.2	1.07
Total	272.724*	456.029*	58.194*	1.07

Table 8.8: After 100 days of simulation using reactive scheduling with distance-rule on expected values for PHEV demand ($theta$=1.000 and PHEV learning window of 40)

Desc.	μ	Max	σ	γ
	Peak-shaving			
Peak	3.363*	5.499*	728.2	0.89
PAR	1.185	1.255	0.031	0.17
	PHEV fairness			
PHEV(Ux)	9.659*	10.640*	497.3	-4.16
PHEV(Avg)	0.423	1.0	0.332	-1.41
	Transformer stability			
Trf(Exc)	502.6	13654.9	2013.4	5.49
Trf(Flt)	116.4	2556.9	373.2	4.87
	Day-ahead imbalance			
Dayh.	0.000*	0.000*	0.000*	0.00
Dayh.(Dx)	272.254*	455.329*	58.226*	1.07
DDx/Tot.	1.000*	1.000*	0.00	0.00
Daily.Avg	2.836*	4.743*	606.5	1.07
Total	272.254*	455.329*	58.226*	1.07

Table 8.9: After 100 days of simulation using the Uniform strategy with learning window 40

average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 15,2%, and a maximum of 21,4%. Average maximum peak during the simulation was 3.27 MWh, while the maximum daily peak was 5.39 MWh.

Average dayahead volume per day was 278.44 MWh and the average distance between the dayahead profile and the dayahead predictions were 5.71 MWh.

Experiment 8.3e *Uniform strategy*

The results of running the simulator for 100 days using the uniform strategy mechanism can be seen in Table 8.9. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The results show a peak-to-average ratio (PAR) with an average of 18,5%, and a maximum of 25,5%. Average maximum peak during the simulation was 3.363 MWh, while the maximum daily peak was 5.50 MWh.

Experiment 8.3f *Mixed strategy*

The results of running the simulator for 100 days using the mixed strategy mechanism can be seen in Table 8.10. From these results, the most interesting variables to note are the values for peak-to-average ratio (PAR), transformer capacity exceeded, and average uncharged capacity of the PHEVs at time of departure. The

Desc.	μ	Max	σ	γ
		Peak-shaving		
Peak	3.271*	5.388*	686.0	1.02
PAR	1.147	1.206	0.024	0.05
		PHEV fairness		
PHEV(Ux)	7.310*	8.238*	314.8	-0.19
PHEV(Avg)	0.548	1.0	0.412	1.31
		Transformer stability		
Trf(Flt)	126.4	2739.6	361.1	4.85
Trf(Exc)	502.6	13654.9	2013.4	5.49
		Day-ahead imbalance		
Dayh.	0.000*	0.000*	0.000*	0.00
Dayh.(Dx)	273.934*	457.497*	58.224*	1.07
DDx/Tot.	1.000*	1.000*	0.00	0.00
Daily.Avg	2.853*	4.766*	606.5	1.07
Total	273.934*	457.497*	58.224*	1.07

Table 8.10: After 100 days of simulation using the Mixed strategy with learning window 40

results show a peak-to-average ratio (PAR) with an average of 14,7%, and a maximum of 20,6%. Average maximum peak during the simulation was 3.27 MWh, while the maximum daily peak was 5.39 MWh.

8.4 Transformer constraints

Experiment 8.4a *Proactive scheduling with peak-shaving*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.5. The results show that transformer capacity was exceeded by an average of 502,6kW each day, with a maximum of 13654,9kW. Average demand filtered by the Transformer agent was 9,9 kW.

Experiment 8.4b *Proactive scheduling with distance-rule*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the distance-rule algorithm for calculating the day-ahead profile, can be seen in Table 8.6. The results show that transformer capacity was exceeded by an average of 502,6 kW each day, with a maximum of 13654,9kW. Average demand filtered by the Transformer agent was 0.0 kW.

Experiment 8.4c *Reactive scheduling with peak-shaving*

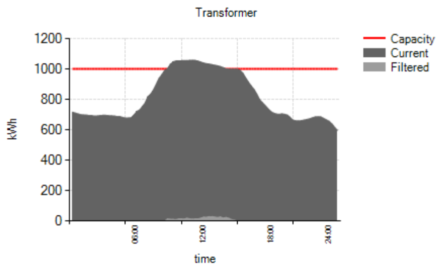

(a) Experiment 8.4f: Mixed strategy

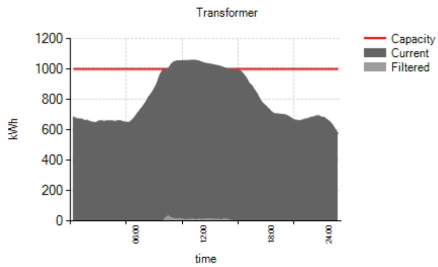

(b) Experiment 8.4e: Uniform strategy

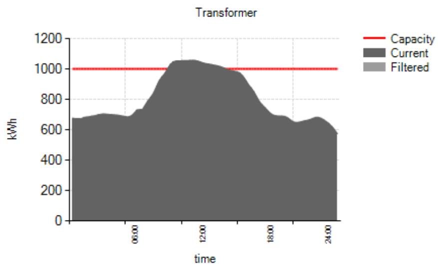

(c) Experiment 8.4d: Reactive scheduling, distance-rule

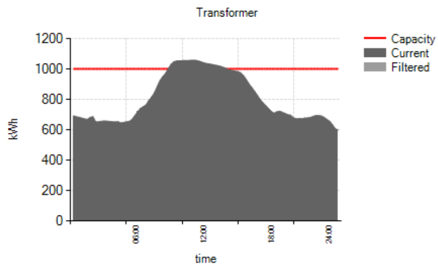

(d) Experiment 8.4c: Reactive scheduling, peak-shaving

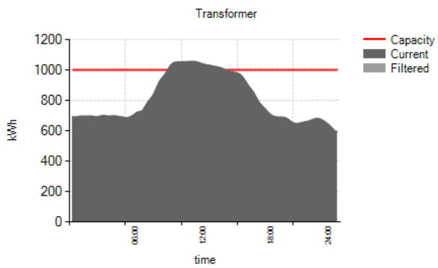

(e) Experiment 8.4b: Proactive scheduling, distance-rule

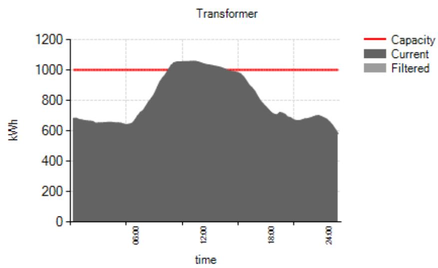

(f) Experiment 8.4a: Proactive scheduling, peak-shaving

Figure 8.5: Sample graphs for a single Transformer agent. The red line represent the maximum capacity of the transformer, while the dark gray area represent the current demand in the grid. The light gray area is the demand filtered by the Transformer agents, and is barely visible in the Mixed strategy and Uniform strategy, but absent in the other graphs. Further discussion of these results can be found in Section 9.3

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.7. The results show that transformer capacity was exceeded by an average of 502,6 kW each day, with a maximum of 13654,9kW. Average demand filtered by the Transformer agent was 10.0 kW.

Experiment 8.4d *Reactive scheduling with distance-rule*

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the distance-rule for calculating the day-ahead profile, can be seen in Table 8.8. The results show that transformer capacity was exceeded by an average of 502,6 kW each day, with a maximum of 13654,9kW. Average demand filtered by the Transformer agent was 0.0 kW.

Experiment 8.4e *Uniform strategy*

The results of running the simulator for 100 days using the uniform strategy mechanism can be seen in Table 8.9. The results show that transformer capacity was exceeded by an average of 502,6kW each day, with a maximum of 13654,9 kW. Average demand filtered by the Transformer agent was 116,4 kW, with a maxmimum of 2556,9 kW.

Experiment 8.4f *Mixed strategy*

The results of running the simulator for 100 days using the mixed strategy mechanism can be seen in Table 8.10. The results show that transformer capacity was exceeded by an average of 502,6kW each day, with a maximum of 13654,9kW. Average demand filtered by the Transformer agent was 126,4 kW, with a maxmimum of 2739,6 kW.

8.5 PHEV fairness

Experiment 8.5a *Proactive scheduling with peak-shaving*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.5. Total uncharged capacity of the PHEVs at time of departure was 10,99 MWh per day, which gives an average of $\frac{10{,}99e3kWh}{616} = 17,84kWh$ per PHEV. Average battery capacity at time of departure was 34,6%

Experiment 8.5b *Proactive scheduling with distance-rule*

The results of running the simulator for 100 days with a proactive scheduling mechanism, using the distance-rule algorithm for calculating the day-ahead profile, can be seen in Table 8.6. Total uncharged capacity of the PHEVs at time of departure was 8,68 MWh per day, which gives an average of $\frac{8{,}68e3kWh}{616} = 14,09kWh$ per PHEV. Average battery capacity at time of departure was 48,2%

Experiment 8.5c *Reactive scheduling with peak-shaving*

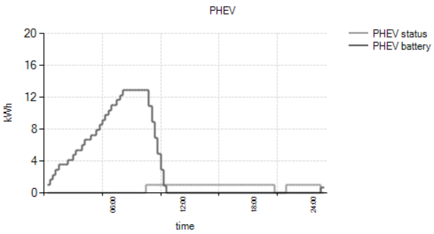

(a) Experiment 8.5a: Proactive scheduling, peak-shaving

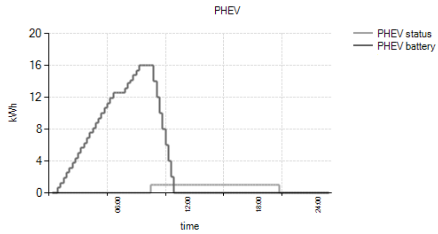

(b) Experiment 8.5b: Proactive scheduling, distance-rule

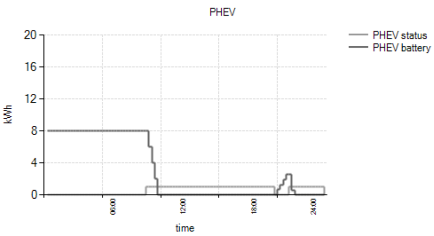

(c) Experiment 8.5c: Reactive scheduling, peak-shaving

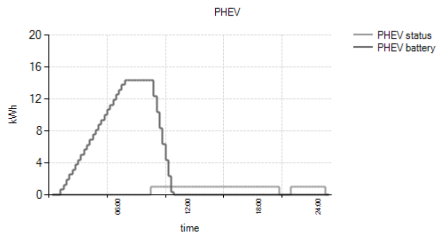

(d) Experiment 8.5d: Reactive scheduling, distance-rule

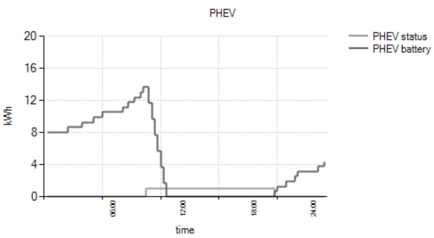

(e) Experiment 8.5e: Uniform strategy

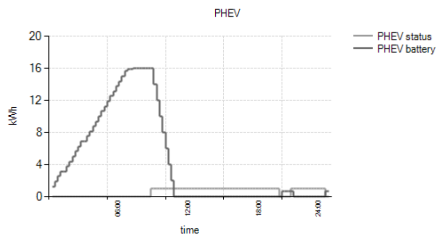

(f) Experiment 8.5f: Mixed strategy

Figure 8.6: Sample graphs of the daily battery state for a single PHEV. The dark gray line represent the current battery capacity of the PHEV while the light gray line represent a step-function which determined whether a PHEV is disconnected from the grid. These samples are taken from a PHEV with the commuter profile, and since the commuter profile spends much of its time driving, the batteries will often be empty by the time it returns home.

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the peak-shaving algorithm for calculating the day-ahead profile, can be seen in Table 8.7. Total uncharged capacity of the PHEVs at time of departure was 10,126 MWh per day, which gives an average of $\frac{10{,}126e3kWh}{616} = 16,42kWh$ per PHEV. Average battery capacity at time of departure was 39,6%

Experiment 8.5d *Reactive scheduling with distance-rule*

The results of running the simulator for 100 days with a reactive scheduling mechanism, using the distance-rule for calculating the day-ahead profile, can be seen in Table 8.8. Total uncharged capacity of the PHEVs at time of departure was 9,3MWh per day, which gives an average of $\frac{9{,}3e3kWh}{616} = 15,1kWh$ per PHEV. Average battery capacity at time of departure was 44,5%

Experiment 8.5e *Uniform strategy*

The results of running the simulator for 100 days using the uniform strategy mechanism can be seen in Table 8.9. Total uncharged capacity of the PHEVs at time of departure was 9,66 MWh per day, which gives an average of $\frac{9{,}66e3kWh}{616} = 15,68kWh$ per PHEV. Average battery capacity at time of departure was 42,3%

Experiment 8.5f *Mixed strategy*

The results of running the simulator for 100 days using the mixed strategy mechanism can be seen in Table 8.10. Total uncharged capacity of the PHEVs at time of departure was 7,31 MWh per day, which gives an average of $\frac{7{,}31e3kWh}{616} = 11,87kWh$ per PHEV. Average battery capacity at time of departure was 54,8%

9 Discussion

In this Section we will discuss the results from Section 8 that were obtained from running the experiments defined in Section 7. First we will discuss the results from running the experiments on how to calculate the day-ahead portfolio in 9.1. After which we will discuss the results from the scheduling experiments, the stability experiments and the PHEV fairness experiments.

9.1 Day-ahead portfolio

From the results in Section 8.2, we see that by calculating the dayahead portfolio using expected values we get an error of 1,19% from the dayahead predictions to the actual energy consumption. Similarly, for calculating the dayahead portfolio using simulated value, with get a negligible error of 0,17%. While the accuracy of these number largely comes from the fact that the historical data for the PowerNodes are used, both as the basis for the dayahead predictions, and as input for the real-time energy consumption in the simulator, we are only interested in getting a reasonably accurate method for calculating the day-ahead portfolio for the simulator, and not in finding a realistic method for calculating it in a real life scenario. In fact, a small error in calculating the expected PHEV demand is preferable to an exact result. The reason for this is that it will make it easier to separate the results of the reactive and the proactive scheduling algorithms. Since both these algorithms will schedule the PHEV demand until imbalances between the day-ahead portfolio and the real-time demand is canceled, this also means that if both algorithms are able to fully cancel these imbalances, or if the day-ahead portfolio is insufficient to schedule all the PHEVs, then they will also perform identically. However, if the day-ahead portfolio exceeds real-time demand, then we should be able to see some differences between the performance of the to algorithm. To this end, the small error that we got from the results of doing the day-ahead experiments show that using Expected values is preferred over Simulated values.

