

ENERGY OPTIMIZATIONS FOR SMART BUILDINGS AND SMART GRIDS

A Dissertation Presented

by

ADITYA K. MISHRA

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

July 2015

College of Information and Computer Sciences

© Copyright by Aditya K. Mishra 2015

All Rights Reserved

ENERGY OPTIMIZATIONS FOR SMART BUILDINGS AND SMART GRIDS

A Dissertation Presented

by

ADITYA K. MISHRA

Approved as to style and content by:

Prashant Shenoy, Chair

Jim Kurose, Member

Ramesh Sitaraman, Member

David Irwin, Member

Ting Zhu, Member

James Allan, Chair
College of Information and Computer Sciences

ACKNOWLEDGMENTS

The process of PhD is challenging in many ways, and it would have been difficult to endure without vital support from everyone around me. This is a heartfelt attempt to express my gratitude to everyone who helped me in this journey.

I am thankful to my advisor Prof. Prashant Shenoy for his support and guidance. This thesis would not have been completed without his help. I have learned some key research skills from him: how to pick and define a problem, set its scope, investigate, and present my work to the world. He has an excellent knack for presenting scientific work and his inputs on my presentations have been invaluable to my learning. I especially admire the efforts he makes to bring grant money for supporting his students.

I have enjoyed interacting with my committee members and have learned a great deal in the process. I would like to thank them for the same. I was fortunate to be able to work with Prof. Jim Kurose on my synthesis project. He always provided me broader perspective about my work and systems research in general. I am also thankful to him for his excellent career advice and encouraging words. I am glad to have had the opportunity to work with Prof. Ramesh Sitaraman. He always provided me very relevant and insightful comments and pointers about my research problems. In spite of his busy schedule he was very kind to speak to me on Skype, phone, and provide the much needed urgent inputs. I am also grateful to Prof. David Irwin for his help and contributions in several research projects, especially in the beginning. Further, I am thankful to Prof. Ting Zhu for thought provoking discussions and his inputs on my work.

My sincere thanks to my family members especially my parents, wife (Bhavana Dalvi), and in-laws. I am grateful to my parents for their unconditional love and support. Their love and support has been vital for surviving the arduous doctorate journey. I am grateful to my wife, Bhavana, for being alongside me in this journey through graduate school. She patiently listened to my ideas, questions, concerns and always gave me the most pertinent advice. In spite of her own research commitments, she was available whenever I needed her. For that, no amount of gratitude can suffice. I am equally grateful to my in-laws for encouraging me to continue my research work during the post marriage years. They always lifted my spirits which helped me give my best at research.

I would also like to thank my friends Abhigyan, Abhinav, Abhishek, Amulya, Anand, Andres, Anita, Boulat, Junghee, Kaituo, Navin, Olaitan, Pan Hu, Pengyu, Rahul, Rick, Sandeep, Srinivasan, Sunil, Supriya, Swagatika, Upendra, Vikas, Xiaozheng for their support and help during these years. They provided a stimulating environment making my graduate student life enjoyable and truly memorable. Finally my heartfelt thanks to Karren, Leeanne, and the rest of the staff at the College of Information and Computer Sciences for their invaluable help and assistance over the years.

ABSTRACT

ENERGY OPTIMIZATIONS FOR SMART BUILDINGS AND SMART GRIDS

JULY 2015

ADITYA K. MISHRA

B.E., RAJIV GANDHI TECHNICAL UNIVERSITY, INDORE, INDIA

M.Tech., INDIAN INSTITUTE OF TECHNOLOGY BOMBAY

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Prashant Shenoy

Modern buildings are heavy power consumers. For instance, of the total electricity consumed in the US, 75% is consumed in the residential and commercial buildings. This consumption is not evenly distributed over time. Typical consumption profile exhibits several peaks and troughs. The peakiness, in turn, dictates the electric grid's generation, transmission and distribution costs, and also the associated carbon emissions.

This thesis discusses challenges involved in achieving the sustainability goals in buildings and electric grids. It investigates building and grid energy footprint optimization techniques to achieve the following goals: 1) making buildings energy efficient, 2) cutting building's electricity bills, 3) cutting utility's costs in electricity generation and distribution, 4) reducing carbon footprints, and 5) making the aggregate electricity consumption profile grid-friendly.

In this thesis, we first design SmartCap, a system to enable homes flatten their consumption/demand by scheduling background loads (such as A/Cs, refrigerator), without causing user discomfort and without direct user involvement. Demand flattening facilitates aggregate peak reduction, which in turn enables grids to 1) reduce carbon emissions, and 2) cut installation and operational costs. Our results demonstrate that SmartCap can decrease the average deviation from mean power by over 20% across all periods with “high” deviation, thereby flattening the “peaky” demand. Next, we present SmartCharge, an intelligent battery charging system that shifts a building’s electricity consumption to off-peak periods by storing low-cost energy for use during high-cost periods, without active user involvement. We show that SmartCharge can typically save 10-15% in bills and can reduce the grid-wide peak demand by up to 20%. We then extend SmartCharge to GreenCharge, which integrates on-site renewables in a building’s electricity consumption. Our experiments show that GreenCharge can cut user electricity bills up to 20%. After GreenCharge, we investigate the use of large-scale distributed energy storage at buildings throughout the grid to flatten grid demand, while 1) maintaining end-user incentives for storage adoption at grid-scale, and 2) ensuring grid stability. We design PeakCharge, an online peak-aware charging algorithm to optimize the use of energy storage in the presence of a peak demand surcharge. Empirical evaluations show that total storage capacity required by PeakCharge to flatten grid demand is within 18% of the capacity required by a centralized system. Finally, we examine the efficacy of employing different combinations of energy storage technologies at different levels of the grids distribution hierarchy to cut electric utility’s daily operational costs. We present an optimization framework for modeling the primary characteristics of various energy storage technologies and important tradeoffs in placing different storage technologies at different levels of the distribution hierarchy. We show that by employing hybrid

storage technologies at multiple levels of the distribution hierarchy, utilities can reduce their daily operating costs due to distributing electricity by up to 12%.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	iv
ABSTRACT	vi
LIST OF TABLES	xiv
LIST OF FIGURES	xv
 CHAPTER	
1. INTRODUCTION	1
1.1 Background and Motivation	1
1.2 Contributions	3
1.2.1 Summary of Contributions	4
1.2.2 SmartCap	5
1.2.3 SmartCharge	6
1.2.4 GreenCharge	7
1.2.5 Scaling Distributed Energy Storage	7
1.2.6 Integrating Energy Storage in Electricity Distribution Networks	8
1.2.7 Thesis Outline	9
2. BACKGROUND AND RELATED WORK	10
2.1 Green Computing	10
2.2 Smart Buildings	11
2.3 Internet of Things in Smart Homes	12
2.4 Variable Electricity Pricing Plans	13
2.5 Energy Storage in Buildings and Electric Grids	15

3. SMARTCAP: FLATTENING PEAK ELECTRICITY DEMAND IN SMART HOMES	17
3.1 Introduction and Motivation	17
3.2 Related work	20
3.3 Background and Problem Formulation	21
3.4 Load Analysis and Observations	23
3.4.1 Interactive vs. Background Loads	23
3.4.2 Interactive Variability	26
3.4.3 Background Variability	27
3.5 Load Scheduler	29
3.5.1 Load Controllers	30
3.5.2 Scheduler	31
3.6 Prototype: Design and Implementation	35
3.7 Evaluation	37
3.7.1 Simulation Results	39
3.7.2 Impact of Electric Vehicles	40
3.7.3 Testbed Results	41
3.8 Conclusion	42
4. SMARTCHARGE: CUTTING THE ELECTRICITY BILL IN SMART HOMES WITH ENERGY STORAGE	43
4.1 Introduction and Motivation	43
4.2 Related Work	47
4.3 SmartCharge Architecture	49
4.4 SmartCharge Algorithm	52
4.4.1 Potential Benefits	52
4.4.2 Problem Formulation	54
4.5 ML-based Demand Prediction	57
4.6 Experimental Evaluation	60
4.6.1 Household Savings	62
4.6.2 Grid Peak Reduction	65
4.6.3 Lab Prototype Results	67
4.7 Cost-Benefit Analysis	69
4.7.1 Return-on-Investment	69

4.7.2	Distributed vs. Centralized	73
4.8	Conclusion	74
5.	GREENCHARGE: MANAGING RENEWABLE ENERGY IN SMART BUILDINGS	76
5.1	Introduction and Motivation	76
5.2	Related Work	79
5.3	GreenCharge Architecture	81
5.3.1	Network Communication and Sensing	82
5.3.2	Market-based Electricity Pricing	84
5.4	GreenCharge Algorithm	87
5.4.1	Potential Benefits	88
5.4.2	Problem Formulation	90
5.5	Predicting Consumption and Generation	93
5.5.1	ML-based Demand Prediction	94
5.5.2	Predicting Energy Harvesting from Weather Forecasts	96
5.6	Experimental Evaluation	97
5.6.1	Household Savings	99
5.6.2	Grid Peak Reduction	102
5.7	Cost-Benefit Analysis	104
5.7.1	Return-on-Investment	105
5.7.2	Distributed vs. Centralized	109
5.8	Conclusion	110
6.	SCALING DISTRIBUTED ENERGY STORAGE FOR GRID PEAK REDUCTION	111
6.1	Introduction and Motivation	111
6.1.1	Contributions	115
6.2	Related Work	117
6.3	Overview and Approach	118
6.3.1	PeakCharge Architecture	119

6.3.2	The Storage Adoption Dilemma	122
6.3.3	An Optimal Approach	125
6.4	Scalable Design	126
6.4.1	The Effect of a Peak Demand Surcharge	127
6.4.1.1	Benefits	127
6.4.1.2	Drawback	130
6.5	Peak-aware Charging	131
6.5.1	Optimizing for the Peak	132
6.5.2	Optimizing for Peaks and Variable Rates	134
6.5.3	Summary	136
6.6	Evaluation	138
6.6.1	Grid-scale Effects	140
6.6.2	Consumer-scale Effects	143
6.7	Conclusion	144
6.A	Appendix	146
7.	INTEGRATING ENERGY STORAGE IN ELECTRICITY	
	DISTRIBUTION NETWORKS	148
7.1	Introduction and Motivation	148
7.2	Related Work	151
7.3	Background	152
7.3.1	Electric Grid	152
7.3.1.1	Distribution Network	153
7.3.1.2	T&D Losses in the Grid	153
7.3.2	Electric Utility's Generation Costs	154
7.3.3	Energy Storage to Lower Utility's Costs	155
7.3.3.1	Energy Storage Technologies	156
7.3.3.2	Energy Storage Characteristics	157
7.4	Problem Statement	159
7.5	Energy Storage Provisioning and Control Framework	160
7.5.1	Inputs	161
7.5.2	Optimization Problem Formulation	163

7.6	Evaluation	167
7.6.1	Experimental Setup and Methodology	167
7.6.2	Potential Savings from Storage	170
7.6.3	Longer-term Savings	177
7.7	Conclusion	178
7.A	Appendix	179
8.	SUMMARY AND FUTURE WORK	182
8.1	Thesis Summary	182
8.2	Future Work	184
	BIBLIOGRAPHY	186

LIST OF TABLES

Table	Page
3.1 In the summer, background loads in our home account for 59% of the total energy consumption.	24
4.1 Average prediction error (%) over 40 day sample period for SVM with different kernel functions.	57
4.2 Estimated cost breakdown for installing SmartCharge’s supporting infrastructure.....	71
5.1 Average prediction error (%) over 40 day sample period for SVM with different kernel functions.	93
5.2 Estimated cost breakdown for installing SmartCharge’s supporting infrastructure.....	106
6.1 Parameter definitions for linear program.	146
7.1 ESD Parameters [40] [35] [111] [94].....	157
7.2 Experiment Parameter Values.	167
7.3 Storage Configuration and Placement (Long-term Contract, Medium capex): (Savings(\$/day), Storage costs (\$/day)). Total cost without storage is \$21.5k/day.	177
7.4 Optimization framework notations	181