9.2 Scheduling experiments

In this section we will discuss the performance of the different scheduling mechanisms with respect to Hypothesis 3.1 and 3.2. The mechanisms will first be discussed individually, comparing their performance to the baseline (minimal extreme experiment), after which we will do a general comparison of the different mechanisms against each other. The most interesting variable to the performance of the scheduling mechanisms with respect to these hypotheses, is the peak-to-average ratio (PAR) and the average maximum peak.

Proactive scheduling with peak-shaving

The results from this experiment can be found in Table 8.5, while the powergraph from a sample day can be seen in Figure 8.4f. With the proactive scheduling mechanism, and the peak-shaving algorithm, an average PAR of 1.160 was obtained.

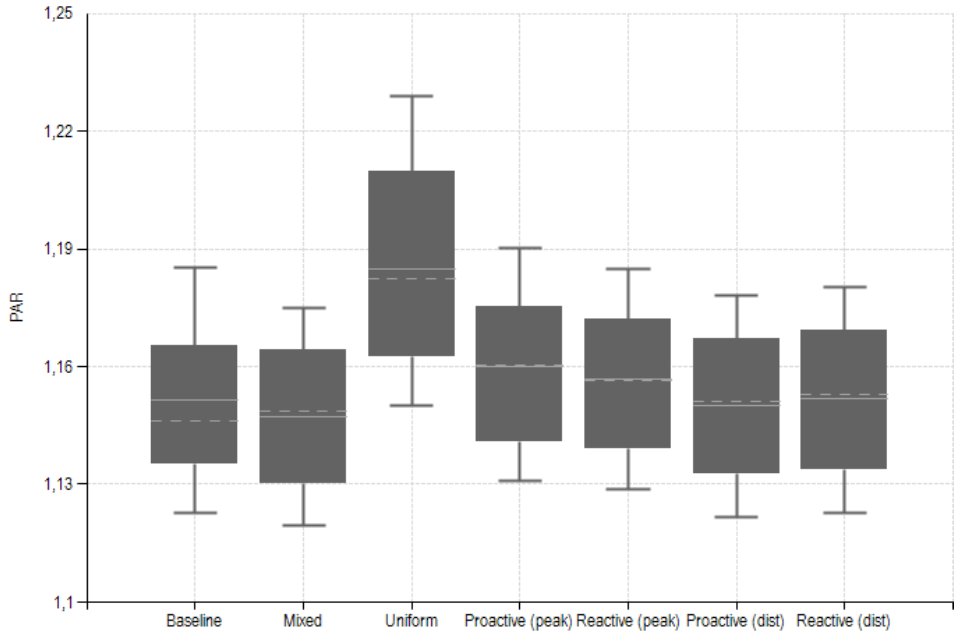

Figure 9.1: A boxplot graph showing the average PAR for each of the different mechanisms. Lower is better.

The average peak in the simulation was 3.27 MW, while the maximum peak was 5.39 MW. The average PAR was almost the same as with the baseline, while the maximum PAR remained the same, showing an negligible reduction from 1.226 to 1.223.

Proactive scheduling with distance-rule

The results from this experiment can be found in Table 8.6, while a powergraph from a sample day can be seen in Figure 8.4e. With the proactive scheduling algorithm, and using the distance-rule algorithm an average PAR of 1.15 was obtained. The average peak in the simulation was 3.27 MW, while the maximum peak was 5.39 MW. Compared to the baseline, with a PAR of 1.152 and average and maximum peak of 3.31 and MW in the baseline, this may seem to be a modest improvement. Perhaps more significantly, the mechanism was able to reduce the maximum PAR from 1.226 in the baseline results to 1.21, and the average peak from 3.31 MW to 3.269 MW which is the same result as in the minimal extreme experiment.

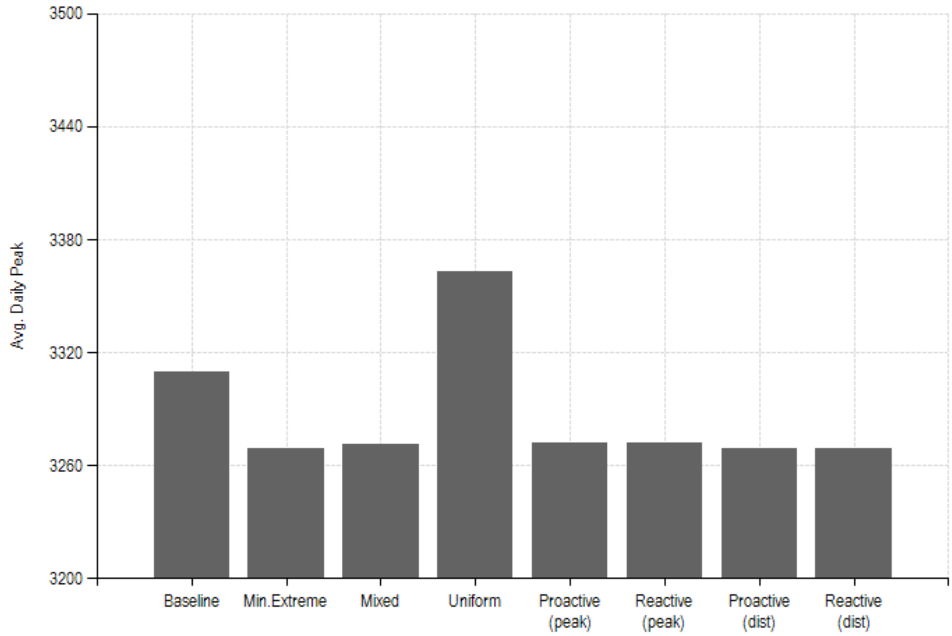

Figure 9.2: A graph showing the average daily peak for each of the different mechanisms. Lower is better.

Reactive scheduling with peak-shaving

With the reactive scheduling algorithm, and using the peak-shaving algorithm, an average PAR of 1.157 was obtained. The average peak in the simulation was 3.272 MW, while the maximum peak was 5.388 MW. The average PAR was almost the same as with the baseline, while the maximum PAR showed a minor reduction from 1.226 to 1.219.

Reactive scheduling with distance-rule

With the reactive scheduling algorithm, and the distance-rule algorithm, an average PAR of 1.152 was obtained. The average peak in the simulation was 3.27 MW, while the maximum peak was 5.39 MW. The average PAR was almost the same as with the baseline, while the maximum PAR showed a minor reduction from 1.226 to 1.214.

Uniform strategy

With the Uniform strategy mechanism, an average PAR of 1.185 was obtained. The average peak in the simulation was 3.36 MW, while the maximum peak was 5.50 MW. This is an increase over the baseline by .033 in average PAR, and an increase by 32000 kW in average peak and 111000 kW in maximum peak. Compared to the minimal extreme experiment, the results show an increase in average and maximum peak by 94000 kWh and 111000 kWh respectively.

Mixed strategy

With the Mixed strategy mechanism, an average PAR of 1.145 was obtained. The average peak in the simulation was 3.27 MW, while the maximum peak was 5.39 MW. This is a minor reduction from the baseline by .005 in average PAR, and a reduction of 39000 kW in average peak. Additionally, the results show not much difference in average and maximum peak compared to the minimal extreme experiment. This is a positive result for the mixed strategy mechanism.

Comparisons

The overall goal with these experiments were to investigate the Hypotheses 3.1 and 3.2. To measure the performance of the mechanisms with respect to those hypotheses, we investigated the result of running the simulations with respect to the values for PAR, average and maximum peak. The results are interpreted as positive if they can outperform the baseline simulations, but they are expected to be no better than the minimal extreme experiment when considering average maximum peak.

The results from running the experiments show that all of the different mechanisms demonstrate a positive result compared to the baseline when considering average peak, where each mechanism except the Uniform strategy managed to obtain a similar average maximum peak to the *minimum extreme experiment*, which is illustrated in Figure 9.2. This is a *positive* result, considering that the results from the minimum extreme experiment represent the best case scenario with respect to average maximum peak.

Why the Uniform strategy performed so poorly is illustrated in Figure 8.4b, where a uniform distribution of the PHEV demand can be seen. Because of the nature of the Uniform strategy mechanism, the distribution of the PHEV demand is equally probable for all timeframes. This means that some of the distributed PHEV demand will inevitably be allocated to the early morning peak hours, which is the reason for the high PAR and average maximum peak values for the Uniform strategy. This is in contrast to the other mechanisms where the PHEVs seem to be fully charge within that time, and where each of the other mechanisms managed to obtain positive results for the average maximum peak.

With the exception of the Uniform strategy mechanism, and with respect to the average maximum peak and the average PAR values, none of the scheduling mechanisms managed to outperform each other. Instead, all the other mechanisms managed to perform well in maintaining a low average peak. And while the Mixed strategy mechanism showed some signs of performing better with respect to the average PAR, the gains were relatively small. As to why no improvement in average PAR value was seen, this discussion is deferred to Section 9.4. This is because these results are best explained when seen in perspective of the results of the other experiments which will discussed subsequently.

9.3 Stability experiments

In this section we will discuss the performance of the different scheduling mechanisms with respect to Hypothesis 3.3. The performance of the mechanisms will be evaluated with respect to how much the capacity is exceeded by on average. This will be compared against the baseline, which will represent the worst case scenario, and the minimal extreme experiment which will represent the best case scenario. The mechanisms will first be discussed individually, comparing their performance to the baseline (minimal extreme experiment), after which we will do a general comparison of the different mechanisms against each other.

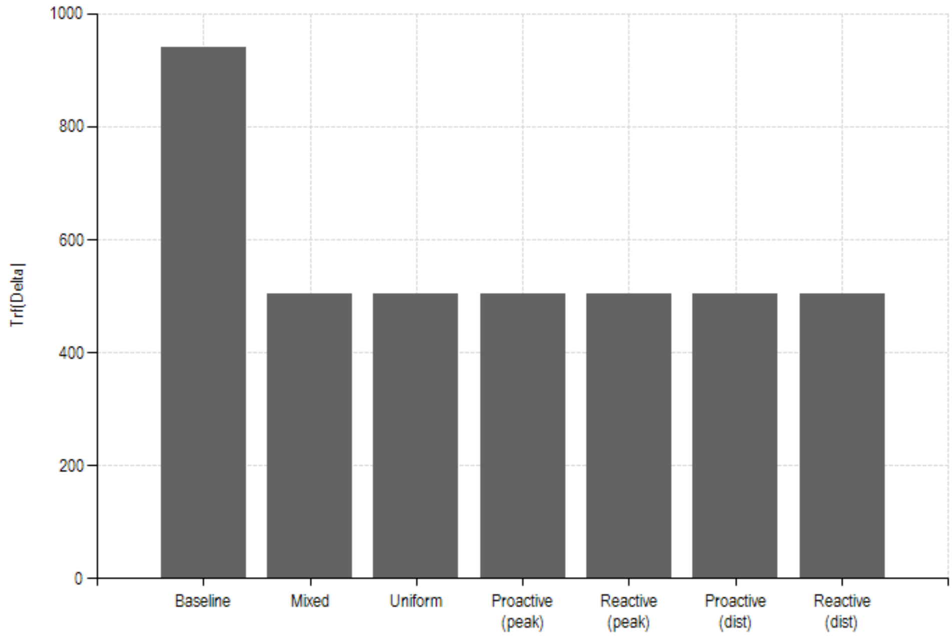

Figure 9.3: A graph showing the average capacity exceeded over all transformers, for each of the different mechanisms. Lower is better.

Proactive scheduling with peak-shaving

With proactive-scheduling and the peak-shaving algorithm, transformer capacity was exceeded by 502,6 kWh each day. This is equal to the minimal extreme experiment, and seem to suggest that the scheduling mechanism is able to schedule the PHEV demand so that no extra demand is experienced during times when trans-

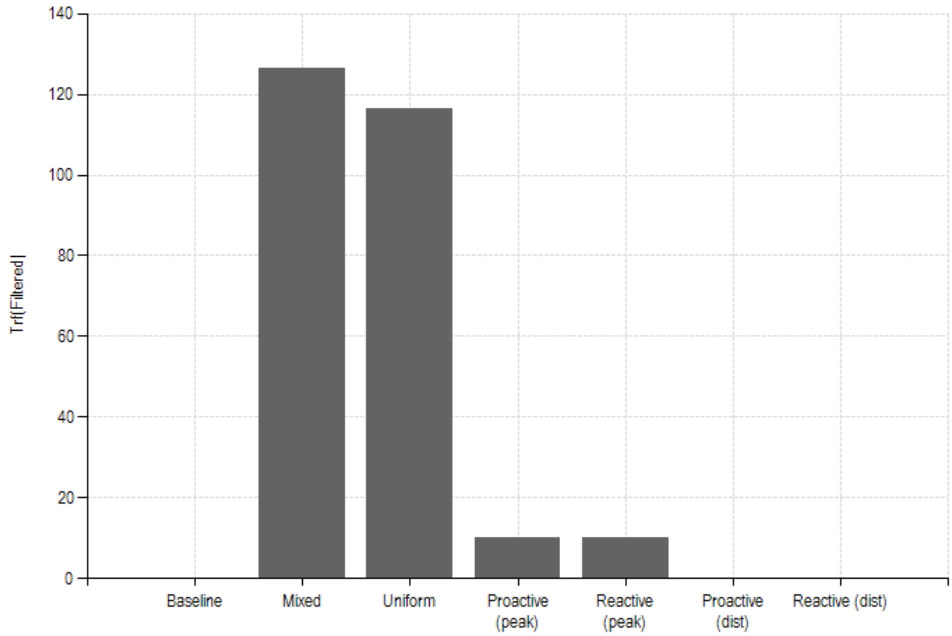

Figure 9.4: A graph showing the average PHEV demand filtered, for each of the different mechanisms. Lower is better.

former capacity is exceeded. Additionally, 9,9 kWh was filtered by the Transformer agent.

Proactive scheduling with distance-rule

With proactive-scheduling and the distance-rule algorithm, transformer capacity was exceeded by 502,6 kWh each day. This is equal to the minimal extreme experiment, and seem to suggest that the scheduling mechanism is able to schedule the PHEV demand so that no extra demand is experienced during times when transformer capacity is exceeded.

Reactive scheduling with peak-shaving

With reactive-scheduling and the peak-shaving algorithm, transformer capacity was exceeded by 502,6 kWh each day. This is equal to the minimal extreme experiment, and seem to suggest that the scheduling mechanism is able to schedule the PHEV demand so that no extra demand is experienced during times when transformer capacity is exceeded. Additionally, 9,9 kWh was filtered by the Transformer agent.