LIST OF FIGURES

Figure	Page
3.1 A graphical depiction of a SmartCap-enabled home.	22
3.2 The power consumption of interactive loads is highly variable throughout the day. As expected, peak power consumption occurs around mealtimes in the morning, early afternoon, and early evening.	25
3.3 Power data for example interactive loads. Occupant behavior, which is not readily predictable, determines when these loads draw power.	28
3.4 Power signatures for four background loads in our home. The on-off period varies with environmental conditions, and is not regular.	29
3.5 A depiction of slack in our refrigerator’s simple on-off control loop. The compressor turns on once the internal temperature reaches an upper threshold, and turns off once it reaches a lower threshold.	31
3.6 A background load scheduler is capable of flattening demand, but must account unpredictable interactive and background loads.	33
3.7 Example of how LSF flattens demand.	35
3.8 LSF decreases the absolute average deviation from the mean power (with no scheduling) on the vast majority of days (91%), as well as over peak 4-hour periods with mid-range and high-range deviations.	36
3.9 Load duration curves for a typical summer day with and without scheduling when using an electric vehicle.	40
3.10 Power usage with and without LSF scheduling using using our smart home testbed with real background loads over a 4-hour period.	42

4.1	The marginal cost to generate electricity increases as utilities dispatch additional generators to satisfy increasing demand. Data from [55].	46
4.2	A depiction of SmartCharge’s architecture, including its battery array and charger, DC→AC inverter, power transfer switch, energy/voltage sensors, and gateway server.	49
4.3	Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.	51
4.4	Example from January 3rd with and without SmartCharge using Illinois prices from Figure 4.3.	53
4.5	Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.	59
4.6	Average dollar savings per day for both real-time and TOU prices in our case study home.	60
4.7	Average percentage savings for both real-time and TOU prices in our case study home.	61
4.8	SmartCharge’s savings as a function of the charging rate for a 12kWh storage capacity.	62
4.9	Varying the average electricity price (a) and the peak-to-off-peak price ratio (b) impacts savings.	63
4.10	Additional savings (in % and \$) from sharing 12kWh and 24 kWh between homes.	65
4.11	With 22% of homes using SmartCharge, the peak demand decreases by 20% (a) and demand flattens significantly (b).	66
4.12	Our UPS-based prototype reduces peak usage by 69% when using a few common appliances.	68
4.13	Amortized cost per kWh as a function of depth of discharge.	72
4.14	Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [93].	73

4.15	SmartCharge’s projected yearly expense and savings assuming recent battery advancements.	74
5.1	A depiction of GreenCharge’s architecture, including its battery array and charger, DC→AC inverter, solar and/or wind energy sources, power transfer switch, energy/voltage sensors, and gateway server.	82
5.2	Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.	85
5.3	Example solar harvest data from a day in August.	87
5.4	Example from January 3rd with and without GreenCharge using Illinois prices from Figure 5.2.	89
5.5	Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.	96
5.6	Average dollar savings per day for both SmartCharge and GreenCharge in our case study home.	97
5.7	Average percentage savings for both SmartCharge and GreenCharge in our case study home.	98
5.8	SmartCharge’s and GreenCharge’s savings as a function of the charging rate for a 24kWh storage capacity.	99
5.9	Varying the average electricity price (a) and the peak-to-off-peak price ratio (b) impacts savings.	100
5.10	Additional savings (in % and \$) from sharing 12kWh and 24 kWh between homes.	101
5.11	With 25% of homes using GreenCharge, the peak demand decreases by 22.5% (a) and demand flattens significantly (b).	102
5.12	Demand flattening with Net Metering.	103
5.13	Amortized cost per kWh as a function of depth of discharge.	108
5.14	Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [93].	109

6.1	Prior switch-based architectures do not significantly lower an individual building’s peak demand. Figure from [78].	119
6.2	PeakCharge architecture, which includes a battery array capable of programmatically controlling the rate of discharge wired in parallel with the grid.	121
6.3	The model we use in our simulator of the marginal cost to generate electricity as demand increases. The fitted function we use is based on scaled data from a recent FERC report [55].	123
6.4	An idealized depiction of the cycle that causes the storage adoption dilemma.	124
6.5	Load oscillations in our simulated microgrid, in presence of day-ahead real time pricing.	126
6.6	Instantaneous and average grid demand for 194 homes in our trace.	128
6.7	While 12kWh of energy storage is capable of shifting only a fraction of demand to the low price period, it is more than enough to completely flatten the demand from Figure 6.1.	129
6.8	Generation cost savings compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans. Zoom-in of generation cost savings for peak-aware algorithm (c).	137
6.9	Peak reduction as a percentage compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans.	138
6.10	Time series of aggregate grid demand for TOU (a) and DART (b) pricing for both without energy storage and using our peak-aware algorithm with each home having energy storage.	138
6.11	Generation cost savings (a) and grid peak reduction (b) as we vary the size of each home’s energy storage capacity.	139
6.12	Increasing the peak demand surcharge prevents rebound peaks in the grid by incentivizing consumers to flatten their demand.	144

6.13	Percentage peak reduction (a), percentage cost savings (b), and dollar cost savings (c) for an individual home using our peak-aware algorithm as the home’s energy storage capacity varies using our peak-aware algorithm under a peak-demand surcharge.	145
7.1	Typical electricity distribution hierarchy.....	154
7.2	Illustrative graph depicting distribution hierarchy	161
7.3	Composition of capex and peak penalty costs for Long-Term Contract.....	167
7.4	Electricity demand on a representative weekday.	169
7.5	Savings from deploying lead-acid battery storage at (a) single level and (b) multiple levels under the long-term contract pricing plan.	171
7.6	Savings from deploying hybrid storage technologies at a single-level for low and high cap-ex costs under long-term contract.	172
7.7	Multi-Level Hybrid ESDs v/s: (a) Single-level Hybrid ESDs and (b) Multi-level lead-acid batteries under long-term contract.	173
7.8	Savings under day-ahead pricing for multi-level lead acid batteries and multi-level hybrid storage.	174
7.9	Aggregate Peak Reduction from Lead-Acid only and Hybrid ESDs for CapEx(Medium).	175
7.10	Average daily savings for March (Day-Ahead).....	178

CHAPTER 1

INTRODUCTION

Modern buildings are heavy power consumers. For instance, of the total electricity consumed in the US, nearly 75% is consumed in the residential and commercial buildings [7]. This consumption is not evenly distributed over time. Typical consumption profile exhibits several peaks and troughs. The peakiness, in turn, dictates the electric grid's generation, transmission and distribution costs, and also the associated carbon emissions.

Therefore, making buildings energy efficient is an important problem. This thesis discusses challenges involved in achieving the sustainability goals in buildings and electric grids. We investigate building and grid energy footprint optimization techniques to achieve following goals: 1) making buildings energy efficient, 2) cutting building's electricity bills, 3) cutting utility's costs in electricity generation and distribution, 4) reducing carbon footprints, and 5) making the aggregate consumption profile grid-friendly.

1.1 Background and Motivation

A significant fraction of global energy consumption is contributed by the buildings. Buildings in the United States (US) account for nearly 40% of total US energy consumption [7]. Approximately 70% of this consumption is from electricity. Besides, this electricity, due to environmental concerns, is largely generated at "dirty" power plants far from populated areas. As a result, a large fraction of electricity is lost in transmission. For example, roughly 6.5% of total electricity in the US is lost in

transmission and distribution [51]. Due to their substantial demand and associated losses, making buildings energy efficient is an urgent need for restricting global energy consumption.

Making individual buildings efficient will help bring down the total energy use. However, it is not the only way of reducing losses and making the grid sustainable. Since transmission losses are proportional to the square of the current, a reduced peak demand can significantly reduce these losses. Further, electricity generation costs are affected disproportionately by the peaks because of the fuel costs. Operation and capital costs associated with electricity distribution are also affected by the peak demands. Besides, peak demand dictates the installed generation, and transmission and distribution capacity of the grid. Therefore, peak reduction will result in reduced losses, decreased fuel requirements, smaller capacity grid installations, and reduced utility operational costs.

In an attempt to tap into significant benefits of peak shaving, utilities are transitioning from flat pricing to variable time-of-use or peak-load electricity pricing plans [41], [8], [31], [99]. These dynamic pricing plans penalize/reward the electricity consumption by raising and lowering the prices during peak and off-peak periods, respectively. The variable pricing is in use for residential customers at several places. For instance, Illinois requires utilities to provide residential customers the option of using hourly electricity prices based directly on whole-sale prices [101]. Ontario charges residential customers based on a time-of-use scheme with three price tiers each day [87].

Variable electricity pricing provides users an opportunity to cut their electricity bills by shifting their consumption to low demand off-peak periods. Thereby reducing the operational costs for the grids by reducing the peaks. Yet, inciting users to respond and participate by cutting or shifting their demand to make it grid-friendly involves various challenges. Shedding unnecessary loads or shifting usage to off-peak hours

requires active user involvement and may require a change in living patterns, which may cause user inconvenience. Moreover, there needs to be enough economic incentive for the users to get involved and stay motivated. Automating the process can avoid direct user involvement, but it presents several challenges of its own. The automation should not adversely affect home-appliance health or lifetimes. Most importantly, when adopted at large scale, it should make the aggregate grid consumption profile sustainable and help cut costs for electric utilities.

Thesis Statement: *Energy footprint optimization in buildings and distribution networks can make them energy efficient, cut end user’s electricity bills, and cut electric utility’s expenses in electricity distribution. These optimizations are achievable without active user involvement and user inconvenience, while making aggregate electricity consumption profile grid-friendly.*

1.2 Contributions

This thesis designs and evaluates systems to shape electricity demand in buildings and electricity distribution networks so as to make the demand more sustainable and grid-friendly, cut end user’s electricity bills, and reduce expenses in electricity distribution. To perform these tasks efficiently and in a user-friendly manner, we devise novel algorithms drawing from diverse fields of mathematical modeling, machine learning, and optimization. We evaluate the proposed algorithms and systems using simulations based on real power consumption data—either collected from real homes, or obtained from a local electric utility company—and by building research prototypes.

1.2.1 Summary of Contributions

This section presents a summary of the contributions presented in this thesis. The contributions can be classified in two broad classes: energy optimizations in buildings and energy optimizations across electric grids.

The key contributions and systems developed for energy optimizations in buildings are as follows:

- *SmartCap*: A system for transparently flattening a home’s electricity demand profile by scheduling background loads. Examples of such loads include A/Cs, refrigerators, dehumidifiers.
- *SmartCharge*: A system for cutting home’s electricity bill in presence of variable/dynamic pricing plans by leveraging energy storage. SmartCharge stores low-cost energy for use during high-cost periods.
- *GreenCharge*: An extension of *SmartCharge* for integrating on-site renewables in home’s consumption profile to further cut electricity bills.

Although energy optimizations in buildings can cut their electricity bills, they may not necessarily make the grid-wide consumption profile more sustainable. Hence, we investigate energy optimization techniques across the grid such that we can cut end user’s electricity bills, cut aggregate peak demands, and cut utility’s costs in electricity distribution. Key contributions and systems developed for energy optimizations across electric grids are as follows:

- *Scaling Distributed Energy Storage*: Two part solution for scaling energy storage adoption across grid. First, proposes to augment traditional variable electricity pricing plans with peak demand surcharge. Second, presents PeakCharge, an online algorithm for making battery charging-discharging decisions in presence of variable rates and peak demand surcharge.

- *Integrating Energy Storage in Electricity Distribution Networks*: Proposes to deploy hybrid combinations of different energy storage technologies across multiple levels of grid hierarchy to cut electricity distribution costs. Presents a framework for modelling major characteristics of various storage technologies, and for capturing the trade-offs of placing them at different levels of the distribution hierarchy to minimize a utility’s expenses in the distribution side.

1.2.2 SmartCap

We designed techniques to enable homes flatten their consumption/demand by scheduling background loads (such as A/Cs, refrigerators) that run transparent to home occupant’s knowledge. No interactive loads (such as microwave, lights) are affected by the scheduling. Demand flattening spreads demand to reduce difference between peaks and troughs in usage. Since demand peaks have disproportionate affect on grid’s operational and capital costs, flattening reduces overall electricity generation costs, which in turn makes electricity cheaper for end users. Further, it is an effective means of reducing carbon emissions associated with electricity generation. To identify the opportunities for demand flattening at homes, we analyzed power consumption data from real homes. The analysis showed that background electrical loads consume more than half of the total energy. Hence, we designed a load scheduler that flattens demand by scheduling only the background loads without causing user discomfort or requiring active user involvement. The scheduler, whenever possible, schedules the background loads such that, first, they do not all come up simultaneously thereby increasing peak, second, they avoid stacking on other “interactive” loads (such as microwave, etc). We evaluated the proposed system by simulations using consumption data from a real home, and deployment over smart home testbed comprising real home-appliances. We found that by scheduling background loads we were able to flatten the fluctuating demand profiles considerably. Our results demonstrate that

SmartCap can decrease the average deviation from mean power by over 20% across all periods where deviation is at least 400 watts, thereby flattening the “peaky” demand.

1.2.3 SmartCharge

Although SmartCap flattens a home’s consumption profile, most of the electricity pricing plans do not incentivize customers for flat demand profile. Hence, an obvious next question is: how can customers cut demand peaks and also save money in electricity bills? This is where SmartCharge comes in.

Besides scheduling background loads, another way of optimizing building’s energy footprint is by using energy storage. When employed in the presence of variable pricing plans such as time-of-use (ToU) pricing, storage can help cut the building’s electricity bill and demand peaks. Energy is stored in the battery during low-cost off-peak periods and the stored energy is used during high-cost peak periods to avoid expensive draw from the grid. Here we investigated the effectiveness of energy storage at lowering building’s electricity bill and peak demands without direct user involvement. We designed an intelligent battery charging system, called SmartCharge. SmartCharge needs two inputs, first, next day’s expected electricity prices, second, the home’s expected electricity consumption. Although electricity prices are usually known in advance, in order to accurately predict the home’s consumption we developed a machine learning (ML) based consumption prediction model. Empirical evaluation of SmartCharge showed that, with today’s battery technology and pricing, it can realize modest savings of 10-15% in electricity bills and can reduce the grid-wide peak demand by up to 20%. Further analysis showed that advancements in battery technology combined with expected rise in electricity prices will significantly increase the savings in near future.

1.2.4 GreenCharge

As one might expect, integrating on-site renewables with energy storage and variable pricing can further boost savings in end-user electricity bills and cut demand peaks. However, since renewables are intermittent and uncontrollable, buildings must still rely, in part, on the electric grid for power. While renewable deployments today use net metering to offset costs and balance local supply and demand, scaling net metering for intermittent renewables to a large fraction of buildings is challenging. In this work, we explore an alternative approach that combines market-based electricity pricing models with on-site renewables and modest energy storage (in the form of batteries) to incentivize end-user renewable deployments. We propose a system architecture and optimization algorithm, called GreenCharge, to efficiently manage the renewable energy and storage to reduce a building’s electricity bills. To determine when to charge and discharge the battery each day, the algorithm leverages prediction models for forecasting both future energy demand and future energy harvesting. We evaluate GreenCharge in simulation using a collection of real-world data sets, and compare with an oracle that has perfect knowledge of future energy demand/harvesting and a system that only leverages a battery to lower costs (without any renewables). We show that GreenCharge’s savings for a typical home today are near 20%, which are greater than the savings from using only net metering.

1.2.5 Scaling Distributed Energy Storage

Although optimizing building energy footprints reduces electricity bills, it does not necessarily make the aggregate grid-wide demand profile sustainable. Utilities need to shave the peak demands on their grids so as to make generation more environment-friendly, and optimize grid’s operational and capital costs. Therefore, utilities are transitioning to variable pricing plans that are engineered for incentivizing customers to shift consumption to off-peak hours and reduce peak load on the grid. Prior

research has proposed energy storage adoption by the end users to exploit these emerging plans. Even though homes can cut their bills using storage with variable pricing, using consumption data from hundreds of real homes we showed that energy storage adoption can worsen the aggregate peak on grid. Simultaneous battery charging across several homes during low price periods leads to the formation of “rebound peaks”. Tall rebound peaks form even when modest fraction of homes use energy storage. Thus, today’s variable pricing plans cannot sustain energy storage adoption at scale. We designed a two part solution to address this problem. First, we augmented variable pricing with a surcharge based on consumer’s peak demand. The surcharge encourages users to flatten their demand, instead of shifting most of it to low-price period, thereby preventing rebound peaks. Second, we proposed PeakCharge, an online peak-aware charging algorithm to optimize the use of energy storage in the presence of a peak demand surcharge. Extensive empirical evaluations on real-world traces showed that the proposed solution is effective at, first, maintaining incentives for consumers to use energy storage at large scale, second, reducing peak demands and ensuring grid stability. Further, our solution requires much less storage per consumer to maximize their savings, thereby significantly reducing the storage installation costs. Empirical evaluations show that total storage capacity required by PeakCharge to flatten grid demand is within 18% of the capacity required by a centralized system.

1.2.6 Integrating Energy Storage in Electricity Distribution Networks

So far, we have looked at employing energy storage at customer premises to cut electricity bills for the users, and reduce operational costs for the utilities. However, can utility companies directly employ energy storage in the grid and cut their costs without relying on the end users? Thus far, we have looked at deploying a single type of energy storage, i.e., lead-acid batteries, at a single level of the electric grid hierarchy,

i.e., homes. We now examine the efficacy of employing different combinations of storage technologies at different levels of the grid’s distribution hierarchy in cutting a utility’s operational costs. We present an optimization framework for modeling the primary characteristics that dictate the lifetime cost of many prominent energy storage technologies. The framework captures the important tradeoffs in placing different technologies at different levels of the distribution hierarchy with the goal of minimizing a utility’s operating costs. We evaluate the framework using real smart meter data from 5000 customers of a local electric utility. We show that by employing hybrid storage technologies at multiple levels of the distribution hierarchy, utilities can reduce their daily operating costs due to distributing electricity by up to 12%.

1.2.7 Thesis Outline

Chapter 2 includes background information needed to set context for the contributions in this thesis. Chapter 3 presents SmartCap, a system for flattening home electricity consumption by scheduling background loads. SmartCharge, a system for cutting home’s electricity bill in presence of variable/dynamic pricing plans by leveraging energy storage is presented in chapter 4. Chapter 5 talks about GreenCharge, which extends SmartCharge to integrate on-site renewables in home’s electricity consumption. Chapter 6 details our solution for scaling energy storage adoption across grid, including PeakCharge, an online algorithm for making battery charging-discharging decisions in presence of variable rates and peak demand surcharge. Chapter 7 presents the framework for modelling major characteristics of various storage technologies, and for capturing the trade-offs of placing them at different levels of the distribution hierarchy to minimize a utility’s expenses in the distribution side. Finally, we conclude the thesis with a summary of findings and future work in Chapter 8.

CHAPTER 2

BACKGROUND AND RELATED WORK

This chapter presents background information on green computing and energy storage to set the context for our contributions. More detailed related work sections are included in the relevant chapters.

2.1 Green Computing

Green Computing is concerned with designing systems with low energy consumption and low carbon footprints. There are two important aspects of green computing: first, *greening of computing*, i.e., making computing devices energy efficient; second, application of *computing for greening*, i.e., employing computer science methods to make physical systems green or energy-efficient.

Greening of computing, as previously mentioned, is concerned with making computing devices energy efficient. For example: making energy efficient mobile devices so that they consume less energy and have longer battery life; making energy efficient server clusters so that they consume less energy while executing jobs; building green or energy efficient data centers, for example by employing energy efficient blade servers. Several techniques are employed to achieve these energy efficiency goals such as frequency and voltage scaling, where the power consumption of a computing device is reduced by reducing the frequency and/or voltage, thereby also reducing the computation speed [112], [114]. Load consolidation is also applied in server clusters for energy efficiency; it is an approach to efficient usage of server resources by reducing the total number of servers required.

Computing for greening, on the other hand, is concerned with employing computer science techniques—such as sensor networks, optimization, artificial learning, and so forth—to make physical systems green or energy-efficient. For example, making manufacturing processes resource and energy efficient by network control and management of equipment. Similarly, employing computing to make buildings energy efficient is another example of computing for greening.

The second aspect of green computing, i.e., computing for greening, especially computing for achieving energy efficiency goals in *Smart Buildings* and electric grids is the focus of this thesis.

2.2 Smart Buildings

A building that can autonomously manage its energy footprint is a smart building [26]. These buildings, typically, have sensors deployed to track the energy consumption, building occupancy, and other building conditions such as temperature, humidity, etc. These buildings can also take automated actions based on the collected sensor data. For instance, if the occupancy sensors in a room detects that the room is unoccupied, the lights in the room can be turned off. This is a simple example of how smart buildings can manage their energy footprint. Besides, smart buildings are capable of incorporating user comfort preferences and integrating renewable energy sources in building’s consumption.

Modern buildings come in many types: office buildings, commercial buildings, industrial buildings, residential buildings. Most office and commercial buildings employ sensors for tracking energy consumption, occupancy, and building conditions and are managed by commercial Building Management Systems (BMS) [113] such as [52], [42], [100]. The BMSs provide some built-in automations, e.g., they are pre-programmed to turn on and off the lights and ventilation system of the building at scheduled times. Although these automations help in cutting a building’s energy con-

sumption, they lack intelligence. For instance, if the occupancy sensors detect the building is unoccupied, a BMS cannot turn down the cooling out of the scheduled time. In contrast, in this thesis we present autonomous, intelligent systems that can monitor, analyze, and control a building’s electricity consumption based on inference from sensor data, without active user involvement.

As opposed to the office buildings, residential buildings employed very little sensors, automation, and control until recently. However, with the advent of Internet-connected smart appliances, especially *Internet of Things* for home appliances, this is changing.

2.3 Internet of Things in Smart Homes

Internet of Things (IoT) is a proposed development of the Internet in which everyday objects—such as home appliances—have network connectivity, allowing them to send and receive data [116]. Equipped with Internet enabled sensors and coupled with intelligent cloud back end, these devices can cut user’s energy consumption. Several such sensors and appliances are available off the shelf. For example, eGauge energy meters([50]) connect to the building’s electric panel and measure its aggregate electricity consumption every second. These energy meters can upload the measured data to cloud, which stores, analyzes, and displays the data for the users, thereby helping them monitor and curtail their usage.

Nest’s smart learning thermostat [84] is an excellent example of Internet enabled smart home appliance. The Nest thermostat can learn home occupancy patterns and program itself. If the user forgets to turn off heat before leaving, the thermostat can take care of it. The thermostat can be programmed from anywhere over the Internet. The manufacturers claim, by doing all these smart things, it can cut heating and cooling costs by up to 15%. Besides, there are several other smart home appliances and sensors available in the market, such as [74], [66], [2].

Besides employing Internet enabled smart appliances, there has been work on programmatically regulating home’s electricity consumption profile. Authors in [29], [69], [106] recognize that home appliances with on-off controllers present a unique scheduling opportunity. For example, [106] presents an algorithm for scheduling a refrigerator with thermal slack off wind power. Authors in both [69] and [29] present offline optimization approaches to schedule multiple on-off loads, assuming that loads have well known and regular periods. Authors in [103] study the potential of optimizing load profiles by exploiting the elastic load components of common household appliances. For example by decreasing instantaneous power draw of an appliance at the expense of increasing its duration of operation.

Although Internet enabled smart appliances and proposed techniques for programmatic appliance control in the literature are beginning to employ intelligence to achieve energy efficiency goals and cut electricity bills, most of them have limited scope and work in isolation to each other. In this thesis, we devise techniques for programmatically controlling home appliances collectively, to regulate home’s consumption profile. Also, as opposed to some of the prior work, we do not assume that the appliance duty cycles are predictable and hence we devise online energy optimization techniques. In general, we employ several computing techniques—including machine learning, optimizations, online algorithms—to enable global energy optimizations across several appliances in a building, and across thousands of homes in a grid, while ensuring the aggregate grid-wide consumption profile becomes more sustainable.