Reactive scheduling with distance-rule

With reactive-scheduling and the distance-rule algorithm, transformer capacity was exceeded by 502,6 kWh each day. This is equal to the minimal extreme experiment, and seem to suggest that the scheduling mechanism is able to schedule the PHEV demand so that no extra demand is experienced during times when transformer capacity is exceeded.

Uniform strategy

With the Uniform strategy mechanism, transformer capacity was exceeded by 502,6 kWh each day. However, 116,4 kWh of demand was filtered by the Transformer agent, which suggests that some of the PHEV demand was scheduled at times when demand was already high. This seem like a reasonable result, considering that the uniform strategy will distribute PHEV demand based on an equal probability distribution.

Mixed strategy

With the Mixed strategy mechanism, transformer capacity was exceeded by 502,6 kWh each day. However, 126,4 kWh of demand was filtered by the Transformer agent, which suggests that some of the PHEV demand was scheduled at times when demand was already high.

Comparisons

The overall goal with these experiments were to investigate the Hypotheses 3.3. To measure the performance of the mechanisms with respect to this hypothesis, we investigated the results of running the simulations with respect to the values for average transformer capacity exceeded. The results are positive if they can outperform the baseline simulations, but they were expected to be no better than the minimal extreme experiment.

The results from running the experiments show that all of the different mechanisms demonstrate a positive result compared to the baseline. All of the mechanisms showed equal performance to the minimal extreme experiment which is a *positive* result, considering that the minimal extreme experiment represent the best case scenario. However, in the reactive scheduling with peak-shaving mechanism and the uniform- and mixed strategies, some of the PHEV demand was filtered by the Transformer agents. This does not affect the performance of the mechanisms with respect to Hypothesis 3.1, as the stability of the grid remains well preserved within the definition that the stability is preserved relative to the baseline; if no change in average Transformer capacity exceeded is observed. Still, it might suggest a weakness for these mechanisms with respect to PHEV fairness experiments and Hypothesis 3.4 which will be discussed in Section 9.4.

9.4 PHEV fairness experiments

In this section we will discuss the performance of the different scheduling mechanisms with respect to Hypothesis 3.4. The performance of the mechanisms will be evaluated with respect to the average uncharged battery capacity of a PHEV at time of departure. This will be compared against the baseline, which will represent the worst case scenario, and the minimal extreme experiment which will represent the best case scenario. The mechanisms will first be discussed individually, comparing their performance to the baseline (minimal extreme experiment), after which we will do a general comparison of the different mechanisms against each other.

Baseline

From the baseline experiment in Section 8.1, it can be seen that the average uncharged battery capacity of the PHEVs at time of departure was 77,61%. This is the value that will be used to compare with when evaluating the performance of the other mechanisms.

Proactive scheduling with peak-shaving

With proactive scheduling and the peak-shaving algorithm, average battery capacity by the PHEVs at time of departure was 34,57%. This is a difference of -43,04% from the baseline.

Proactive scheduling with distance-rule

With proactive scheduling and the distance-rule algorithm, average battery capacity by the PHEVs at time of departure was 48,20%. This is a difference of -29,41%.

Reactive scheduling with peak-shaving

With reactive scheduling and the peak-shaving algorithm, average battery capacity by the PHEVs at time of departure was 39,55%. This is a difference of -38,31%.

Reactive scheduling with distance-rule

With reactive scheduling and the distance-rule algorithm, average battery capacity by the PHEVs at time of departure was 44,53%. This is a difference of -33,08%.

Uniform strategy

With the Uniform strategy mechanism, average battery capacity by the PHEVs at time of departure was 42,34%. This is a difference of -35,27%.

Mixed strategy

With the Mixed strategy mechanism, the average battery capacity by the PHEVs at time of departure was 56,37%. This is a difference of -21,24%.

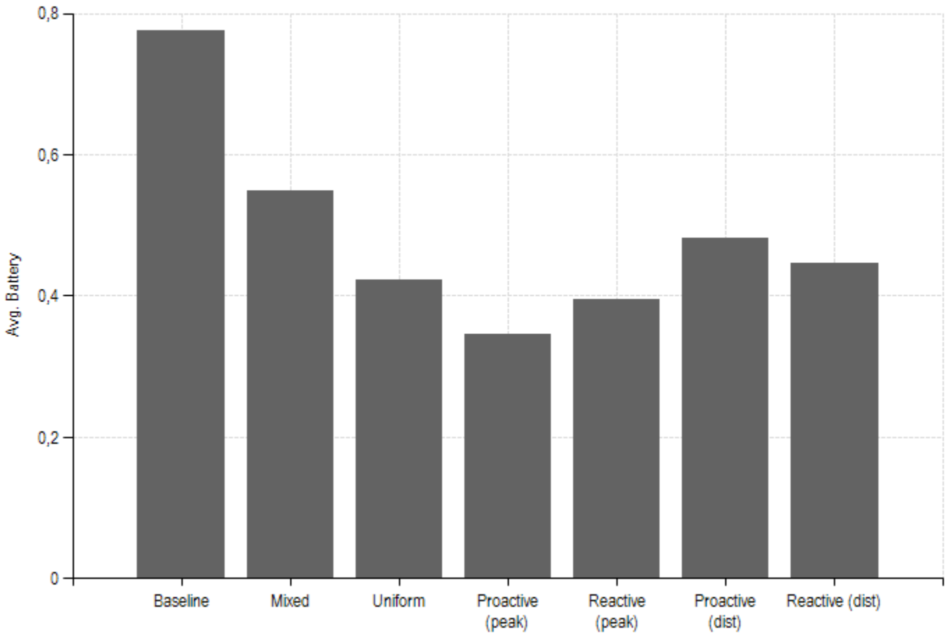

Figure 9.5: A graph showing the average battery capacity at time of departure, for each of the different mechanisms. Higher is better.

Comparisons

The overall goal with these experiments were to investigate Hypothesis 3.4. To measure the performance of the mechanisms with respect to this hypothesis, we investigated the result of running the simulations with respect to the values for average transformer capacity exceeded. The results are interpreted as positive if they can outperform the baseline simulations, but they were expected to be no better than the minimal extreme experiment.

The results from the PHEV fairness experiments seem to suggest that there is a cost to the average battery capacity of the PHEVs when using the scheduling mechanisms. From Figure 9.5, we see that the Mixed mechanism is the overall best performing scheduling mechanism with a difference of -21.24% compared to the baseline experiment. More troubling is the results from calculating the day-ahead portfolio using the peak-shaving algorithm in the centralized scheduling mechanisms, which showed performance penalties of -43,04% for the proactive scheduling mechanism, and -38,31% for the reactive scheduling mechanism. With the distance-rule algorithm, the situation was somewhat better with -29,41% and -33,08%, respectively, but both were still lower than the Mixed stratey mechanism.

An illustration of how the average battery capacity at time of departure varied depending on time of day can be seen in Figure 9.7. This illustrates why none of

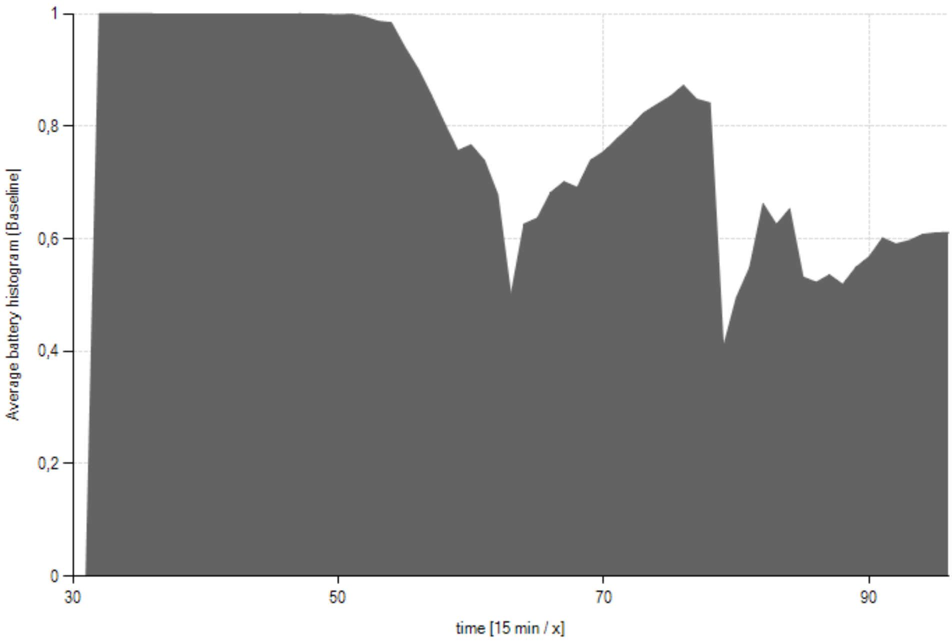

Figure 9.6: A graph showing the ratio of average battery capacity at time of departure, where the Y axis is the ratio of average battery capacity and X is the time of departure. Higher is better

the mechanisms, including the baseline experiment, managed to obtain an average value closer to 100%. The reason for this is illustrated by how the average value drops during the course of the day, as the PHEV profiles are designed so that the PHEVs are most active during the afternoon and evening hours. This means that many of the PHEVs when returning home from work, are connected to the grid for such a short time that even if all the time is spent charging, there is insufficient time to fully charge before they leave home again.

Also, returning to the concern that was raised about the values for the PAR in Section 9.2: Considering the results from the PHEV fairness experiments, we see that most of the different mechanisms performed worse with respect to charging the PHEVs compared to the baseline experiment. Since the PHEVs are charging less in the scheduling mechanisms, this means that the average consumption will be lower. And if the average peak decreases less than the total average decreases, then the PAR will increase compared to the baseline. And while all the mechanisms showed *postive* results, considering that they all performed similar to the best case scenario, there is a natural limit to the potential for peak-reduction because of the morning peak hours. This can be seen in Figures 8.3f, 8.3d and 8.3b, where all of the PHEV demand during the afternoon peak hours have been moved to low-demand hours during the night. However, because the early morning peak hour

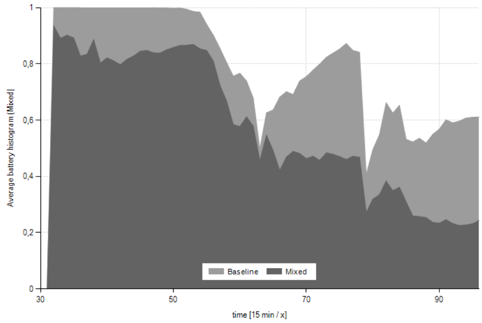

(a) Mixed strategy

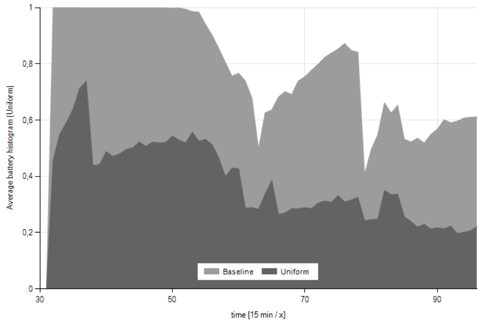

(b) Uniform strategy

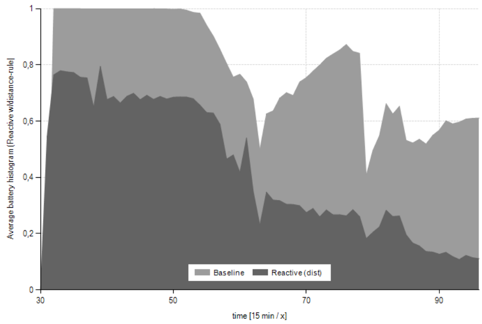

(c) Reactive scheduling, distance-rule

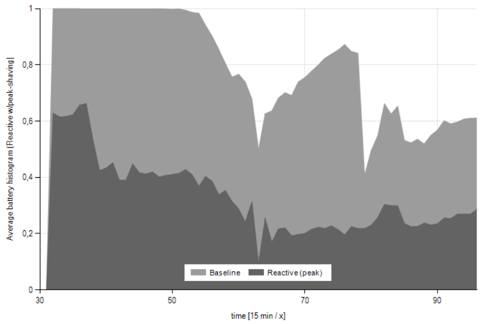

(d) Reactive scheduling, peak-shaving

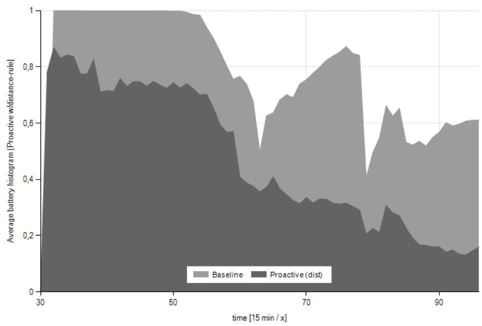

(e) Proactive scheduling, distance-rule

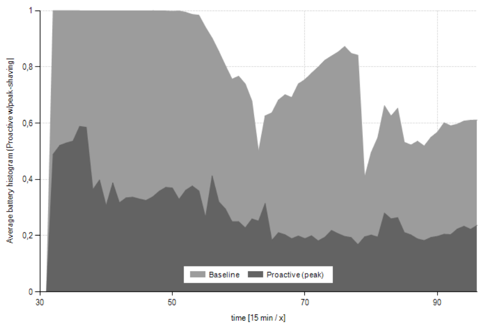

(f) Proactive scheduling, peak-shaving

Figure 9.7: These graphs show the ratio of average battery capacity at time of departure for each of the different mechanisms. Each mechanism is contrasted with the results from the baseline experiment for comparison.

is higher than the afternoon peak for that particular day, the peak value for this day will not be lower. This explains much of the lack of results when considering the PAR value. Still, since all the mechanisms managed to perform similar to the best case scenario when considering the average peak value, this suggests that the only limitation to the peak-shaving effect for these mechanisms depend on factors beyond the control of the mechanisms.

Reactive scheduling vs proactive scheduling

While both the reactive and proactive mechanisms are very similar in technique, there are some minor differences. Where the reactive scheduling mechanism will try to minimize the imbalances as soon as possible, the proactive mechanism will try to minimize the average imbalance. However, as to which of these perform better at scheduling the PHEVs, the results were inconclusive. With the peak-shaving algorithm, the reactive scheduling mechanism seemed to perform better, obtaining a 4,98% improvement over the proactive mechanism, while with the distance-rule algorithm, the proactive mechanism obtained a 7,34% higher average than the reactive mechanism. This seems to suggest that which mechanism is better depends on how the day-ahead imbalances are distributed. In the distance-rule algorithm, the day-ahead imbalances are distributed mostly during night hours, as seen in Figures 8.4c and 8.4e. This seems to favor a proactive mechanism where PHEVs are charged so that the average imbalances are reduced. Meanwhile, in the peak-shaving algorithm, seen in Figures 8.4d and 8.4f, the day-ahead imbalances are distributed closer to peak hours and are less uniformly distributed, which seems to favor a reactive mechanism.

For the stability experiments, the results show almost no difference in the amount of demand filtered by the Transformer agents. Instead, the difference seem to be more dependent on which algorithm is used to calculate the day-ahead portfolio. In neither the reactive- or the proactive mechanism did the Transformer agent have to filter demand when the distance-rule algorithm had been used. With the peak-shaving algorithm, a small amount of filtering was done, but the amount was almost similar both for the reactive- and the proactive mechanism suggesting that determining factor was which day-ahead algorithm was used.

Finally, in the scheduling experiments where we investigated the performance of the different mechanism with respect to lowering peak demand and the peak-to-average ratio, the results are inconclusive as to which mechanism perform better. As can be seen in Figures 9.2 and 9.1, all of the mechanisms except the Uniform strategy solution showed very similar results. With the exception of the Uniform strategy, all the mechanisms showed *positive* results toward maintaining a low average peak, showing similar performance to the minimum extreme experiments. This seem to suggest that the potential for reducing the average peak was only limited by the average peak of the PowerNode demand, which represent the non-schedulable loads.

Distance-rule vs peak-shaving

Besides the two scheduling mechanisms, the experiments were also done using two different methods for calculating the day-ahead profile. These were included to see how the composition of the day-ahead profile would affect the performance of the scheduling mechanisms.

With respect to the fairness experiments, where the mechanisms are measured by how successful they are at charging the PHEVs in time, the proactive and reactive mechanisms were able to achieve an average battery capacity of 34,57% and 39,55% using the peak-shaving algorithm for constructing the day-ahead profile, compared to 48,20% and 44,53% with the distance-rule. These results were somewhat unexpected, as it was expected that by constructing the day-ahead profile so that the mechanisms would schedule the PHEVs closer to the time they would normally have been recharging. However, these results can be explained by considering the PHEV profiles and how the centralized scheduling mechanisms work. The centralized scheduling mechanisms will schedule the charging of PHEVs until the imbalances between the day-ahead profile and the real-time demand is minimized. Also, because the peak-shaving algorithm will move the demand to nearby hours, the extra demand in the day-ahead profile is likely to be placed during the evening hours. However, when we consider that the PHEV profiles that were designed for the experiments, they ensure that most of the PHEVs will be most active during the evening hours. This means that since the scheduler will try to schedule much of the PHEV demand for those hours when the PHEVs are most active, this means that the scheduling plans will be largely incompatible with the activity of the PHEVs when the peak-shaving algorithm is used. This means that the performance of the peak-shaving algorithm and the distance-rule algorithm, with respect to charging the PHEVs, seem to depend highly on how active the PHEVs are. Overall, these results seem to favor the distance-rule which will distribute most of the imbalances during night hours when the PHEVs are least active.

For both the scheduling and the stability experiments, the results did not seem to favor any one particular algorithm. Each of the algorithms performed well in maintaining a low average peak value, and neither of the algorithms in the centralized mechanisms had any problem ensuring the stability of the grid compared to the minimal extreme experiment.

Centralized scheduling vs decentralized scheduling

So far in this Section, we have compared the performance of the two different centralized scheduling mechanisms and how their performance differs depending on how the day-ahead profile is constructed. In this Section, we will discuss the performance between the centralized.

The best results of the centralized scheduling mechanisms came from both the proactive and reactive scheduling mechanism where the distance-rule was used. These mechanisms showed an average battery capacity at time of departure of 48,20% and 44,53% respectively. In contrast, for the decentralized mechanisms, the best performer was the Mixed strategy mechanism which had an average battery capacity of 56,37%, while the Uniform strategy managed an average battery capacity of 42,34%. This means that using the Mixed strategy, the decentralized mechanism was able to outperform both of the best performing centralized mechanisms. Even still, the Uniform strategy mechanism managed to obtain a performance close to the best performing centralized schedulers. This argues in favor of the decentralized scheduling mechanisms when considering charging of the

PHEVs and the fairness experiments. This is impressive considering the relative simplicity of the decentralized mechanisms compared to the centralized.

For the stability experiments, we find that both the Mixed and the Uniform strategy solution manage to preserve the stability of the grid relative to the minimal extreme experiments. However, compared to the centralized mechanisms, both the decentralized mechanisms show a higher degree of filtered demand by the Transformer agent, with an average of 116.4 kWh demand filtered in the Uniform strategy solution, and 126.4 kWh filtered in the Mixed strategy solution. This means that both decentralized mechanisms performed worst in these experiments with respect to Hypothesis 3.3. It should, however, be noted that in both decentralized mechanisms the filtered demand was relatively low, with only a daily average of 117.5 kWh and 131.0 kWh filtered for the Uniform and Mixed strategy solutions. This is a relatively low number compared to the number of PHEVs used in the simulations. It would be interesting to see how this would scale with the number of PHEVs.

9.5 Summary

In this section we have discussed the results that we obtained by running the experiments defined in Section 7. From these results, we found that all the mechanisms performed well in maintaining a low average peak value compared to the minimum extreme experiment. This seems to *verify* Hypothesis 3.1 and 3.2 that it *is* possible to obtain a peak-shaving effect by using a centralized scheduling mechanism in combination with a carefully calculated day-ahead portfolio, and that the same effect is possible with a decentralized mechanism (depending on the distribution).

While the improvement to PAR over the baseline was not large for any of the mechanisms, this seemed to be partly because of early morning peak hours, which the scheduling mechanisms were unable to affect since the demand during those hours were mostly due to non-schedulable loads. But also because the scheduling mechanisms could not the performance of the baseline experiment with respect to charging the PHEVs, leading to a lower overall average consumption. Some of this effect can perhaps be seen in Figures 9.5 and 9.1, where a high column in the chart over the average battery comparison seems to indicate a lower column in the chart over the average PAR values. The exception is the Uniform strategy mechanism, which because of its poor performance in lowering the average peak in Figure 9.2 is exempt from this pattern, while the other mechanisms perform almost identically in this respect.

For the centralized mechanism, it is also important to note that the composition of the day-ahead profile can make a significant impact on the results. This is apparent from the different results obtained by the centralized mechanism when the day-ahead profile was constructed using the peak-shaving algorithm, compared to when the day-ahead profile was constructed using the distance-rule algorithm. Especially when the average battery capacity is concerned. From the results, we found that if the centralized scheduler was given incentive to schedule the PHEV demand during night hours, the PHEVs were able to capitalize on this. While if the demand was scheduled during evening hours, when PHEV activity was higher, the PHEVs suffered some loss to the average battery capacity. For both other experiments, all mechanisms showed positive results, as all mechanisms were able to maintain a low average peak, and they were all able to maintain stability in the power grid.

10 Conclusion

In this thesis we have considered the impact that plug-in hybrid electric vehicles may have on the power grid, and how this may be handled in a way that does not compromise the stability of the grid or the end-user quality of comfort, while simultaneously scheduling PHEV loads so that the extra demand does not coincide with peak hours. This is a complex, multi-faceted problem that is not easily solved. To address this problem, we posed several questions at the beginning of this thesis, in Section 3.

To verify or falsify the hypotheses, we developed a Smart Grid simulator where we could experiment with a multi-agent system composed of four different types of agents: A Transformer agent, a PHEV agent, a PowerNode agent and a BRP agent, each with its own responsibilities. To help these agents solve the problems mentioned earlier, we considered several mechanisms which would help them achieve their goals. The mechanisms that were considered, were two centralized solutions where a centralized scheduler would help to schedule the PHEV demand, with the ultimate goal of lowering the PAR. In addition to these two centralized mechanisms, two decentralized mechanisms were also considered, where the PHEV agents generated their own charging strategies. In the Uniform strategy mechanism, the PHEV agents generated their charging plans based on a uniform distribution, while in the Mixed strategy mechanism the PHEV agents would generate their charging plans based on a weighted distribution generated from predictions it received from the BRP agent. Finally, to evaluate the performance of the different mechanisms, we defined a set of experiments in Section 7 that would help us to verify or falsify our hypotheses.

For the centralized mechanisms, the overall best results came from the reactive and proactive scheduling algorithms where the day-ahead profile had been constructed with the distance-rule algorithm. Both of these were able to obtain an average maximum peak similar to the minimal extreme experiment, the best case scenario, and while the positive results were limited with respect to the average PAR, this seemed to because of the early morning peak hours, which were largely caused by non-schedulable loads. But also because there seemed to be some cost associated with using a scheduling algorithm concerning the average PHEV battery capacity, which would lead to a lower overall average consumption. However, since the results for the average maximum peak was *positive*, this *verifies* the Hypothesis 3.1, that by carefully calculating the day-ahead algorithm, a centralized scheduling mechanism can be used to obtain a peak-shaving effect. Additionally, the Hypothesis was also *verified* for the Mixed strategy mechanism, which presented equally positive results compared to the centralized scheduling mechanisms in this regard.

Concerning the stability of the grid, both the centralized and decentralized mechanisms managed to obtain results similar to the minimum extreme experiment, presenting no signs of exceeding the transformer capacity more than the best case scenario. And while the Uniform- and Mixed strategy mechanisms presented some signs of demand filtering by the Transformer agents, this was relatively small amount and seemed to have little or no impact on the ability of the PHEV agents in the decentralized mechanisms to recharge their batteries. Especially when considering that the Mixed strategy mechanism was the overall best performer in this regard. This *verifies* Hypothesis 3.3, that a scheduling mechanism can be designed that will not compromise the stability of the grid.

While all the mechanisms performed well with respect to the average maximum peak and grid stability, most of the mechanisms showed signs of weakness concerning average battery capacity at time of departure. For the centralized mechanisms, this seemed to depend on how the day-ahead portfolio was calculated. If the imbalances were distributed during times when the PHEVs were most active then this seemed to negatively affect their ability to charge. For instance, the peak shaving algorithm would shift loads away from peak hours to hours that were relatively close. Because of the way the centralized scheduling mechanisms worked, this meant that the scheduler would try to fit PHEV demand into those adjacent hours. While this seemed like a good idea in theory, the experiments showed that for highly active PHEV profiles, if the PHEVs are not present in the grid then they miss this opportunity to charge. Additionally, neither the peak-shaving algorithm, nor the distance-rule algorithm considers whether there are enough PHEVs in the grid to be able to fully capitalize on the demand that is shifted. This means that there is a possibility that the much of the day-ahead imbalances could be concentrated over such short periods of time that the PHEVs are unable to capitalize on it. While this does *not* verify Hypothesis 3.4 that a scheduling mechanism can be designed that will be fair with respect to the end-user, it does not directly disprove it either. This is because the performance varied much depending on how the day-ahead portfolio was calculated, and it could be that a method exists which can calculate an optimal portfolio that will schedule the PHEV charging plans without compromising the average battery capacity. Also, the decentralized Mixed strategy mechanism presented *promising* results, outperforming all the centralized mechanisms in average battery capacity. Considering the relative simplicity of this mechanism, there could be significant potential for improvement for that mechanism.

When considering which mechanism that performs best, it is worth noting that the performance of the centralized mechanism is highly dependent on how the day-ahead profile is constructed. Similarly with the decentralized mechanism, the performance will vary much depending on how the strategy is generated. However, with the Mixed strategy solution discussed in this thesis, the decentralized mechanism proved to be a consistently good performer, performing equally well or better than the centralized mechanisms, thus *verifying* Hypotheses 3.2 and 3.5, that a decentralized mechanism can be used to obtain a peak-shaving effect, and that it can perform at least as well as a centralized scheduling mechanism. Additionally, by letting the PHEV generate the strategies themselves, the necessity of a centralized scheduler is removed. This increases the robustness of the system, as the dependency on a single component in the system is reduced. The most apparent downside to the Mixed strategy solution is that it can not be used to minimize the imbalances between the day-ahead profile and the real-time consumption, such as the centralized scheduling mechanisms can. However, assuming that the performance of charging the PHEVs in the Mixed strategy mechanism can be improved, the benefit of having a decentralized and robust mechanism may outweigh the benefit of reducing day-ahead imbalances.

11 Future work

In this section, some ideas for future work will be presented. These ideas were not implemented or experimented with in this thesis, either because it was necessary to limit the number of experiments, or because the ideas were outside the scope of the theme of the thesis.

While some positive and some negative results were presented for the different mechanisms that were used, it would be beneficial to perform further experiments on different methods for calculating the day-ahead portfolio, or for generating the strategy probability distributions. As became clear from the results, the different scheduling mechanisms varied in performance depending on whether the peak-shaving algorithm or the distance-rule algorithm was used. The same was also the case for the decentralized mechanism, depending on whether the Uniform strategy or the Mixed strategy distributions were used to generate the charging plan. And while the performance of these were analyzed and compared against each other, it is likely that these mechanisms can perform differently depending on what parameters are used, how the day-ahead portfolio is calculated and how the strategy probability distributions are determined.

Another typical concern in multi-agent systems is the incentive an agent has to be untruthful. For instance, in the centralized mechanisms discussed in this thesis, there is an obvious advantage for the PHEV agents to report their expected time of departure untruthfully. By reporting their expected time of departure to be earlier than the agents actually believes, it is more likely to be prioritized when the BRP creates the charging plans, or when the Transformer agent decides on which plans to filter. Note, however, that this concern does not arise in the decentralized mechanisms with respect to the scheduling of the charging plans, as each PHEV agent is itself responsible for generating its own charging plan. Assuming that the PHEV agent believes that the BRP agent is reporting its predictions about future energy demand truthfully, then the PHEV agent can do no better than to generate a plan according to the predictions it receives. Even so, the PHEV agent still has an incentive, also in the decentralized mechanism, to report its expected time of departure untruthfully to the Transformer agent. This decreases the likelihood that its charging plan will be filtered.