2.4 Variable Electricity Pricing Plans

Cutting electricity consumption of buildings—for example by using smart appliances like Nest’s thermostat—makes them energy efficient and cuts their electricity bills. However, this is not the only way of achieving energy efficiency in buildings

and the grid. Another equally effective way of making buildings energy efficient is by reducing their consumption during peak demand periods. During these periods, a lot of electricity is being used everywhere and electric grid is under stress. Electric utilities have to dispatch extra peaking generators to satisfy this demand, these generators typically operate on expensive fossil fuels and have a greater carbon footprint. Therefore, cutting a building’s consumption during peak periods not only makes the building energy efficient, but also cuts the overall carbon footprint of electricity generation. Furthermore, peak demands dictate the grid’s capital expenses—as the generation and distribution infrastructure has to be provisioned for the peaks—hence reducing peaks also cuts capital costs.

Therefore, to incentivize peak reduction, many utilities are transitioning from conventional fixed-rate pricing models, which charge a flat fee per kilowatt-hour (kWh), to new market-based schemes, e.g., real-time or time-of-use pricing. These market-based schemes have a higher price for electricity during high demand peak periods and a lower price during low demand off-peak periods. For instance, Illinois already requires utilities to provide residential customers the option of using hourly electricity prices based directly on wholesale prices [101], while Ontario charges residential customers based on a time-of-use scheme with three different price tiers (off-, mid-, and on-peak) each day [87].

Although dynamic electricity pricing plans incentivize customers to shift their consumption to low-price off-peak periods, naive ways of doing this can lead to user inconvenience and increase demand peaks on the grid, as shown in Chapters 4, 6. Therefore, in this thesis we present solutions to automatically (without active user involvement) shift a building’s consumption to off-peak periods and cut their electricity bills while ensuring that the grid-wide demand peak becomes more sustainable.

2.5 Energy Storage in Buildings and Electric Grids

Energy storage in the grid is used to store excess electricity when production (from power plants especially intermittent renewable electricity sources such as wind power, tidal power, solar power) exceeds consumption [115]. The stored energy is used in times when the electricity consumption exceeds production and cannot be deferred/delayed. Therefore, energy storage helps in smoothing out the electricity production by avoiding drastic scaling up and down due to momentary fluctuations in consumption.

Another important function of storage in the grid is integration of renewable energy. As renewable energy from solar, wind, and tidal sources varies inherently with time and weather, it is seldom reliably available during peak demand periods. However, excess renewable energy can be stored in batteries and used during peak demand periods.

Several types of energy storage technologies have been deployed in the grid, for example: compressed air storage, flywheels, pumped water storage, super conducting magnetic energy storage, etc. Recently, several battery startup companies have also been coming up with new storage technologies for the grid such as [82] [53] [90] [27].

More recently, researchers have proposed using energy storage in presence of real-time electricity prices to cut costs—for example [43], [46]. The basic idea in most of these works is to store cheap electricity during low-price off-peak periods and then later use the stored energy during peak periods, thereby cutting costs. There has also been work on economics of scaling energy storage across a number of homes with today's variable pricing plans [37] and [110]. These studies have shown, although energy storage under variable pricing can cut costs, it may increase the aggregate grid-wide peak, instead of reducing it.

In this thesis, we too investigate employing energy storage to cut customer's bills and optimize energy footprints. However, in contrast to the existing literature, we

employ energy storage not only to maximize cost savings, but also ensure that the grid-wide consumption profile becomes more sustainable. Further, we evaluate our techniques using real-world power consumption data from homes and electric utility companies. Besides, as opposed to the existing work, we investigate deploying hybrid combinations of various storage technologies across the grid hierarchy to cut utility's daily expenses in electricity distribution.

CHAPTER 3

SMARTCAP: FLATTENING PEAK ELECTRICITY DEMAND IN SMART HOMES

Modern buildings are heavy power consumers, hence making buildings efficient and shaping their electricity demands can help in making the electric grid more sustainable. In this chapter we design techniques to enable homes flatten their electricity consumption by scheduling background loads (like A/Cs, refrigerator) without causing user discomfort or requiring active user involvement.

3.1 Introduction and Motivation

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [7]. As a result, designing new “green” buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home’s energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect *loads*¹ in empty rooms [64]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [63]. These intelligent load management schemes reduce a home’s energy footprint and its bill by automatically disconnecting loads from power when necessary

¹We use the term load throughout the chapter to refer to any appliance or device in the home that draws electricity.

or convenient. This chapter focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home’s electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [7], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [8, 32, 41, 99]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer’s peak usage. For example, a utility may determine the bill, in part, based on a customer’s largest half-hour of electricity demand within a day, regardless of the total day’s energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challeng-

ing [45]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap’s design, this chapter studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility (lights, TVs, microwaves, etc.), a substantial portion of home electricity demand derives from loads with some limited flexibility. These flexible loads, such as air conditioners (A/Cs), refrigerators, freezers, dehumidifiers, and heaters, typically operate in the background: while the result of their power draw is readily apparent, e.g., a comfortable room temperature and frozen food, *when* they draw power and the *magnitude* of this power draw is not important. Note that flattening demand is distinct from, and orthogonal to, conservation efforts that reduce total energy consumption over long periods. Instead of reducing total energy usage, flattening demand redistributes consumption by shifting load to decrease demand peaks while filling in troughs. A goal of our work is to quantify when and how much demand flattening is possible from background loads.

We hypothesize that homes are capable of flattening electricity demand during peak load times by intelligently scheduling only background loads. To evaluate our hypothesis, we analyze power data gathered from a real home at outlets, switches, and panels over three months. Our data shows that while background loads account

for under 10% of the loads on a typical summer day, they consume nearly 60% of the energy. SmartCap’s load scheduler flattens demand by scheduling background loads according to a Least Slack First (LSF) policy, inspired by the Earliest Deadline First algorithm in computing systems, where slack is a measure of how long each background load is able to remain off without affecting its objective, e.g., maintaining an environmental setpoint or fully charging a battery. We evaluate SmartCap by simulating background load scheduling using data from our home deployment. We also implement SmartCap in a smart home testbed we have built, which uses similar background loads as our home. We leverage our testbed to experiment with SmartCap using real appliances. As an example of our results, we show that LSF decreases the average absolute deviation from the mean power (a measure of flatness) by over 20% for all 4-hour periods (over the 82 day period) where the deviation is greater than 400W.

3.2 Related work

Increasing the penetration of demand-side load management in residential settings is a key goal of smart grids. Thus, SmartCap’s general architecture, which includes the home gateway, an array of real-time power meters, and programmable switches, is similar to other proposed architectures for programmatically regulating home electricity demand [29, 30, 69, 95]. While space constraints preclude a full survey of prior work, past approaches focus on using these architectures for a range of scheduling objectives, such as reducing total consumption, reducing costs based on variable prices, or varying usage to match renewable generation or make use of a battery. Our work differs in its focus on flattening demand without affecting occupants by scheduling background loads. We do not explore scheduling to satisfy other objectives, since it requires disturbing occupants by periodically disconnecting interactive loads.

Prior work also recognizes that loads with on-off controllers present a unique scheduling opportunity [29, 69, 106]. For instance, Taneja et al. [106] present an algorithm for scheduling a single refrigerator with slack that operates off wind power. Both Keshav and Rosenberg [69] and Bakker et al. [29] present offline optimization approaches to scheduling multiple on-off loads in isolation, assuming that loads with on-off controllers have well-known and regular periods. In contrast, our work quantifies the benefits of scheduling background loads in a real home. Data from our home reveals that background loads do not exhibit regular periods, due to environmental changes, while interactive loads are difficult to predict in advance. As a result, we eschew offline optimization scheduling algorithms in favor of an online approach that uses each load’s current slack as a heuristic to determine its priority at any time.

3.3 Background and Problem Formulation

The focal point of SmartCap’s architecture is an intelligent smart home gateway. The home gateway serves as the interface between a smart home and the smart grid. As shown in Figure 3.1, the gateway receives information from multiple potential sources, including real-time electricity prices and demand-response signals from the grid, generation data from on-site renewables, and consumption data from each household load. The gateway’s data sources inform its load scheduling policy. This policy determines *which* loads to power and *when* by issuing actuation commands to loads. While we focus on the problem of scheduling background loads to flatten demand without affecting occupants, our home gateway is capable of implementing scheduling policies with other objectives, such as ensuring home power demands are always less than supply when using intermittent on-site renewables [118]. SmartCap’s architecture depends on loads that expose programmatic control to turn them on and off. While today’s “dumb” appliances generally do not expose such remote actuation capabilities, utilities are currently testing such smart appliances for demand response




Figure 3.1. A graphical depiction of a SmartCap-enabled home.

initiatives [63]. As appliances begin to allow remote actuation at finer granularities, advanced techniques for controlling power will be possible. Given current standards for remote actuation, connecting loads to external programmable switches and outlets using home automation protocols, such as X10 or Insteon, is sufficient in many cases to provide programmatic load actuation in today’s appliances. We currently use Insteon-enabled outlets and switches in our home deployment and testbed [118].

We divide electrical loads into two broad groups: interactive and background loads. Household occupants directly control interactive loads by toggling switches, and actively observe their behavior; examples include lights, TVs, computers, microwaves, and vacuums. The vast majority of household loads are interactive. Our LSF scheduling policy assumes that only the occupants are able to control interactive loads. In contrast, household occupants do not directly control background loads, and only passively observe their behavior; examples include refrigerators, dehumidifiers, and A/Cs. As long as these loads satisfy occupant expectations, e.g., a target temperature or humidity level, their usage pattern is neither important nor noticeable. SmartCap monitors background loads and controls when they consume power. We view transparently flattening demand from background loads as an important prerequisite in satisfying many other demand-side scheduling objectives. While dis-

connecting interactive loads may be necessary at certain times to strictly cap power usage, scheduling background loads without affecting occupants should always be the first priority under constraint. Note that SmartCap’s architecture explicitly does not permit utilities to monitor or control household loads, since such capabilities represent an invasion of privacy [81].

3.4 Load Analysis and Observations

To study the extent to which scheduling background loads is able to flatten demand, we collect fine-grained power data from a real home that houses three occupants. We have collected the home’s aggregate power for the last 12 months and power at each outlet and switch for the past 82 days. Since our monitoring did not affect the occupants’ daily routine, our data reveals realistic home power usage patterns over the monitoring period. Our home deployment continuously gathers power usage data for the entire home every second and 30 individual outlet loads every few minutes; our prototype maintains a record of the on-off state of 30 of the home’s wall switches at every instant in time. SmartCap’s gateway is also able to remotely (and programmatically) control the home’s outlets and wall switches. More details about our home deployment are available in prior work [118].

3.4.1 Interactive vs. Background Loads

To quantify the potential benefits of scheduling background loads, we separate the power consumption of background loads from that of interactive loads. In our prototype home, we monitor seven background loads at outlets: a refrigerator, a freezer, a dehumidifier, three window air conditioning units (A/Cs), and a heat recovery ventilation (HRV) system. By contrast, we estimate that the home used 85 distinct interactive loads over the past year. Thus, SmartCap does not attempt to schedule the vast majority of household loads, since it would affect the home’s occupants.

<i>Load</i>	<i>Peak</i>	<i>Average</i>	<i>Quantity</i>
Refrigerator	456W	74W	1
Freezer	437W	82W	1
HRV	1129W	24W	1
Dehumidifier	505W	371W	1
Main A/C	1046W	305W	1
Bedroom A/C 1	571W	280W	1
Bedroom A/C 2	571W	141W	1
Background	4715W	1277W	7
Interactive	9963W	887W	85

Table 3.1. In the summer, background loads in our home account for 59% of the total energy consumption.

Interactive loads that we do not schedule include lights, entertainment appliances (e.g., TV, cable box, gaming console), computing equipment (e.g., routers, laptops, desktops), kitchen appliances (e.g., microwave, toaster oven, espresso maker, garbage disposal), and miscellaneous devices (e.g., clocks, vacuums, hair dryers). In most cases, disconnecting any of these loads from power when in use is readily apparent to occupants. We also group clothes dryers, washing machines, and dishwashers with interactive loads. While we could schedule the start time of these appliances, we do not include them because adjusting the start time affects occupants. To see why, consider that a scheduler may be able to decide when an appliance executes, but occupants must ultimately initialize the appliance, e.g., fill it with clothes or dishes, before its scheduled start time. Changing the start time may force occupants to initialize the appliance at an inopportune time. Further, for clothes dryers and washing machines, their operation is often pipelined, with households washing multiple laundry loads back-to-back.

Observation #1: While background loads comprise 7.5% of the total loads over our monitoring period, they account for 59% of the average energy consumption. Table 3.1 shows the peak and average power consumption for each background load we monitor during a representative week in the summer, as well as the peak and average power




Figure 3.2. The power consumption of interactive loads is highly variable throughout the day. As expected, peak power consumption occurs around mealtimes in the morning, early afternoon, and early evening.

consumption for all background and interactive loads. During this week, background loads consume 209 kWh, while interactive loads consume 146 kWh. The three window A/C units significantly increase the fraction of energy consumed by background loads, since each A/C draws between 400W and 1kW when the compressor is on. On hot days, the compressor may run as much as half the day, depending on the comfort level the occupants desire. Note that during the winter the A/Cs do not run, since the home uses a gas furnace for heat. As a result, background load is lower in the winter. In this case, the duct heater for the HRV system, which heats incoming air from the outside, dominates background energy consumption, accounting for 70% of the total, while the refrigerator, freezer, and dehumidifier account for the remaining 30%. Below, we highlight other observations from our home’s data that influences our approach to scheduling.

3.4.2 Interactive Variability

Observation #2: The power consumption of interactive loads varies due to the actions of occupants throughout the day, and is not readily predictable. Figure 3.2 highlights this point by showing the power consumption of the interactive loads in isolation on a typical day. Additionally, Figure 3.3 shows consumption patterns for four interactive loads. Notice that the power draws of these loads vary considerably throughout the day, with the peak periods occurring during the morning between 6am and 10am and in the early evening between 5pm and 9pm. These periods coincide with food preparation and are partially the result of using high-power kitchen appliances, such as a coffee pot, garbage disposal, microwave, dishwasher, or toaster oven. During the night, the minimum steady state power consumption is roughly 200W, while during the morning and evening it frequently rises above 2kW for frequent short periods.

The kitchen appliances tend to induce peaks by using large amounts of power for relatively short time periods, such as the coffee pot in Figure 3.3. Our observation also holds for meal preparation at breakfast, lunch, and dinner. Accurately predicting the power consumption of interactive loads at fine time scales is difficult. While the home’s occupants typically eat dinner between 4pm and 8pm, if and when they use a microwave, toaster oven, dishwasher, or garbage disposal is highly variable during this four hour time window each day. Additionally, the occupants have flexible work schedules, and often work from home during the day—on this day one of the occupants ate lunch at home, which accounts for the spike in power around noon. Since interactive loads are not readily predictable, our scheduler must be able to react to drastic and sudden changes in their power consumption.

3.4.3 Background Variability

Observation #3: The operating period of background loads varies due to both environmental conditions and external events, and is also not readily predictable. Figure 3.4 highlights the point by graphing the power consumption of four of the background loads we monitor. Each background load is clearly periodic: it alternates between distinct ‘on’ and ‘off’ states. While it is possible to design these loads with variable drive controllers, all the background loads in our home use simple on-off controllers that toggle between an on and off state [102]. In this case, the on-off periodicity is a result of each background load maintaining an environmental *setpoint*: in this example, the refrigerator and freezer maintain their internal temperature within a fixed guardband, the dehumidifier maintains a humidity level within a fixed guardband, and the HRV heats outside air to a pre-specified temperature. The *guardband* defines the acceptable maximum and minimum levels for the load’s target environmental metric. Common household loads use simple control loops to stay within the guardband. For example, when the load’s metric reaches a maximum allowable value, the load turns on until the metric reaches a minimum value, at which point the load turns off.

Since environmental conditions vary, neither the length nor the magnitude of a load’s on-off period is entirely regular. To illustrate, the figure shows that the refrigerator (upper-right) and freezer (upper-left) exhibit longer on periods in the early evening between 5pm and 9pm, along with some transient usage spikes. In both cases, the longer on periods are the result of the occupants opening the refrigerator and freezer doors, which increases the internal temperature and causes them to turn on their compressors to lower the temperature. Tasks other than maintaining temperature also contribute to the transient spikes in power consumption. For example, both the refrigerator and freezer power multiple 60W incandescent light bulbs when the door is open and also periodically make ice; the refrigerator also cools a separate




Figure 3.3. Power data for example interactive loads. Occupant behavior, which is not readily predictable, determines when these loads draw power.

freezer compartment. The refrigerator exhibits a much more irregular consumption pattern, since it resides in the kitchen and the occupants open its door more frequently than the basement freezer. The HRV and dehumidifier exhibit irregular periods for similar reasons.

The dehumidifier’s operating cycle dictates that it runs until it reaches a setpoint humidity—in our case 50%—or until it has run for two consecutive hours, at which point it remains off for 2 hours to cool down. Thus, on hot and humid summer days, the dehumidifier will run for 2 hours every 4 hours if it cannot reach its setpoint humidity, and consume a significant fraction of power (1.8 kWh). On moderately humid days, the dehumidifier will come on and off according to its setpoint humidity, causing an irregular on-off period. On this day, the environmental humidity was high, so the dehumidifier ran regularly. Similar to the refrigerator/freezer, the window




Figure 3.4. Power signatures for four background loads in our home. The on-off period varies with environmental conditions, and is not regular.

unit A/Cs exhibit irregular periods based on changing outdoor temperatures and the frequency with which exterior doors open and close. While some environmental factors may be partially predictable, such as temperature or humidity, interactive events such as doors opening and closing also affect the period and power consumption of background loads. Thus, scheduling background loads must take into account these difficult to predict changes in their periodicity.

3.5 Load Scheduler

SmartCap’s background load scheduler leverages the well-known concept of *slack*, which quantifies the extent to which a scheduler is able to advance, defer, raise, or lower a load’s power consumption without affecting its operational goal [29, 69, 70,

106]. Before detailing the LSF algorithm, we first discuss different types of load controllers to understand the available dimensions of scheduling freedom.

3.5.1 Load Controllers

Simple on-off controllers encompass the vast majority of controllers found in residential loads, since they are cheap and reliable. As discussed earlier, on-off controllers often maintain an environmental metric, e.g., temperature or humidity, within a specified guardband. For these loads, slack arises from the fact that the load is able to remain off until its metric reaches the guardband’s maximum (or minimum) value, at which point the load must turn on. In effect, these loads indirectly store power in their contained environment by increasing (or decreasing) a target metric, which then slowly decreases (or increases) due to leakage with the outside environment. On-off controllers are also commonly driven by timers, which dictate fixed-length on-off periods. While a scheduler is able to advance or defer when these loads turn on or off, as long as they do not violate their guardband or fixed-length on-off period, it is not able to raise or lower power consumption when the loads are on.

Battery chargers are another example of a load with slack, since they are capable of raising or lowering their power consumption by adjusting the charging rate. While most household batteries are small, e.g., phones, laptops, and tablets, the emergence of plug-in electric vehicles (EVs) is poised to introduce a large load with substantial slack to homes. EVs that plug into standard 120V/15A outlets are able to charge at a rate of up to 1.8kW, while a dual-pole 240V/30A circuit that uses both legs of a home’s split-phase input power is able to charge at a rate of up to 7.2kW. In either case, advanced chargers are capable of varying the rate of charge up to these maximums. For battery chargers, the primary scheduling constraint is fully charging the battery over some duration, or charging to an acceptable capacity,

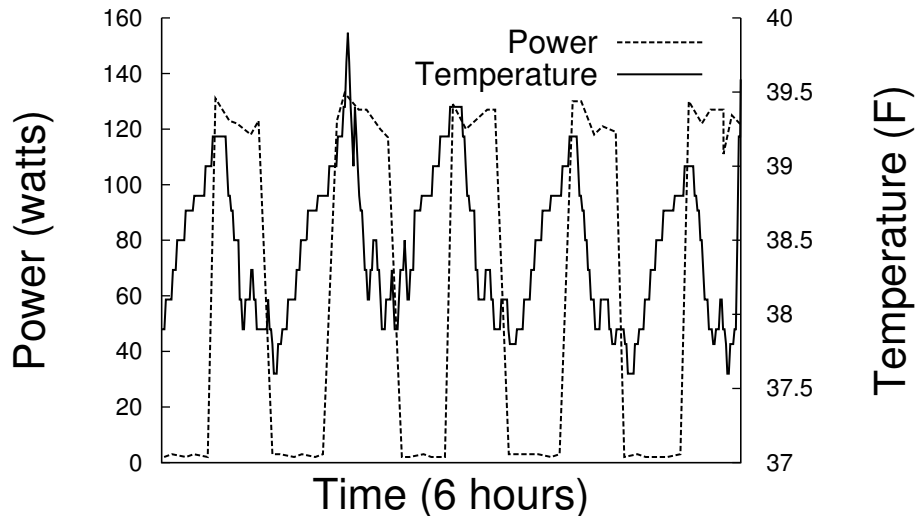


Figure 3.5. A depiction of slack in our refrigerator’s simple on-off control loop. The compressor turns on once the internal temperature reaches an upper threshold, and turns off once it reaches a lower threshold.

While not present in our prototype, variable drive controllers are capable of raising and lowering their power consumption when on. These controllers offer clear benefits over on-off controllers, but they are typically not found in household appliances due to cost and reliability issues. As a result, our experiments do not study their impact.

3.5.2 Scheduler

We define a load’s slack at any time t as the remaining length of time the load can be off, i.e., disconnected from power, without assuring that it will violate its objective. For a load that maintains an environmental condition with an on-off controller, it must turn on when its environmental metric reaches a guardband boundary. For a battery charger, it must turn on when only the maximum charging rate over the remaining plug-in duration is sufficient to fully charge the battery, or to charge it to an acceptable capacity. We define slack in units of time, rather than energy as in [106], only for ease of exposition—slack time is proportional to slack energy for stable load and environmental conditions. We assume each load is able to maintain an estimate

of its remaining slack time based on its current power state and by monitoring the state of its internal and external environment. As shown in Figure 3.5, slack estimates may change over time based on both the load’s power state—when the load is off slack increases—and environmental conditions, such as a refrigerator door opening or the humidity increasing. Since these changes in slack may be unpredictable, our scheduler is reactive and online, continually adjusting which loads receive power based on their available slack. Finally, we assume that our gateway is able to query the slack of each load at any time using simple models as in [106].

Before describing our scheduler, we first illustrate a simple example using ideal background loads with well-defined on and off periods in isolation, and without uncontrollable interactive loads. The illustration demonstrates how shifting power usage is able to flatten demand. Figure 3.6(a) depicts an extreme example, where the slack for three window A/C units that draw 1kW when on dictates that they must turn on for 15 minutes anytime within each hour to maintain their respective setpoint temperatures. In the worst case, without any scheduling, these units may be nearly synchronized and cause power usage to reach 3kW for close to 15 minutes over the hour, while drawing 0W for the remaining 45 minutes. In the best case, with appropriate scheduling, it is possible to shift the on periods such that only a single A/C is on at any given time, resulting in a peak usage of only 1kW (Figure 3.6(b)); since the on periods of the A/Cs interleave with room to spare, we are able to perfectly flatten demand.