While the Mixed strategy mechanism performed better than the centralized reactive and proactive mechanisms with respect to charging the PHEV batteries and for maintaining a low PAR, it did show some signs of weakness with respect to the stability hypothesis as the Transformer agents had to resort to filtering the PHEV charging plans. While the filtered amount was low, it was more than compared to the centralized mechanisms. To address this concern, it could be interesting to modify the decentralized mechanism slightly. Currently, the PHEV agents interact with the BRP agent in acquiring predictions about the future energy demand. However, assuming that the Transformer agents were able to make predictions about the future energy demand in their responsible sub-grids, it could be possible to alter the mechanism so that the PHEV agents got their predictions from their parent Transformer agent instead. This would have the effect that their charging plans would be based on the future energy demand of their local sub-grid, and thereby the charging plans of the PHEVs would be targeted at maintaining a low PAR in their own local grid. The hypothesis is that, if the PHEV agents manage to coordinate their charging plans so as to maintain the stability of their local-grid, then it seems reasonable to infer that if the stability of all local sub-grids are maintained, so is the stability of the entire grid itself.

Finally, in Section 2, we discussed the existence of load-scheduling problems in Game Theory. Remember the proposition that every load scheduling game admits at least one Nash equilibrium. Further, given any optimal allocation $A : [n] \to [m]$, there exists a finite sequence of improvements steps that the agent can make in order to find a Nash equilibrium. Assuming that the problem of load scheduling PHEVs in the Smart Grid is reducible to a load scheduling game in Game Theory, whereby; instead of allocating $[n]$ tasks to $[m]$ machines, the allocation is $[n]$ units of demand over $[m]$ units of time, with the additional constraint that only one allocation is possible per unit time. Then for any allocation such as the one made with the Mixed strategy mechanism discussed in this thesis, there exists a finite sequence of improvement steps that the PHEV agents can make to improve their charging plans until they find a *pure* Nash equilibrium. This means that it could be possible to extend or modify the decentralized mechanism discussed in this thesis so that the PHEV agents – after choosing the initial allocation – can iteratively try to improve upon this allocation, by following a sequence of improvement steps in order to maximize their average battery capacity while still maintaining a low average peak-to-average ratio.

References

Alglib. Spline interpolation. http://www.alglib.net/interpolation/spline3.php, Apr. 2012.

S. M. Amin and B. F. Wollenberg. Toward a Smart Grid. *IEEE power & energy magazine*, (october):34–41, 2005.

N. Boucké and T. Holvoet. Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a Smart Grid. *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, (section 3):803–810, 2011.

H. Farhangi. The Path of the Smart Grid. *IEEE power & energy magazine*, (february), 2010.

S. W. Hadley and A. Tsvetkova. *Potential Impacts of Plug-in Hybrid Electric Vehicles on Regional Power.* Number January. 2008. ISBN 1800553684.

E. D. Haukedal. SimCar: Open source simulator. https://github.com/eirikdal/SimCar, Apr. 2012.

B. V. Herbruggen and B. V. Zeebroeck. Emissions of road traffic in belgium. *Policy*, 2006.

IEEE PES Multi-agent Systems Working Group. Application Areas. http://ewh.ieee.org/mu/pes-mas/agent-technology/applications/, Dec. 2011.

A. Mohsenian-Rad, V. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *Smart Grid, IEEE Transactions on*, 1 (3):320–331, 2010.

N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*, volume 53. Cambridge University Press, 2007. ISBN 9780521872829. doi: 10.1145/1785414.1785439. URL http://portal.acm.org/citation.cfm?doid=1785414.1785439.

NVE. Avanserte måle- og styringssystemer, oppsummering av høringsuttalelser og endelig forskriftstekst. Technical report, 2011.

Office of Electricity Delivery and Energy Reliability. Smart Grid. http://energy.gov/oe/technology-development/smart-grid, Oct. 2011.

S. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based control for decentralised demand side management in the smart grid. pages 2–6, 2011.

P. Rowles. Demand response and demand side management whatâĂŹs the difference? http://www.energyadvantage.com/blog/2010/02/demand-response-demand-side-management-whatâĂŹs-difference/, Nov. 2011.

S. Russell and P. Norvig. *Artificial intelligence: a modern approach*. Prentice hall, 2nd edition, 2010.

R. Shaw, M. Attree, T. Jackson, and M. Kay. The value of reducing distribution losses by domestic load-shifting: a network perspective. *Energy Policy*, 37(8): 3159–3167, Aug. 2009. ISSN 03014215. doi: 10.1016/j.enpol.2009.04.008.

L. C. Stavrogiannis. Electricity Tariff Design and Implementation for the Smart Grid by MSc in Artificial Intelligence. *Change*, (October), 2010.

The Climate and Pollution Agency. Klimautslippsreduksjoner mot 2020. `http://www.klimakur.no/`, Apr. 2012.

U.S. Department of Energy. Argonne Leads DOE's Effort to Evaluate Plug-in Hybrid Technology. `http://www.transportation.anl.gov/phev/index.html`, Apr. 2012.

S. Vandael, N. Boucké, T. Holvoet, and G. Deconinck. Decentralized demand side management of plug-in hybrid vehicles in a smart grid. *status: accepted*, 2010.

S. Vandael, T. Holvoet, and G. Deconinck. A multi-agent system for managing plug-in hybrid vehicles as primary reserve capacity in a Smart Grid. In *Proceedings of the Second International Workshop on Agent Technologies for Energy Systems (ATES 2011)*, number figure 1, pages 3–4, 2011.

F. Wernstedt, P. Davidsson, and C. Johansson. Demand side management in district heating systems. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, volume 5, page 272. ACM, 2007. ISBN 9788190426275.

Wolfram Alpha. Wolfram|Alpha: Computational Knowledge Engine. `http://www.wolframalpha.com/`, June 2012.

M. Wooldridge. *An introduction to multiagent systems*. Wiley, 2nd edition, 2002.

A Article

Load-scheduling and PHEVs in the Smart Grid

Eirik Daleng Haukedal, Boye Annfelt Høverstad, Axel Tidemann, Pinar Öztürk
Norwegian University of Science and Technology
Trondheim, Norway
eirik.haukedal@gmail.com, hoverstad@idi.ntnu.no, axel.tidemann@gmail.com, pinar@idi.ntnu.no

***Index Terms*—multi-agent systems, smart grid, demand-side management, load scheduling, PHEV.**

***Abstract*—Load peaks can have a negative impact on the stability of the power grid and maintenance costs for transmission and generation companies. Currently, increasing use of plug-in hybrid electric vehicles (PHEV) further proliferates the problem because charging patterns are expected to coincide with peak demand hours, especially the afternoon peak hours.**

To avoid the problem that increasing PHEV demand will further aggravate peak demand hours, we have investigated two different mechanisms for scheduling the charging PHEVs; one centralized mechanism, where a centralized scheduler creates the charging plan for the PHEVs, and one decentralized mechanism, where the PHEVs create their own charging plans by randomizing over a strategy probability distribution.

We found that both mechanisms were able to schedule the PHEVs so that aggravating peak load was avoided, but that it came at a cost of the PHEVs capacity to recharge. However, the decentralized mechanism showed promising results that it could be possible to create a multi-agent mechanism that is able to charge PHEVs by avoiding peak hours, and still be fair to the end-user.

I. Introduction

According to official Norwegian estimates [1], plug-in hybrid electric vehicles (PHEVs), electric vehicles and hydrogen-based vehicles will account for 5% of the total car population in Norway by 2020, and that figure is expected to rise up to 26% by 2030. While this is arguably a positive development from a climate perspective, one of the anticipated challenges in the future smart grid is how to efficiently handle the extra load associated with charging the growing number of PHEVs. This is mainly a challenge, not because there is an insufficient overall capacity to accommodate the extra charge, but because the extra demand resulting from PHEVs recharging their batteries is expected to coincide with times at which demand is already at its highest, namely peak hours [2].

Considering the challenges of how to address the future demand caused by PHEVs, multi-agent systems (MAS) may prove to be a promising candidate technology. A multi-agent system consists of intelligent agents interacting in an environment. The agents can be computer software modules, human beings, or anything else capable of autonomous and rational actions. Multi-agent technology can be used to build systems that are scalable, fault-tolerant, secure and easy to reason about.

Addressing the problem of peak-load, we have investigated two different approaches to load-scheduling where the real-time consumption of energy is influenced by controlling the charging of PHEVs; a centralized mechanism and a decentralized mechanism.

In the first approach, with the centralized scheduling mechanism, we have investigated an approach to load-scheduling where the real-time consumption of energy is influenced by controlling the charging of PHEVs. If we at any time know 1) the current, real-time demand for energy in the grid, 2) the energy traded on the forward market (day-ahead) at that time, and 3) information about the state of the PHEVs in the grid, then we have sufficient information to schedule the charging of the PHEVs in a way that will help to minimize the difference between the real-time energy demand and the energy traded for on the day-ahead market.

In the second approach to load-scheduling, we have investigated a decentralized mechanism where no centralized schedulers are used. The idea is that that the PHEVs will generate their charging strategies on their own accord, by randomizing over a probability distribution that they create themselves. The important thing to consider with this mechanism is how this probability distribution is created. For this, we have considered two possible methods. One in which the probability distribution is uniform, and another in which the distribution is created the PHEV agent, by interacting with a central agent that provides the PHEV agents with predictions about the future energy consumption in the grid.

II. Related work

The simulator that we developed, is based upon the works presented in [3]–[5]. While they also investigated scheduling mechanisms for the charging of PHEVs, which minimizes the difference between the day-ahead portfolio and real-time consumption, our simulator extends upon their work by considering how such scheduling mechanisms can also be used to influence peak-shaving. Although it is not used yet, the PHEV model in our simulator is capable of learning its usage-patterns from experience. By learning, it is possible for the system to better schedule PHEVs when there is uncertainty about when they will be used.

In [6], McArthur et al. argue why multi-agent systems is a methodology well suited to the power engineering domain. The publication is a result of work done by IEEE Power Engineering Society's Multi-Agent Systems (MAS) Working Group. They identify several key areas where agents are applicable: monitoring and diagnostics, distributed control, modeling and simulation and protection schemes for the network. In [7] they offer guidance and recommendations on the design and implementation of MAS specifically for power engineering. They stress the importance of standards that allow interoperability between systems, such as: 1) using FIPA standards for communication, 2) employing a common "upper" ontology for interoperation between MAS, 3) employing a specific design methodology for MAS, where knowledge engineering and task decomposition allow for ontology and agent design. However, they find it too hard to give a recommendation on the design of the agent anatomy. This highlights the difficulty and importance of good agent design, however the main point of the two papers are how well suited MAS are for power engineering.

[8] describes a system in a micro grid context, that includes distributed energy sources as well as storage devices. Such a micro grid can either operate connected to the main grid or be isolated. A their laboratory micro grid, each connected device also has a corresponding agent in the network. They use a simple centralized algorithm for calculation of the difference between produced and consumed power to correct any discrepancies. In their laboratory experiments, loads are turned on and off and the system responds accordingly. The system is thus designed to be reactive, i.e. responding to changes immediately, without any inherent prediction or learning capabilities. Their work represent no intelligent way of load scheduling, but is an example of the use of MAS in a power control setting.

[10] focus on reducing the peak-to-average (PAR) ratio by using energy consumption scheduling (ECS) devices. The ECS is implemented in smart meters, and the goal is to perform autonomous demand side management in a neighborhood. A distributed algorithm finds the optimal energy consumption for each subscriber. They focus on incentives for reducing PAR, such as lower utility charges. The scenario consists of one energy source (e.g. substation) shared by several subscribers, each having an ECS with communication capabilities. They show that given an appropriate pricing scheme, the subscribers can reach a Nash equilibrium where the overall system performs better, and they also pay less individually. The authors formulate the problem mathematically as a convex optimization problem. The algorithm shows good results, however the system does not learn or predict any load profiles, the assumption is that the energy consumption is pre-determined and set by the subscriber.

In a similar manner to [10], [11] investigate the effect of sharing load profiles among users versus not sharing (for privacy reasons), and device a distributed algorithm (i.e. a game) for users in the first case, and a stochastic strategy in the latter case. The key is a dynamic pricing scheme that encourages users to achieve a desirable load profile, from the perspective of the utility. Their scenario is the same as in [10]. Both approaches reduce PAR, the distributed game approach being the best. What separates the results from [10] is the ability of users to make inferences based on the instantaneous load in the case where the other users do not share their demands. This is benchmarked against a situation where there is no communication between the users.

III. Models

To test our hypotheses, we have developed an open-source simulator [14], built around a hierarchy of models and agents as seen in Figure 1. It is composed of a Balancing Responsible Party agent (BRP), Transformer agents (TRF), PHEV agents, and PowerNode agents. The first three agents are similar to the agents described in [3]. It is the responsibility of the BRP agent to minimize the imbalances between the day-ahead profile and the real-time energy consumption, while the TRF agents and the PHEV agents are responsible for ensuring that demand does not exceed the capacity of its associated transformer, and that the PHEV is adequately charged in time for use. The PowerNode agent in the simulator represent any node in the grid capable of producing or consuming energy, and which is not a PHEV. A typical interaction between the agents can be seen in Figure 2, where the agents are negotiating a charging schedule for the PHEV agent.

A. PowerNode model

The PowerNode model in the simulator represents anything in the grid that can consume or produce power, and which is not a PHEV. For instance, it can represent a household, consuming power, in which case it will contribute a negative flow of energy to the power grid. Or it can represent a distributed energy facility, contributing a positive flow of energy.

In terms of the complexity of the model, the main thing to consider is how the model decides what the flow of energy will be. In this simulator, the flow of energy is determined by historical data collected from smart meters installed in real houses. In essence, this means that each PowerNode model in the simulator will be represented by one household, or equivalent, from the dataset used. The flow of energy at time t in the simulator is then determined by the flow of energy at time t in the dataset, for that household, offset by the initial recording time for the dataset.

B. PHEV model

The most important and complex model in the simulator is the PHEV model. For this model, we tried to implement the PHEV model so that it reflects and captures the uncertainty of human nature. In addition to this, the PHEV model also incorporates models information such as average discharge rate for the PHEV while driving and average recharge rate when connected to the grid. To model its behavior in the grid, we have designed a set of PHEV profiles, where a PHEV profile contains a set of probability distributions that determines whether a PHEV leaves/disconnects from the grid. Mapped to each probability distribution, is information about the duration of that trip, as well as information about how much of the time was actually spent driving. In Figure 3, the four PHEV profiles used for the experiments can be seen, showing a Suburban profile, a Homeworker profile, a City profile and a Commuter profile. These profiles were mostly designed to stress test the system during peak hours and challenge the ability of the PHEVs and the BRP to schedule the charging of the PHEVs. And while some considerations have been made to make the profiles reflect realistic behavior, these considerations are not based on any scientific analysis.