To quantify flattening over an interval, we use the *average absolute deviation from the mean power*, which is an average of the absolute difference between power at every time t and the average power. We use this metric instead of the standard deviation simply because it is more intuitive; standard deviation exhibits the same trends but is greater than or equal to our metric. The magnitude of the deviation quantifies how much demand varies; a lower deviation indicates flatter demand and a better

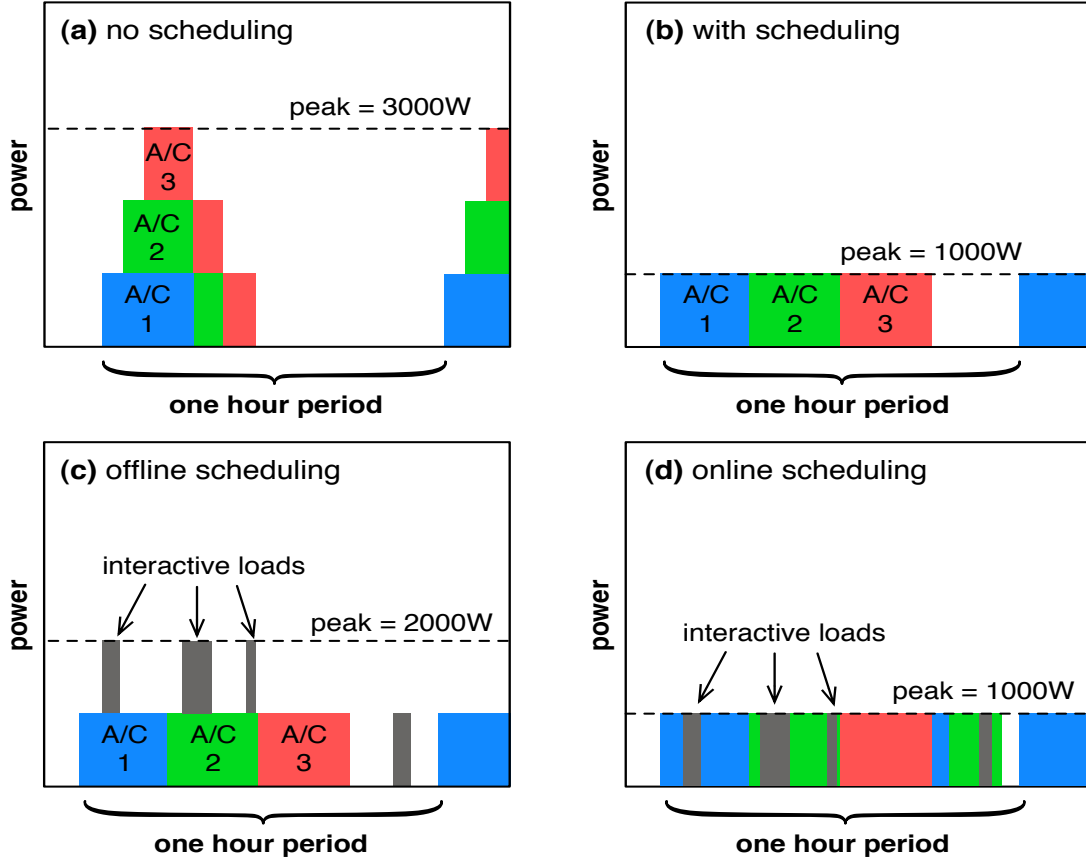


Figure 3.6. A background load scheduler is capable of flattening demand, but must account unpredictable interactive and background loads.

schedule. In our example, the worst-case no scheduling scenario has a deviation of 1125W from the mean power, while the best-case scenario has a deviation of 375W due to 15 minutes of no power consumption at the end of the period. In this scenario, interleaving the A/Cs results in a 3x reduction in the deviation and, thus, a significantly flatter demand profile.

As noted in prior work [29, 69], the scheduling problem for ideal background loads with regular known on-off periods distills to a simple offline optimization problem in the absence of interactive loads. Figure 3.6(c) demonstrates how interactive loads alter scheduling by inserting into our previous example four 5 to 15 minute peaks of 1000W during the hour-long period, as could be expected from heating up food in

a microwave. Even though A/Cs have enough slack within the hour to defer their power consumption whenever the microwave turns on (Figure 3.6(d)), an algorithm that determines the schedule in advance will not know about these microwave events. While this is a simple idealized example, it illustrates that load scheduling in the presence of unpredictable interactive loads is an online, and hence heuristic, process. Sudden and unpredictable changes to a load’s slack, such as from opening doors or changes in weather, introduce similar issues that warrant an online approach. As we discuss in Section 3.7, and in contrast to Figure 3.6(a) and (b), we find that scheduling background loads is most advantageous during “peaky” periods with many short, but high power, interactive loads.

SmartCap’s scheduler executes every interval T to determine which background loads receive power (and how much for the battery charger). In our simulator and testbed, we choose T ’s length to be significantly less (one minute) than the typical on-off periods of our background loads; the setting also ensures that background loads are not quickly turned on and off, which may degrade their reliability. We assume that once a load’s slack reaches zero, the scheduler must provide it the necessary power regardless of the increase in peak usage. We call our basic load scheduling policy *Least Slack First* (LSF), since it supplies power to loads in ascending order of their current slack value. Thus, loads with a lower slack have a higher priority. LSF is a direct adaptation of the Earliest Deadline First (EDF) scheduling policy common in real-time operating systems. We combine LSF with a target capacity threshold to determine how many loads to power, and how much power to supply to battery chargers. Once the sum of the background loads’ power usage reaches the capacity threshold, the scheduler stops powering additional background loads. Figure 3.7 depicts how LSF scheduling flattens demand for a real power signal, assuming three A/Cs turn on near each other as in Figure 3.6. As in our example, LSF flattens the demand profile by interleaving the on periods.




Figure 3.7. Example of how LSF flattens demand.

Our experiments use an adaptive threshold based on an exponentially weighted moving average of the home’s power consumption over the previous hour. Setting the capacity threshold presents a trade-off. A threshold too low causes the scheduler to defer too many loads, resulting in their slack values approaching zero in tandem. This induces large peaks by ultimately forcing the scheduler into simultaneously powering many loads with zero slack. A threshold too high causes the scheduler to power too many background loads at a time, resulting in a peak that is higher than necessary.

3.6 Prototype: Design and Implementation

Our SmartCap home deployment is in an average 3-bedroom, 2-bath house with 1700 sq. ft. and a total of 8 rooms across three floors, including a basement. Since the prototype is a real home with three occupants that went about their daily routines during the monitoring period, our data reflects real-life home usage patterns. The home does not have central air, and its furnace and water heater use natural gas, which removes three potentially large consumers of electricity. During the summer,

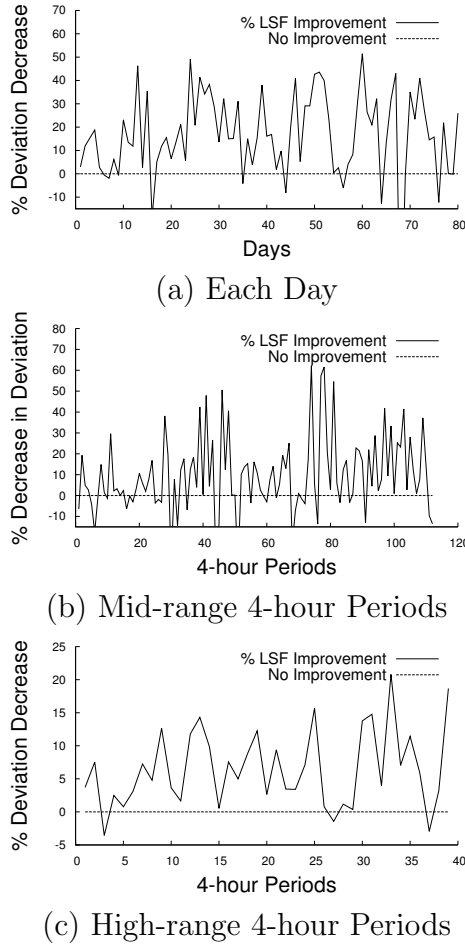


Figure 3.8. LSF decreases the absolute average deviation from the mean power (with no scheduling) on the vast majority of days (91%), as well as over peak 4-hour periods with mid-range and high-range deviations.

the occupants use three window A/C units to cool the home—one large unit in the living room and a smaller unit in each upstairs bedroom.

We provide a brief summary of our SmartCap deployment. A TED 5000 measures power consumption for the entire home every second using meter-like measurements of the wires supplying grid power to the home’s main circuit breaker panel. The TED specification claims accuracy within 2%; we found the TED to be within 1% of the utility power readings during the monitoring period. We use Insteon-enabled switches to monitor and control loads; Insteon is a common, commercially-available home automation protocol that uses power line communication. In particular, we

use the Insteon iMeter Solo to monitor power at background load outlets, and the Insteon ApplianceLine to control power to our background loads from our gateway. Our gateway connects to an Insteon Power Line Modem (PLM), which is able to inject Insteon commands and listen for responses over the home’s power lines. The gateway both polls the iMeters for their power usage and issues on-off commands to the appliances through the PLM. For background load scheduling, SmartCap only requires power data for the whole home and at the seven background loads. However, our prototype is capable of remotely monitoring and controlling each outlet and wall switch in the home [118]. To monitor environmental metrics and compute slack, we deploy eight temperature and humidity sensors inside or near each background load, as well as outside, using an Oregon Scientific WMR200A weather station.

In addition to our in-home SmartCap deployment, we also setup a smart home testbed to mimic our home’s background loads. The testbed enables us to perform repeatable experiments, such that we do not disturb home occupants. It uses the same SmartCap system as our real deployment: Insteon-enabled power meters and switches to monitor and control background loads. The background loads include a humidifier, dehumidifier, multiple electric heaters, a freezer, and a refrigerator – we use heaters rather than A/Cs, since our testbed resides within a window-less room and A/Cs require outside drainage. Since we use external load control switches that are not integrated with the appliance to connect and disconnect power, we use appliances that remember and restart in the same setting after a power outage. For experiments, we are able to replay traces using our home data both with and without LSF scheduling.

3.7 Evaluation

We evaluate LSF in simulation and in our smart home testbed to explore its performance in realistic settings. Our simulator, written in Java, uses input traces of

household load events to simulate background load scheduling using LSF. Each load event corresponds to a change in the power level for a single load. The simulator also associates both a maximum and minimum slack value with each background load every period, which includes a single off interval and its subsequent on interval. At each period boundary, the simulator assumes the load is at its maximum slack value if it has just transitioned to the off state, and assumes the load has zero slack if it has just transitioned to the on state.

The simulator uses a simple linear model for computing per-period slack: when a load is on its slack increases linearly, and when it is off it decreases linearly. To always ensure that the load reaches its maximum slack by the end of each period, the simulator determines the slope of the linear increase or decrease in slack using the ratio of the on and off durations for the current period. Note that, due to environmental changes, each background load may exhibit different period durations, as well as per-period on and off durations, throughout the day. In practice, SmartCap may use an environmental model to compute slack in real time; linear models tend to perform well, as Figure 3.5 demonstrates for the refrigerator and its inside temperature. Since we automatically generate input traces from the home’s power data, our per-period slack computation is an indirect way of accounting for environmental changes in simulation.

Since our scheduler only controls background loads, our input traces represent all interactive loads as a single load with many frequent load events. To get power readings for the interactive loads, we subtract each background load from the home’s aggregate power consumption. Note that since we collect aggregate power consumption every second, our trace includes a new event nearly every second to represent the changing consumption of the interactive loads.

3.7.1 Simulation Results

We first evaluate LSF for flattening peak power usage in our home deployment. We focus on the last 82 days during the summer. Flattening is most important during summer months, since peak demands typically occur in these months [6]. Figure 3.8 shows the percentage decrease in average absolute deviation from the mean power using LSF scheduling for different periods. Recall from 3.5 that we use the average absolute deviation to quantify the flatness of the demand profile. Figure 3.8(a) plots the percentage decrease in deviation over each day, and demonstrates that LSF flattens the profile on over 91% of days, resulting in a 16% flatter profile on average. LSF does not flatten the profile on 9% of days, since those days already have a low deviation without scheduling. On 33% of days, LSF decreases the deviation by more than 20%. These results are significant, since each day includes long periods of relatively little activity, e.g., all night, where the average deviation is not high due to minimal occupant activity. Despite the long periods of inactivity that occur each day, LSF is still able to flatten the day-long demand profile.

We also examine how LSF performs for shorter 4-hour intervals that correspond to peak usage times, since these are the periods where demand flattening is most important. We divide 4-hour periods throughout the summer by the magnitude of their average absolute deviation (or “peakiness”). We find that LSF does not provide much improvement (<3%) for periods that do not exhibit “peaky” behavior, since the demand profile is already flat. We find that over 69% of the 4-hour periods throughout the 82 days have average deviations less than 400W; these periods generally correspond to nighttime or when the home is unoccupied. The remaining 31% of the periods exhibit deviations from 400W to 1000W (22%) and over 1000W (9%). Figure 3.8(b) shows that LSF works well for the mid-range (400W-1000W) and high-range (>1000W) 4-hour periods, decreasing the respective average deviations by 23% and 21%, on average.

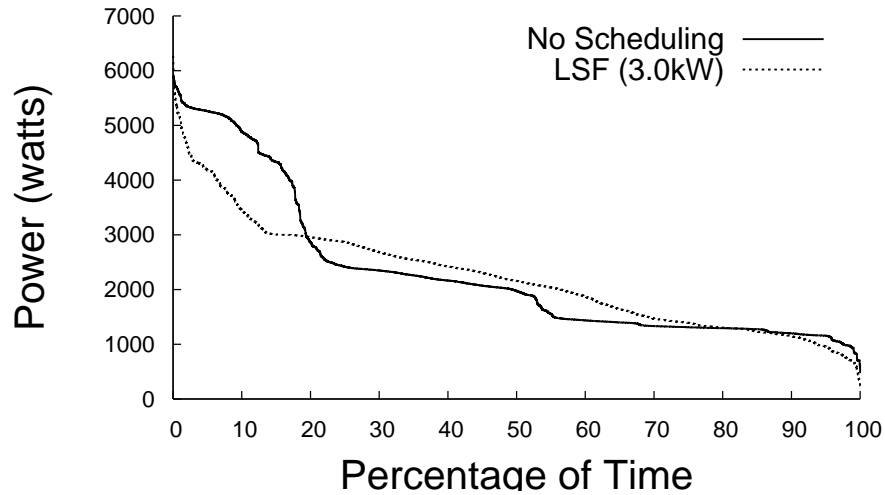


Figure 3.9. Load duration curves for a typical summer day with and without scheduling when using an electric vehicle.

The data indicates that many flat 4-hour periods exist throughout our trace, which suggests that most of LSF’s improvement stems from scheduling background loads around interactive loads that cause brief, but significant, power peaks. If the background loads themselves interleaved to cause significant peaks in power, we would expect more improvement during periods with few interactive loads, e.g., nighttime. Since our home has many background loads that operate based on different environmental conditions, they rarely all turn on simultaneously. Thus, without LSF, the background loads already exhibit a great deal of statistical multiplexing, and there is little LSF can do to flatten their peak usage. Our results also indicate that LSF works well during “peaky” periods, which typically occur during peak demand periods, where the average deviation is high.

3.7.2 Impact of Electric Vehicles

We also studied the impact of EVs on LSF’s ability to flatten peaks. Today’s grid was not provisioned for the increased power consumption from widespread EV adoption. As a result, the grid must either add capacity or use better load scheduling,

e.g., through new pricing models, to force EVs to multiplex their charging over time. For instance, in our home, charging an EV on a typical summer day increases the home’s total power consumption by 52%. SmartCap and LSF represent a possible avenue for flattening demand with EVs. As in the simulations above, we use data from our prototype home, but add an EV charger based on the Chevy Volt, with a battery capacity of 16kWh plugged in at night between 7pm and 6am that takes 5 hours to charge.

Figure 3.9 shows the results for an average summer day by plotting a load duration curve both without scheduling and using our LSF scheduler. Load duration curves are a common method for visualizing the flatness of power distributions. The curve shows the percentage of time on the x -axis during the day that electricity demand was at the corresponding power value on the y -axis. An ideal load duration graph is a completely horizontal line at the average power usage. On this day, LSF reduces the average absolute deviation by 22%. In particular, LSF reduces the peak time periods where demand is highest (on the left side of the graph) significantly, and shifts their power consumption across many of the lower power periods throughout the day.

3.7.3 Testbed Results

Finally, to demonstrate LSF’s performance in a realistic setting we use our smart home testbed. Figure 3.10 shows the power usage, as measured by our Insteon power meters, for a representative 4-hour period. We use data from our home on June 15th from 2pm to 6pm to replicate the same sequence of background load on and off periods in our testbed. As discussed earlier, our gateway sends commands to Insteon ApplianceLincs to connect and disconnect background loads from power. The experiment demonstrates how LSF shifts the power usage of the background loads forward to compensate for the interactive loads early in the period. As a result, on this day, LSF decreases the absolute average deviation from the mean power by 23%.

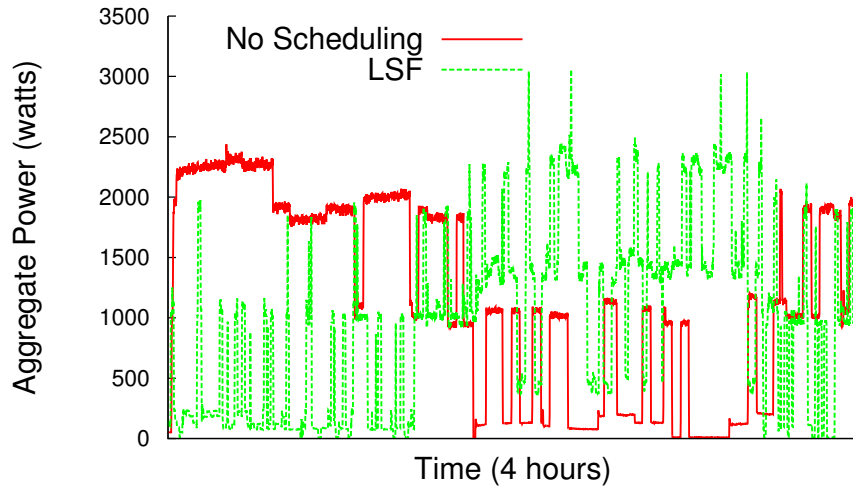


Figure 3.10. Power usage with and without LSF scheduling using using our smart home testbed with real background loads over a 4-hour period.

3.8 Conclusion

Demand-side management is challenging, since it often requires active, and often burdensome, consumer involvement. Forcing people to think about how they use power is simply not effective in encouraging broader adoption of demand-side management. Thus, we focus on quantifying the benefits of scheduling transparent background loads. We show that LSF is able to flatten household demand over each day, despite long periods of inactivity at night. Importantly, we also show that LSF is useful over shorter (4-hour) peak usage periods, where demand is “peaky” and deviates frequently and significantly from the average.

CHAPTER 4

SMARTCHARGE: CUTTING THE ELECTRICITY BILL IN SMART HOMES WITH ENERGY STORAGE

There are limitations to energy optimizations by scheduling background loads—as studied in Chapter 3—such as due to fixed duty cycles there is only so much room for switching off an appliance and flattening the demand. Also, as today’s residential pricing plans do not directly incentivize flat consumption profiles, customers have no monetary incentive for flattening. Therefore, there is need to look for alternatives to load scheduling.

Besides scheduling background loads, another way of optimizing the building’s energy footprint is by using energy storage. When employed with variable pricing plans such as time-of-use (ToU) pricing, storage can also help cut user’s electricity bill. Energy is stored in the storage during low-cost periods and the stored energy is used during high-cost periods to avoid the expensive draw from the grid. In this chapter, we first investigate the effectiveness of energy storage at lowering building’s electric bill without direct user involvement. And then we evaluate the impact of large-scale energy storage adoption on grid electricity demand.

4.1 Introduction and Motivation

The cost of generating electricity is rising. The average price of electricity for residential consumers in the United States has risen 29% over the past five years [108]. Despite energy-efficiency improvements, residential electricity demand in the U.S. has increased 49% over the last twenty years, due to a steady rise in the number

of household electrical devices. These facts combined with stagnant income growth over the past decade—down 1.9% in inflation-adjusted USD [47]—have resulted in electricity costs consuming a growing share of household budgets. The average home electricity bill now accounts for 2.8% of household income, and has risen by \$300 to \$1,419 per year over the last twenty years (in inflation-adjusted USD) [38]. Since today’s prices do not incorporate negative externalities associated with electricity generation, such as air pollution and climate change, its real cost to society is likely much higher than today’s prices reflect. Studies suggest that recent price and demand increases will continue into the foreseeable future.

Of course, the most direct way for consumers to cut their electricity bill is to simply use less electricity. Unfortunately, as the trends above indicate, rising prices have not yet motivated consumers to conserve power. Another important way to cut bills is to reduce demand peaks, which have a disproportionate affect on generation costs. Peak demands drive both capital expenses—by dictating the number of power plants, transmission lines, and substations—and operational expenses, since “peaking” generators are generally dirtier and costlier to operate than baseload generators [68]. To illustrate the impact of peak demands, Figure 4.1 shows the marginal cost of operating generators in the southeast U.S., and demonstrates that the marginal cost for generating electricity is non-linear and increases rapidly as utilities move up the dispatch stack to satisfy increasing demand [55]. Peak demands also result in significantly higher transmission losses, since these losses are proportional to the square of current. Thus, even small reductions in peak usage have a significant impact on generation costs. Recent estimates attribute 10-20% of generation costs in the U.S. to servicing only the top 100 hours of peak demand each year [83].

In an attempt to reduce peak demand, many utilities are transitioning from conventional fixed-rate pricing models, which charge a flat fee per kilowatt-hour (kWh), to new market-based schemes, e.g., real-time or time-of-use pricing, which more ac-

curately reflect electricity’s cost by raising and lowering prices during peak and off-peak periods, respectively. For instance, Illinois already requires utilities to provide residential customers the option of using hourly electricity prices based directly on wholesale prices [101], while Ontario charges residential customers based on a time-of-use scheme with three different price tiers (off-, mid-, and on-peak) each day [87]. We envision utilities widening the use of market-based pricing in the future to reduce generation costs, as demands and prices increase.

Unfortunately, market-based electricity pricing places a significant burden on consumers to continuously monitor prices, and then alter their usage to reduce costs without disrupting normal daily activities. The task is challenging, since most consumers have no idea how much power individual devices consume, and generally do not want to think about or plan their electricity usage. Thus, consumers may not respond appropriately to price changes, and the grid may not gain the cost-saving benefits of peak reduction. Further, as we show in Section 4.6.2, even if consumers respond appropriately, today’s market-based pricing plans may actually increase grid peaks (and costs) if demand is highly elastic and responsive to price changes. The difficulty in regulating demand may also discourage consumers from opting into market-based pricing plans. For instance, in Illinois, less than one percent of consumers have opted to switch from fixed-rate to market-based pricing [36].

To address the problem, we propose *SmartCharge*, an intelligent charging and discharging system that determines when and how much to store low-cost energy for use during high-cost periods based on expectations of future demand. SmartCharge’s primary benefit is that it does not require consumers to alter their electricity usage to reduce their electric bill under market-based pricing plans. Instead, SmartCharge reduces costs by determining when to switch a home between using (and storing) grid power and using previously stored power from a battery array. We frame the cost-minimization problem as a linear optimization that leverages knowledge of next-day




Figure 4.1. The marginal cost to generate electricity increases as utilities dispatch additional generators to satisfy increasing demand. Data from [55].

electricity prices and usage patterns. Since electricity prices are largely set in day-ahead markets [85], next-day prices are well-known. We predict next-day consumption by developing statistical machine learning (ML) to build a model based on important predictive metrics, such as weather, time-of-day, day-of-week, etc.

Our hypothesis is that combining SmartCharge with market-based pricing is capable of reducing electricity costs for consumers over the short- and long-term. Over the short-term, consumers save by storing energy during low-cost periods for use during high-cost periods. Over the long-term, as SmartCharge penetration increases, average prices will fall due to significant reductions in peak demand. However, as we discuss in Section 4.6.2, to attain maintain peak reduction at scale using SmartCharge, utilities will need to modify today’s market-based electricity pricing plans, which do not properly incentivize energy storage at scale. In evaluating our hypothesis, we make the following contributions:

SmartCharge Design. We detail SmartCharge’s architecture, which includes a battery array and charger, DC→AC inverter, and power transfer switch, as well as a gateway server and energy/voltage sensors to monitor home electricity consumption

and the battery array’s state of charge. We then outline the linear optimization problem the gateway server solves each day to reduce costs by switching the home’s power source between the grid and a battery array.