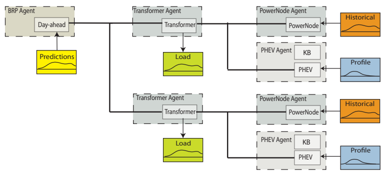

Fig. 1. An overview of the agent-model hierarchy, showing the different agents in the system and the main channels of communication.

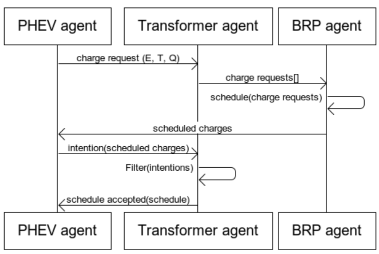

Fig. 2. An interaction showing the PHEV agent requesting a schedule from the BRP agent.

C. Transformer model

While the PowerNode- and PHEV models in the simulator will provide electricity throughput, the Transformer model will act as a constraint as to the capacity of that throughput in the different parts of the grid. This is to prevent the BRP agents or PHEV agents from scheduling the charging of the PHEVs in such a way that it violates the capacity constraints of the transformers. This means that the Transformer model will need a defined value for its maximum capacity of energy that it can handle at any given time. For these values, we have used standardized values for power grids in Norway, with capacities in the range of 220-240V for low-voltage transformers, 11kV for high-voltage transformers and 66kV for regional voltages.

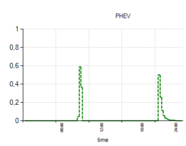

(a) Commuter

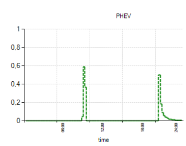

(b) City

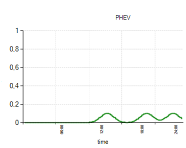

(c) Homeworker

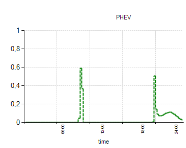

(d) Suburban

Fig. 3. The different, discretized PHEV profiles used in the simulator.

D. Grid model

The Grid model is the model for the structure and hierarchy of the grid itself. For the purpose of the experiments, the grid was modeled so that the capacity of the grid would be challenged during times of high demand. Using the defined values for transformer capacity mentioned earlier, the grid was generated by finding the peak load for each of the power nodes from the historical data, assigning power nodes to each transformer until the sum of the peak loads exceeded the capacity of the transformers by a scaling factor. After a low voltage transformer had been assigned sufficient power nodes, the low voltage transformers were assigned to high voltage transformers by following the same principle. This process was iterated until the set of power nodes were empty.

E. Determining the day-ahead profile

The last model in the simulator is the Day-ahead model, which is the model used by the centralized mechanisms when scheduling the loads of the PHEVs. The model is essentially a prediction about the next day energy consumption, which we will call the day-ahead profile. To determine the day-ahead profile, a two-step process was used. First, the baseline was determined as the expected contribution from all of the non-schedulable loads. For this, we used the historical data sampled from houses with smart meters installed. While this has the effect that the day-ahead profile will be perfectly accurate with respect to the power nodes, it is the demand from the PHEVs that we are ultimately interested in.

To determine the expected contribution from the PHEVs, we implemented and tested two different methods. The first method estimated the PHEV contribution to demand by running the simulation in advance, without using any of the scheduling mechanisms. This let the PHEVs charge on demand, whenever it was possible to so. The other method calculated the expected contribution from the PHEVs statistically by using known values for the probabilities that a PHEV would leave at a given time, plus the duration of the travel during which it would be discharging its batteries, the average rate with which it would be discharging, and the rates with which it would recharge upon return.

After the baseline and the expected contribution of the PHEVs was determined, these were added together in order to yield the day-ahead profile. This could have been used as a basis for the day-ahead portfolio in the simulator. However, considering that we also desire to influence the reduction of peak-loads, an additional step was performed on the portfolio, so that the PHEV loads were shifted away from the peak-hours. This was also done by investigating two slightly different methods. In one, the baseline and the PHEV contributions were added together, after which one of the two algorithms was used. One algorithm was a peak-shaving algorithm, loosely based on the Kohonen algorithm for self-organizing maps, and the other was a distribution algorithm,

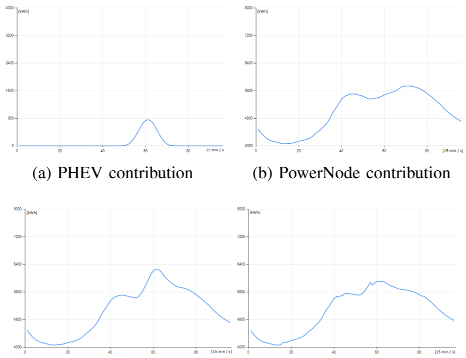

(a) PHEV contribution (b) PowerNode contribution

(c) PHEV + PowerNode (d) PHEV + PowerNode – after peak-shaving

Fig. 4. Sample images from the simulator showing the different steps for calculating the day-ahead profile

where the PHEV contributions were distributed onto the day-ahead portfolio based upon a distance rule. Common to both algorithms is that they shift loads away from peak-hours while the overall volume of the portfolio is preserved. Both algorithms result in a portfolio that can be used to schedule PHEV loads to help reduce peak-load.

In Figure 4, the 4 steps in the process of calculating the day-ahead profile are shown with 4 graphs illustrating the result of each step in the process. (4a) shows the expected contribution from all the PHEVs in the simulator, while (4b) shows the baseline with the contribution from all of the PowerNodes. In (4c) the baseline and expected PHEV contribution are added together, leading to an aggravation of the afternoon peak-load, and finally, the end result in (4d) shows the result of applying the Kohonen-inspired, peak-shaving algorithm.

IV. Multi-agent system

There are four types of agents in the multi-agent system that we designed for these experiments, a PHEV agent, a Transformer agent, a BRP agent and a PowerNode agent. Each of these agents are responsible for their respective models. The PHEV agent will act and react upon changes made to its PHEV model, the Transformer agent will act and react upon changes to its Transformer model, and similarly for the PowerNode agent. The model for the BRP agent is the day-ahead model and its predictions about the electricity demand, which it will use, for instance, in scheduling the PHEVs charging plans in any of the centralized mechanisms. Also, depending on the mechanism used, each of these agent will also interact with each other and their environment/models differently.

A. PHEV agent

In both mechanisms, the PHEV agent have two main responsibilities. Its primary responsibility is to maximize the battery levels within time of departure. Its secondary responsibility is to achieve its primary responsibility in the most socially economical way possible. How it fulfills these responsibilities, however, depend upon which mechanism is used. In the decentralized mechanisms, the PHEV agent chooses its charging strategies by itself, while in the centralized mechanism it defers control of creating the charging plan to a centralized scheduler.

B. Transformer agent

The transformer agent has one main responsibility, to ensure that the capacity of the physical transformer it is assigned to is never exceeded. To enforce this constraint, it can filter demand messages made by PHEV agents that choose to participate in the centralized scheduling mechanisms. After the PHEV agent have negotiated a charging plan with the BRP agent, or if the PHEV agent have already done so previously, the PHEV agent will send a Demand message containing their current energy demand to the Transformer agent. The Transformer agent will collect these Demand messages from all of its connected PHEV agent and PowerNode agents. If the sum of the demand of all the Demand messages exceeds its capacity, then it will filter away Demand messages that are coming from PHEV agents until its capacity constraint is satisfied. Note, however, that it only filter demand messages sent by the PHEV agents. It does not filter messages from the PowerNode agents, as these are assumed to be non-deferrable loads.

C. Balancing responsible party

Depending on which mechanism is used, the BRP agent has one of two responsibilities. In the centralized mechanisms, the BRP agent is responsible for scheduling the PHEV charging plans for all the PHEV agents in the grid. If any of the two decentralized mechanisms is used instead, then the BRP agent is responsible for providing its predictions about the future energy demand to the PHEV agents.

D. Centralized mechanisms

In the centralized mechanisms, the PHEV agent coordinates its charging plan by interacting with the BRP agent and the Transformer agent. The BRP agent communicates with its Transformer agent, which it can announce its intention to charge to. In turn, the Transformer agent which is responsible for keeping peak load within the constraint of its transformer, forwards the aggregated intentions of all its PHEV agents to the BRP agent which performs load scheduling on the intentions. After having scheduled the load profiles, the BRP agent announces its charging plans back to the Transformer agent. Depending on whether the constraints of the transformer is satisfied, the Transformer agent will either request a new plan from the BRP agent or forward the message back to the intended PHEV agents.

1) The scheduling algorithms: How the PHEV agents charging plans are scheduled in the centralized mechanisms depend upon which scheduling algorithm that is used. The BRP agent can choose between two centralized scheduling algorithms: A reactive scheduling algorithm and a proactive scheduling algorithm [3]. The goal of both algorithms is to schedule the charging plans so that the overall difference between the dayahead predictions and the realtime energy demand is as low as possible.

Reactive scheduling: In the reactive profile, energy is balanced in a way that attempts to maintain a perfect balance between the day-ahead profile and the actual demand, for as long as possible. In Algorithm 1, pseudocode for the scheduling algorithm is shown [3], showing how energy is reserved for as long as the prediction of a given time is less than the dayahead consumption for that timeperiod, and while there is still energy left over to assign.

Algorithm 1 Pseudocode for a reactive scheduling strategy, adapted from the algorithm described by [3]

```
function CREATE-PLAN
    energyLeft ← sum(intentions)

    for each t ∈ now ... ttd do
        if sum(prediction) < sum(dayahead) & energyLeft > 0 then
            plan_t ← χ.rate
            prediction_t ← prediction_t + χ.rate
            energyLeft ← energyLeft − χ.rate
        end if
    end for

    return plan
end function
```

Algorithm 2 Pseudocode for a proactive scheduling strategy, adapted from the algorithm described by [3]

```
function CREATE-PLAN
    energyLeft ← sum(intentions)

    while energyLeft > 0 do
        if sum(dayahead) − sum(prediction) > 0 then
            t ← arg max_t(prediction(t) − dayahead(t))
            plan_t ← χ.rate
            prediction_t ← prediction_t + χ.rate
            energyLeft ← energyLeft − χ.rate
        end if
    end while

    return plan
end function
```

Proactive scheduling: In this profile, energy is distributed in a way such that charges are assigned to times at which the imbalance is greatest while there is a positive difference in the dayahead quantity compared to predicted consumption. Otherwise, energy is assigned to the times at which the imbalance is smallest. This is intuitively because it is desirable to minimize the average distance between the dayahead quantity and the real-time predictions. By assigning energy to the time of largest imbalance while the difference is positive, actual real-time consumption is brought closer to the dayahead quantity. If the difference is negative this means that the overall consumption has exceeded the net dayahead quantity. This means that wherever the charge is placed, it will have a negative impact, but it will do the least harm at the time at which the imbalance is the least. This is illustrated in the pseudocode shown in Algorithm 2.

E. Decentralized mechanisms

In the decentralized mechanisms, the idea is that a peak-shaving effect can be obtained by letting PHEVs choose charging strategy on their own. They do this by randomizing over a probability distribution, where the probability at a given time determines the likelihood that a PHEV chooses to charge during that time. This process is illustrated by pseudocode in Algorithm 3.

Algorithm 3 Randomizing over a distribution

```
let remaining = battery.max - battery.current

let plan = [..]

for t ∈ now .. ttd do
    if rand(0,1) < strategy_t then
        plan_t ← rate
        remaining ← remaining - rate
    end if
end for
```

1) Uniform strategy: For this, we investigate two different strategies; one in which PHEVs choose charging times by randomizing over a uniform distribution, under the constraint that they will try to charge fully if they can. In the other strategy, the PHEVs will ask the BRP agent about its predictions on total electricity demand for the next six hours. The PHEVs will then generate a strategy, assigning probabilities of whether it will charge in each of the fifteen minutes leading up to the next six hours. With the uniform strategy, the probability distribution is $\frac{1}{2}$ for each possible timeslot.

2) Mixed strategy: In the mixed strategy, the probabilities will be weighted favorably towards times when total demand is low. The algorithm for generating the mixed strategy can be seen in algorithm 4, and Figure 5 shows a graphic illustration of the process. The end-result is a distribution that is an inverted distribution based on the normalized predictions about the future energy consumption. These predictions come from the BRP agent, and each time the PHEV agent recognizes that it is connected to the grid and that its current battery capacity is below maximum capacity, it will ask the BRP agent for its predictions about the future energy consumption based on a window from the present time and until the PHEV agent believes it will disconnect from the grid again. Also, once the PHEV has generated its charging strategy, randomizing over the distribution, it reports it charging plan back to the BRP agent. The BRP agent will then update its own predictions with the charging plan that it got from the PHEV agent. This prevents the PHEV agents from all generating their charging plans based on the same distribution.

V. EXPERIMENTS

To test the mechanisms and our hypotheses, we designed a set of experiments that would measure the performance of

Algorithm 4 Generating a mixed strategy charging plan

let predictions = request(predictions, ttd)

let max = $\arg\max predictions$
let min = $\arg\min predictions$

for each $x \in predictions$ **do**
    $strategy_t \leftarrow 1.0 - (x - min)/(max - min)$
end for

let plan = create-plan(strategy)

inform(brp, plan)

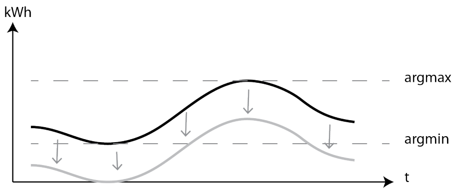

(a) Scaling and normalizing the predictions by $(x - \arg\min)/(\arg\max - \arg\min)$

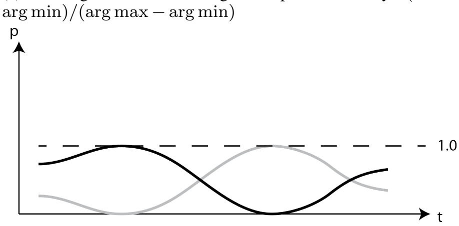

(b) Inverting the normalized predictions

Fig. 5. An illustration of how the mixed strategy is generated

the different mechanisms with respect to stability, fairness and peak-to-average ratio. The experiments consisted of running the simulations for the length of 100 days, where each combination of mechanism and day-ahead algorithm were used. This gives a total number of 18 different experiments, with 6 combinations of day-ahead algorithm and mechanisms, where each was tested with respect to stability, fairness and peak-shaving. The experiments were done on a grid composed of 1756 power nodes and 616 PHEVs. The number of power nodes come from the historical data that was used for the simulations, while the number of PHEVs is reflected from the predictions that by 2030, the ratio of PHEVs in Norway will be 30%. This assumes an average number of 1 vehicle per house. In addition, the time period for the sampled data that was used was from April 1st until July 10th, 2006. This interval selected has hopefully limited the most significant influences of seasonal changing on the experiments.

For the composition of the PHEV profiles used in the experiments, each PHEV was randomly assigned one of the four profiles in Figure **??** with a uniform likelihood of 25% for each of them. The battery capacity of the PHEVs is based on the Chevrolet Volt, with a large 16 kWh lithium-ion battery, and each PHEV was given an average charging rate of 2.5 kWh per hour, and a discharging rate of 8 kWh per hour driving. These values were selected to give an average recharging rate of 6,4 hours, and an average total driving time of 2 hours. Further, the model assumes that the PHEV will charge its batteries fully to maximum capacity, and that it will drive entirely on its batteries until the capacity is depleted. While these may be naive assumptions about how PHEVs work in practice, note that these experiments are designed to stress test the power grid, meaning that we are mostly interested in demand and consumption by the PHEVs when connected to the grid.

To compare and contrast the experiments, we ran two further experiments: The baseline experiment, in which no scheduling mechanisms were used, and the minimal extreme experiment, in which the simulator were run without PHEVs present. These two experiments represent polar extremes in a sense that the baseline represent the worst case scenario and the minimal extreme represent a lower case scenario. Either experiment should not produce a worse average maximum peak than the baseline, otherwise the mechanism would work counter to its purpose. Also, no experiment should produce a lower maximum peak than the minimal extreme, which would be an indication of a flaw with the simulator.

VI. Discussion

For the scheduling experiments, we were interested in the performance of the different mechanisms with respect to lowering the peak-to-average ratio and the average maximum peak. The results of these experiments can be seen in Figure 6 and 7. Figure 6 shows a box plot graph of the peak-to-average ratio. As can be seen, the results showed only minor differences to the maximum, minimum and average values for the PAR. The reason for this is illustrated by Figure 7, where it is apparent that most of the mechanisms are equal to the performance of the minimal extreme with respect to the average maximum peak value. This seems to suggest that any further performance improvement is limited mostly by the average maximum peak of the power nodes. The exception to this case is the Uniform strategy solution, which performs worse than all the mechanisms, including the baseline experiment.

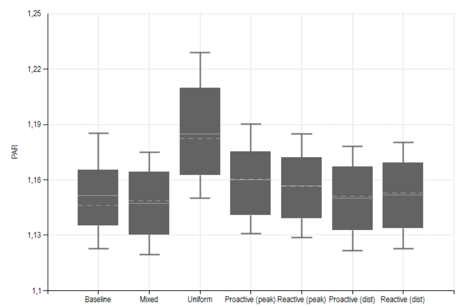

Fig. 6. A boxplot graph showing the average PAR for each of the different mechanisms. Lower is better.

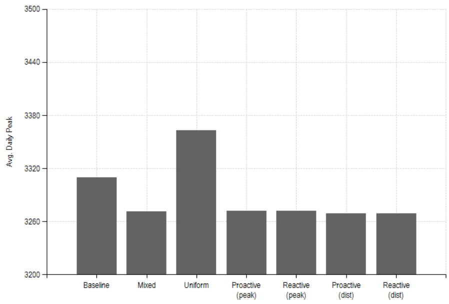

Fig. 7. A graph showing the average daily peak for each of the different mechanisms. Lower is better.

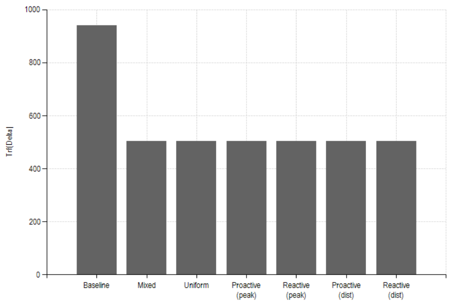

Fig. 8. A graph showing the average capacity exceeded over all transformers, for each of the different mechanisms. Lower is better.

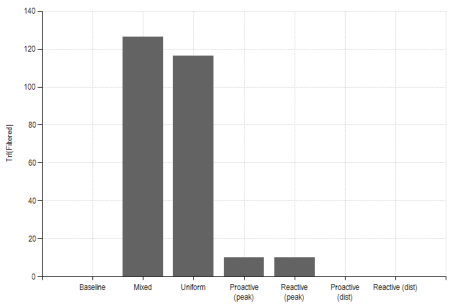

Fig. 9. A graph showing the average PHEV demand filtered, for each of the different mechanisms. Lower is better.

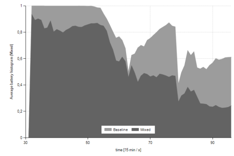

(a) Mixed strategy

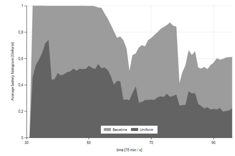

(b) Uniform strategy

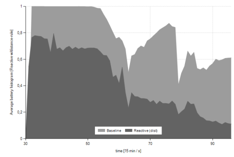

(c) Reactive scheduling, distance-rule

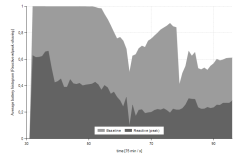

(d) Reactive scheduling, peak-shaving

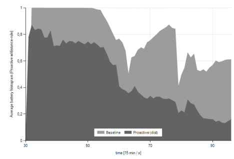

(e) Proactive scheduling, distance-rule

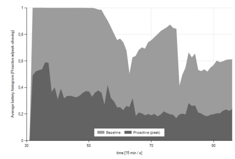

(f) Proactive scheduling, peak-shaving

Fig. 10. These graphs show histograms over the average battery capacity at time of departure for each of the different mechanisms. Each mechanism is contrasted with the results from the baseline experiment for comparison.

While the results were somewhat inconclusive as to which mechanism was best at reducing the average maximum peak and the average PAR, the results were somewhat more convincing in the fairness experiments. In these experiments we compared how the mechanisms performed with charging the PHEV batteries. To do this, we measured the average ratio of battery capacity at time of departure (BCTD) for each of the PHEVs. This is illustrated in 10, which shows a histogram over the average BCTD from 07.30 and until 24.00. The time-period between 24.00 and 07.30 was omitted from the graphs, as no PHEV profiles were given a positive probability for departing during those hours. From this figure, it is apparent that the mechanism that performed best compared to the baseline was the Mixed strategy solution, which showed an average BCTD of 56,37%. This is compared to the baseline, which had an average of 77,61%. It is also apparent that the average BCTD drops steadily from morning until evening. This is consistent with the usage patterns defined by the PHEV profiles, where the PHEVs have most of the night to recharge until the morning hours, leading to a high BCTD. However, from the afternoon and on, when the PHEVs are more active, it is more challenging for the mechanisms to recharge them in time, leading to a steady drop in the BCTD as the likelihood of activity increases.

Finally, for the stability experiments, where we interested in observing the performance of the different mechanisms with respect to the average demand that exceeded the Transformers capacity, or when applicable, the amount of demand filtered by the Transformer agent from the PHEV charging plans. The results for these experiments are illustrated in Figure 9, which show the amount of PHEV demand filtered by the Transformer agent, and in Figure 8, which show the amount of energy exceeding the transformer capacity. The distinction between these variables is that, while both measure excess demand, to an extent, the filtering mechanism is only employed in the scheduling mechanism. This means that if demand exceeds the capacity for any of the transformers, the Transformer agent may intervene by filtering away some demand from the PHEV charging plans. However, only when there is no PHEV demand left to filter, will the excess demand be registered as exceeding the capacity.

From Figure 8 experiments showed that all combinations of the different mechanisms performed equally well to the minimal extreme experiments which is the experiment where no PHEVs are present in the grid. This means that all of the mechanisms were able to prevent the PHEV demand from aggravating the already strained power grid during times of high demand. This is consistent with the observations made in the scheduling experiments, which showed that all of the algorithms managed to maintain an average maximum peak that was identical to the minimal extreme experiment. However, while all of the mechanism showed flawless performance with respect to capacity exceeded, the decentralized mechanisms had to resort to filtering some of the PHEV charging plans, as can be seen in Figure 9. Ultimately, this has no effect on the stability of the system, as the worst case scenario would be that the Transformer agents would have to filter all of the PHEV charging plans. This would lead to a result that would be identical to the minimal extreme experiment, where no PHEVs are used, meaning it would only lead to the PHEVs not being able to charge their batteries. But since the Mixed strategy solution was the mechanism that showed best performance in charging the batteries, it seems fair to assume that the benefit of this charging strategy more than cancels the negative result in demand filtered.

VII. Conclusion

In summary, the mechanism that overall proved to be the best performer was the decentralized, Mixed strategy solution, where the PHEVs generate their own charging strategies based on predictions they receive from the central BRP agent. This does not preclude the centralized mechanisms, however, as the results show that how the day-ahead profile is constructed can have great impact on the efficiency of the mechanism. From the results we found that while the centralized mechanisms performed well with reducing peak-to-average ratio and maximum peak, they showed some weakness in charging the PHEVs. The best results of the centralized mechanisms came when the day-ahead profile was created from the distance-rule algorithm, where the largest imbalances in the day-ahead profile was placed during the night time. In the peak-shaving algorithm, where the imbalances were placed in relative proximity to its expected original time-location, it seemed as if the highly active PHEV profiles caused the scheduling mechanism to be unable to capitalize on these imbalances. This further seems to stress the importance of how the day-ahead profile is calculated in the centralized mechanism, and it is possible that better results can be obtained with other algorithms than described in this article.

However, while the centralized mechanisms showed some difficulty in scheduling based on how the day-ahead profile was created, the Mixed strategy solution in the decentralized mechanism showed positive results in all the experiments. This is a promising result for decentralized scheduling mechanisms, as they have several advantage over centralized mechanisms. Firstly, if the entire mechanism depend on a centralized scheduler, then the system is vulnerable if the BRP agent goes down. Secondly, scheduling algorithms are often computationally demanding so optimal scheduling may not be plausible to do in real-time. Compare this to the decentralized mechanism, where the only dependency is the agent that the PHEV agent receives its predictions from, and even if this agent should become unavailable for some time, the PHEV agent may even default to predictions of its own, removing the need for interaction between agents altogether. However, removing the exchange of information between agents would lead the agents blind to choices made by other agents, which would most likely lead to a less optimal result depending on how each PHEV agent would generate their charging plans. Considering all of this, the decentralized mechanism seem like a promising solution to load-scheduling.

References

[1] The Climate and Pollution Agency, "Klimautslippsreduksjoner mot 2020," http://www.klimakur.no/, Apr. 2012.

[2] S. W. Hadley and A. A. Tsvetkova, "Potential impacts of plug-in hybrid electric vehicles on regional power generation," *The Electricity Journal*, vol. 22, no. 10, pp. 56–68, 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S104061900900267X

[3] N. Boucké and T. Holvoet, "Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a Smart Grid," *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, no. section 3, pp. 803–810, 2011.

[4] S. Vandael, T. Holvoet, and G. Deconinck, "A multi-agent system for managing plug-in hybrid vehicles as primary reserve capacity in a Smart Grid," in *Proceedings of the Second International Workshop on Agent Technologies for Energy Systems (ATES 2011)*, 2011, pp. 3–4.

[5] S. Vandael, N. Boucké, T. Holvoet, and G. Deconinck, "Decentralized demand side management of plug-in hybrid vehicles in a smart grid," 2010.

[6] S. McArthur, E. Davidson, V. Catterson, A. Dimeas, N. Hatziargyriou, F. Ponci, and T. Funabashi, "Multi-agent systems for power engineering applications—part i: Concepts, approaches, and technical challenges," *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1743–1752, Nov. 2007.

[7] ——, "Multi-agent systems for power engineering applications—part ii: Technologies, standards, and tools for building multi-agent systems," *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1753 –1759, Nov. 2007.

[8] J. Oyarzabal, J. Jimeno, J. Ruela, A. Engler, and C. Hardt, "Agent based micro grid management system," in *International Conference on Future Power Systems*, Nov. 2005.

[9] K. Mets, T. Verschueren, W. Haerick, C. Develder, and F. De Turck, "Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging," in *Network Operations and Management Symposium Workshops*, April 2010, pp. 293–299.

[10] A.-H. Mohsenian-Rad, V. Wong, J. Jatskevich, and R. Schober, "Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid," in *Innovative Smart Grid Technologies*, Jan. 2010, pp. 1–6.

[11] S. Caron and G. Kesidis, "Incentive-based energy consumption scheduling algorithms for the smart grid," in *First IEEE International Conference on Smart Grid Communications*, Oct. 2010, pp. 391–396.

[12] T. Logenthiran, D. Srinivasan, and A. M. Khambadkone, "Multi-agent system for energy resource scheduling of integrated microgrids in a distributed system," *Electric Power Systems Research*, vol. 81, no. 1, pp. 138–148, 2011.

[13] M. Hommelberg, C. Warmer, I. Kamphuis, J. Kok, and G. Schaeffer, "Distributed control concepts using multi-agent technology and automatic markets: An indispensable feature of smart power grids," in *IEEE Power Engineering Society General Meeting*, June 2007, pp. 1–7.

[14] Haukedal, Eirik D., "SimCar: Open source simulator." https://github.com/eirikdal/SimCar, Apr. 2012.

B User Manual

B.1 Files and folders

The simulator uses the following files

- /SimCar/SimCar/data/grid.cfg
- /SimCar/SimCar/data/profiles.cfg

The simulator uses the following folders:

- /SimCar/SimCar/data/*
- /SimCar/SimCar/data/img/*
- /SimCar/SimCar/data/log/*
- /SimCar/SimCar/data/log/latex/*
- /SimCar/SimCar/data/log/experiments/*
- /SimCar/SimCar/data/interpol/*

The *data* folder contains the grid- and PHEV profile configuration files. In the *img* folder, the simulator will store images from experimental runs, where images are sampled for each day of the simulations. The images include consumption graphs over the entire grid, and sample battery- and transformer graphs for selected PHEVs and Transformers. In the *log* folder, and subfolders, log files from the experimental results are recorded. Both in text format, LaTeX format and raw binary data collected from the experiments. In the *interpol* folder, the historical data is stored as .dat files, where each .dat file represent the historical used for each of the PowerNodes in the simulator.

B.2 Scripts

In addition to the User Interface, the simulator comes with 4 complimentary scripts: test_powerprofiles.fsx, make_tree.fsx, parse_realtime_data.fsx, parse_results.fsx.

test_powerprofiles.fsx

This script contains unit tests on calculating the expected PHEV demand

make_tree.fsx

This script generates a power grid tree. It has four variables that can be set: phev_ratio, peak_ratio, high_trf, low_trf.

The first variable, phev_ratio, specifies the ratio of PHEV to PowerNodes, while peak_ratio is the factor used when assigning PowerNodes to the Transformers. For instance, a peak_ratio of 1.5 means that PowerNodes will be assigned to the

Transformer until the sum of the maximum peak value of the PowerNodes exceeds 1.5 times the maximum capacity of the Transformer.

The last two variables, high_trf and low_trf specifies the maximum capacity of the low-pass transformers and high-pass transformers.

parse_realtime_data.fsx

This script provides some additional charting functionality that can be used to plot the comparison charts used in the in the Discussion section of the thesis.

parse_realtime_data.fsx

This script was used for parsing the historical data used for the simulator to a format that the simulator could use. The script also performs the Akima spline interpolation on the data.

B.3 User Interface

The graphical user interface, seen in Figure B.1, is divided into three main sections: The main charting area that contains the charting components, seen in Figure B.2, the logging section, seen in Figure B.3, and the sidebar, seen in Figure B.4. The main charting area is only meant to provide visual feedback that the simulator is running, while quantitative results are stored in folder on the disk. Configuration of the experiments are done in the sidebar, and textual information about the progress of the simulations can be seen in logging section.

Sidebar

The sidebar section of the user interface contains controls for setting the day-ahead algorithm, as seen in Figure B.5a, where the day-ahead portfolio can be set to calculated using the Peak-shaving algorithm, the Distance-rule algorithm, None, Superposition and Uniform. And in Figure B.5b, the drop down box for selecting which mechanism to use can be seen. The choices include: Reactive, Proactive, Uniform, Mixed and None. For selecting how to calculate the expected PHEV demand, see Figure B.5c. The choices include: Expected, Simulated and None.

In addition to selecting which combination of mechanism and day-ahead algorithms, it is possible to set the different parameters associated with these. Figure B.5d shows where to set the θ variable in the distance-rule algorithm discussed in Section **??** of the thesis. In addition, Figure B.5e shows where to set the different parameters associated with the peak-shaving algorithm discussed in Section 5.4.3 of the thesis, and Figure B.5f shows where to set the size of the PHEV learning window discussed in Section 6.2.1 of the thesis.

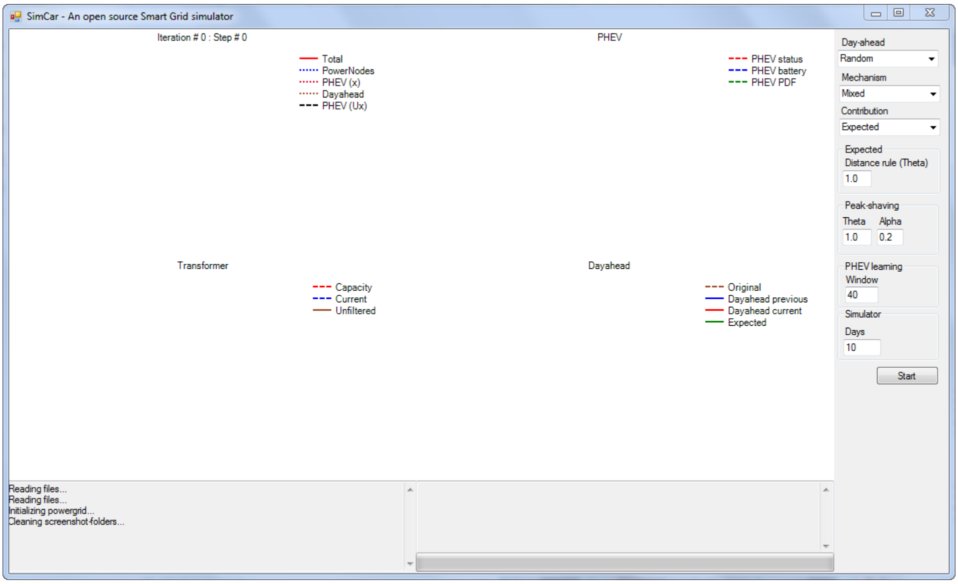

Figure B.1: This screenshot shows the graphical user interface for the simulator that was developed for this thesis.

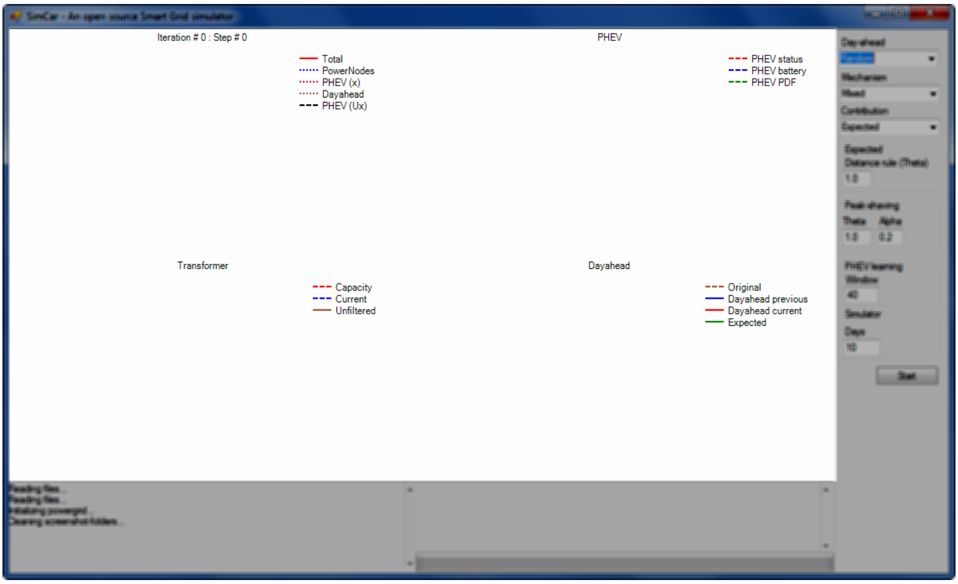

Figure B.2: A screenshot of the graphical user interface for using the simulator. In this screenshot, the different charting areas are highlighted. These are meant to provide the users with graphical feedback to the user of the progress, and are not meant to be interacted with. The results of the simulations are rather stored on drive for easier access.

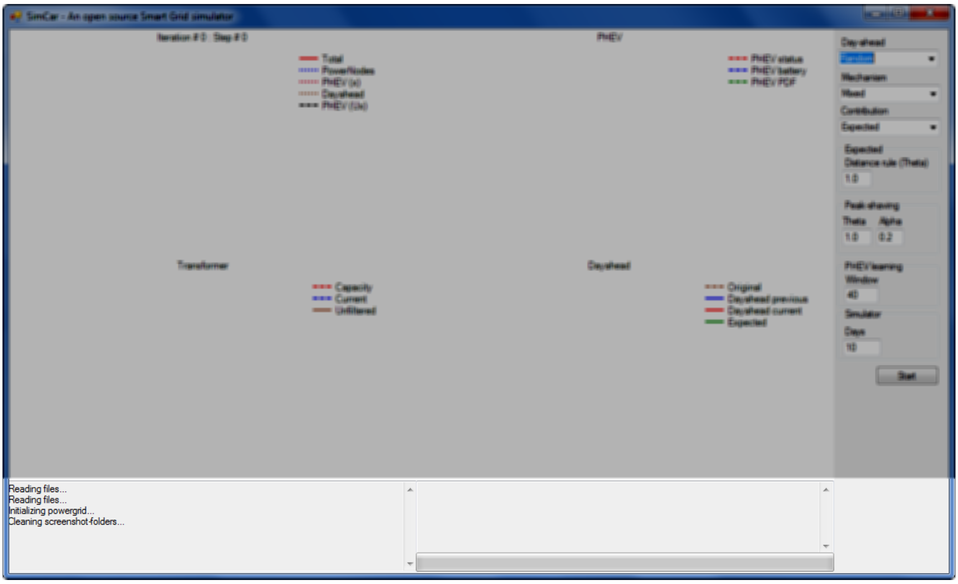

Figure B.3: This screenshot shows the simulator textbox to the bottom left of the screen and the debugging textbox. In the simulator textbox, the daily results will be recorded. The debugging textbox will show any information which is relevant to the simulations, its progress and success or failure, but which is not relevant to the results of the experiments.

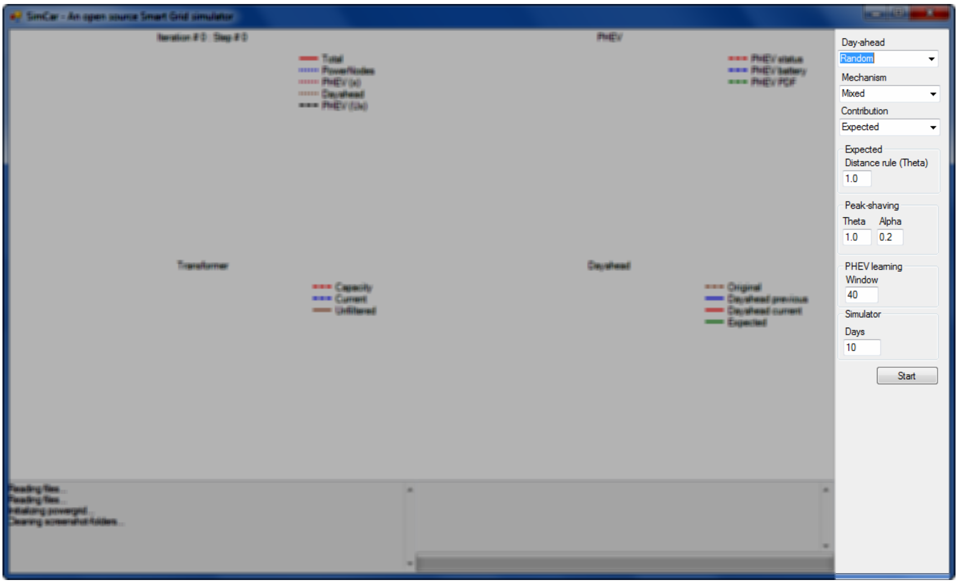

Figure B.4: On the far right in the user interface, the sidebar with the different controls can be found

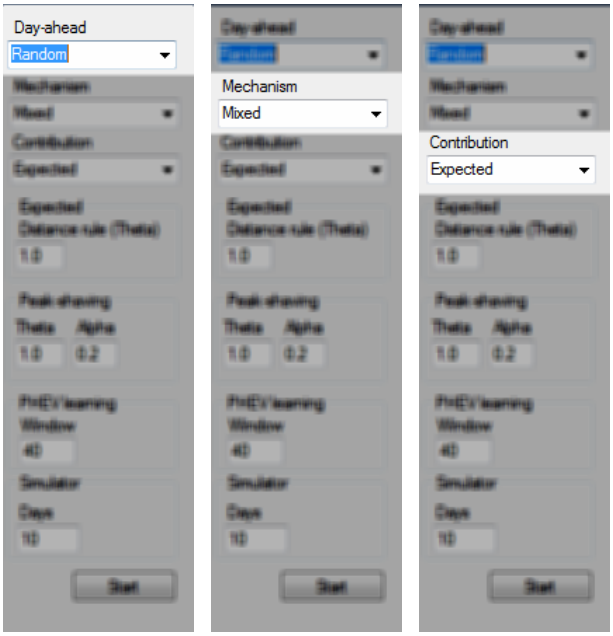

(a) Day-ahead algorithm (b) Scheduling mechanism (c) Expected PHEV demand

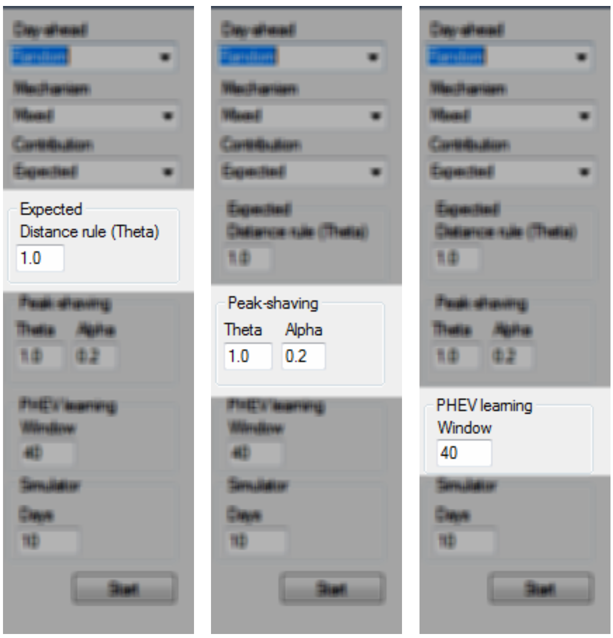

(d) Distance-rule parameters (e) Peak-shaving parameters (f) PHEV learning window

Figure B.5: A snapshot of the simulator user interface, showing the different configuration options available through the simulator.

C Message Protocol

Type	**Contents**	**Description**
Charge	U_χ - uncharged capacity ttl - time to departure $rate$ - charging rate	Used by the PHEV agents to announce their charge intentions
Demand	E_d - current demand by the agent	Used by the PHEV agents and PowerNode agents to announce their demand to the Transformer agent
Consume	E_c - accepted demand to agent	Used by the Transformer agents to inform the PHEV agents and PowerNode agents how much they can consume at this time
Intentions	$\vec{E}$ - list of charge intentions	Wrapper message for the Transformer agents to collect all charge intentions from all their children nodes before propagating upwards
Dayahead	$d(t)$ - dayahead function	Used by the simulator to update the BRP agent with new dayahead predictions
Prediction	$p(t)$ - prediction function	Used by the simulator to update the BRP agent with realtime predictions
Request Predictions	-	Used by the PHEV agents in the Mixed strategy mechanism to request a window of predictions
Request Model	-	Used by the simulator to request the agents model
Request Dayahead	-	Used by the simulator to request the dayahead model from the BRP agent
Predictions	$\vec{p}_t$ - predictions up to time t	Used by the BRP agent when responding to RequestPredictions from the PHEV agents
Model	$\vec{m}$ - a model	Used by the agents when responding to RequestModel
Strategy	$\vec{\chi}_t$ - charging strategy up to expected time to departure	Used by the PHEV agents in the Mixed strategy mechanism when informing the BRP agent about their final strategy
Update	t - current time	Synchronization message from the simualtor to the agents. Informs the agents that time has advanced by a tick
Kill	-	Sends a termination request to the agents
Reset	-	Sends a reset request to the agents
Schedule	$sched$ - the schedule function	Used to send the scheduling algorithm to the BRP agent.

Table C.1: Full list of messages available to the agents in the simulator