ML-based Consumption Prediction. Since solving SmartCharge’s optimization problem requires knowledge of next-day consumption, we develop a ML-based prediction technique that learns the unique characteristics of a home’s usage pattern over time. We show that our approach has an average error of 5.75% in a case study of a real home over a 40 day period. Our evaluation shows that solving the optimization using ML-based predictions comes within 8-12% of an oracle with perfect knowledge of next-day consumption.

Implementation and Evaluation. We evaluate SmartCharge in both simulation, using power data from real homes and existing market-based residential pricing plans, and with a small-scale prototype using a home UPS system and a few household appliances. Our results show that SmartCharge is able to reduce a typical home’s electric bill by 10-15% using realistic battery capacities. We also show that, if widely deployed, SmartCharge reduces grid demand peaks by 20%. Finally, we analyze SmartCharge’s installation and maintenance costs, and show that recent battery advancements combined with modest (and expected) price increases may make SmartCharge’s return on investment positive within the next few years.

4.2 Related Work

Daryanian et al. [43] first identified the opportunity to exploit energy storage in real-time electricity markets using a linear programming formulation similar to ours. However, their problem formulation ignores many of the battery inefficiencies that influence the realizable savings. Further, the work does not address stochastic demand in residential settings, whereas we develop machine learning techniques to accurately predict next-day consumption. Finally, we conduct experiments to analyze the peak

reduction effects of energy storage in the grid using real data, as well as analyze the ROI for installing and maintaining the system.

More recent work explores a similar problem as ours, but from different perspectives. For example, van de ven et al. [46] model the problem as a Markov Decision Process and claim that there is a threshold-based stationary cost-minimizing policy. The policy is optimal assuming that consumption is independent and identically distributed (i.i.d.). A preliminary evaluation with simulated demands following an i.i.d. distribution shows cost savings up to 40%. In contrast, we take a more experimental approach using traces of real home power usage and market-based rate plans. For the home in our case study, which has an aggregate power usage close to the average U.S. home, we show that the optimal savings is never more than 20% with realistic energy storage capacities ($< 60\text{kWh}$). Rather than solving the problem with respect to a particular demand distribution, we distill the problem to a linear program that uses our prediction model of future consumption levels

Vytelingum et al. [110] and Carpenter et al. [37] both focus on the economics of storage at scale, which we also discuss. Vytelingum et al. show that for sufficiently low adoption rates, the difference between the peak and off-peak prices approaches zero, reducing the financial incentives for installing energy storage. Similarly, in parallel with our work, Carpenter et al. also show that today's pricing schemes may increase the grid's peak demand at scale if prices do not adjust to demand. The work studies the profitability of a variety of different pricing schemes, and their effectiveness in decreasing grid demand peaks at scale. Koutsopoulos et al. [72] explore the problem from the perspective of a utility operator. In this case, the utility controls when to charge and discharge battery-based storage to minimize generation costs, assuming the marginal cost to dispatch generators is similar to Figure 4.1. In contrast to our problem, the approach is more applicable to large centralized energy storage facilities. We discuss the trade-offs between distributed and centralized energy storage in §4.7.2.

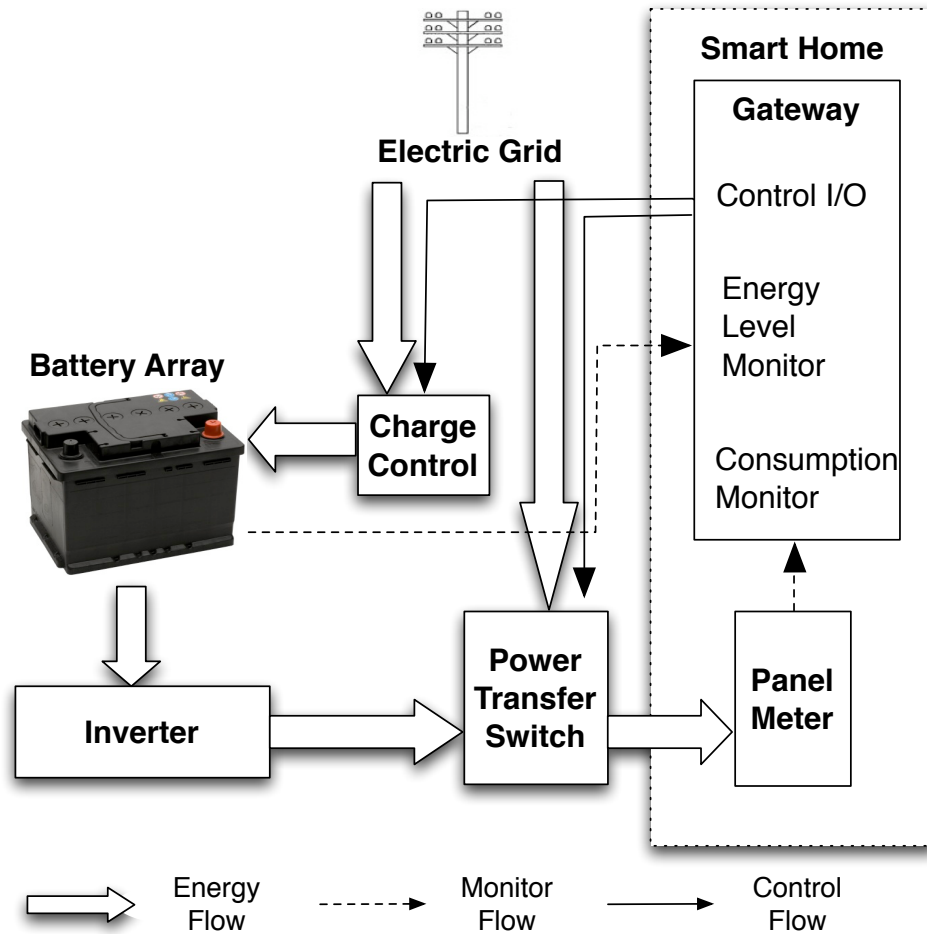


Figure 4.2. A depiction of SmartCharge’s architecture, including its battery array and charger, DC→AC inverter, power transfer switch, energy/voltage sensors, and gateway server.

4.3 SmartCharge Architecture

Figure 4.2 depicts SmartCharge’s architecture, which utilizes a power transfer switch that is able to toggle the power source for the home’s electrical panel between the grid and a DC→AC inverter connected to a battery array. A gateway server continuously monitors 1) electricity prices via the Internet, 2) household consumption via an in-panel energy monitor, and 3) the battery’s state of charge via voltage sensors. Before the start of each day, the server solves an optimization problem based on the next day’s expected electricity prices, the home’s expected consumption pattern, and

the battery array’s capacity and current state of charge, to determine when to switch the home’s power source between the grid and the battery array. The server also determines when to charge the battery array when the home uses grid power. In §4.7, we provide a detailed estimate of SmartCharge’s installation and maintenance costs based on price quotes for widely-available commercial products.

Most utilities still use fixed-rate plans for residential customers that charge a flat fee per kilowatt-hour (kWh) at all times. In the past, market-based pricing plans were not possible, since the simple electromechanical meters installed at homes had to be read manually, e.g., once per month, and were unable to record when homes consumed power. However, utilities are in the process of replacing these old meters with smart meters that enable them to monitor electricity consumption in real time at fine granularities, e.g., every hour or less. As a result, utilities are increasingly experimenting with market-based pricing plans for their residential customers. To cut electricity bills, SmartCharge relies on residential market-based pricing that varies the price of electricity within each day to more accurately reflect its cost. We expect many utilities to offer such plans in the future.

There are multiple variants of market-based pricing. Figure 4.3 shows rates over a single day for both a time-of-use (TOU) pricing plan used in Ontario, and a real-time pricing plan used in Illinois. TOU plans divide the day into a small number of periods with different rates. The price within each period is known in advance and reset rarely, typically every month or season. For example, the Ontario Electric Board divides the day into four periods (7pm-7am, 7am-11am, 11am-5pm, and 5pm-7pm) and charges either a off-peak-, mid-peak, or on-peak rate (6.2¢/kWh, 9.2¢/kWh, or 10.8 ¢/kWh) each period [87]. The long multi-hour periods and well-known rates enable consumers to plan their usage across reasonable time-scales and adopt low-cost daily routines, e.g., running the dishwasher after 7pm each day. However, while TOU pricing more




Figure 4.3. Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.

accurately reflects costs than fixed-rate pricing, it is not truly market-based since actual prices vary continuously based on supply and demand.

TOU pricing is a compromise between fixed-rate pricing and real-time pricing, where prices vary each hour (or less) and reflect the true market price of electricity. Unfortunately, real-time pricing complicates planning. Since prices may change significantly each hour, consumers must continuously monitor prices and adjust their daily routines, which may now have different costs on different days. Illinois was the first U.S. state to require utilities to offer residential consumers the option of using real-time pricing plans. To facilitate planning, Illinois utilities provide simple web pages, e.g., www.powersmartpricing.org/chart, to view next-day prices each evening. While some utilities use real-time prices not known in advance, most utilities use day-ahead market prices, which are set one day in advance. Since utilities purchase most of their electricity in day-ahead markets, e.g., 98% in New York [85], next-day prices are well-known.

SmartCharge works well with both TOU and real-time pricing plans. In either case, SmartCharge solves the optimization problem detailed in the next section at

the end of each day to determine when to switch between grid and battery power to minimize costs, based on next-day prices and expected next-day consumption. The number of periods each day—four in Ontario or twenty-four in Illinois—simply changes a parameter in the optimization’s constraints.

4.4 SmartCharge Algorithm

SmartCharge cuts electricity bills by storing energy during low-cost periods for use during high-cost periods. The total possible savings each day is a function of both the home’s rate plan and its pattern of consumption. Throughout the chapter, we use power data from a real home we have monitored for the past two years as a case study to illustrate SmartCharge’s potential benefits. The home is an average 3 bedroom, 2 bath house in Massachusetts with 1700 square feet. To measure electricity, we instrument the home with an eGauge energy meter [50], which installs in the electrical panel by wrapping two 100A current transducers around each leg of the home’s split-leg incoming power. We have monitored the home’s power consumption every second for the past two years. In 2010, the home consumed 8240kWh at a cost of \$1203.53 (or 22.6 kWh/day), while in 2011 it consumed 9732kWh at a cost of \$1339.51 (or 26.7 kWh/day). The costs are near the \$1419 average U.S. home electric bill.

4.4.1 Potential Benefits

To better understand SmartCharge’s potential for savings, it is useful to consider a worst-case scenario where 100% of the home’s consumption occurs during the day’s highest rate period. Consider our home’s hourly electricity use on January 3rd, 2012, as depicted in Figure 4.4. On this day, the home consumed 43.7 kWh, primarily due to the occupants running multiple laundry loads after returning from a holiday trip. With Ontario’s TOU plan, if the home had consumed 100% of the day’s power during the 10.8¢/kWh on-peak period, and SmartCharge shifted it all to the 6.2¢/kWh off-




Figure 4.4. Example from January 3rd with and without SmartCharge using Illinois prices from Figure 4.3.

peak period, then the maximum savings is 43%, or \$2.01 (from \$4.72 to \$2.71) for the day. Since the home did not consume 100% of its power during the on-peak period, the maximum realizable savings (if we shift all of the on-peak and mid-peak consumption to the off-peak period) is only 30%, a decrease of \$1.14 for the day (from \$3.85 to \$2.71). In practice, battery and inverter inefficiencies, which combined are ~80% efficient, reduce the savings further, to \$0.99 for the day. This per-day savings rate translates to a yearly savings of \$361.35, if the system achieves it every day.

Real-time pricing plans, as in Illinois, offer even more potential for savings, since the difference between the highest and lowest rate is significantly larger than a typical TOU plan. For example, on August 1st, 2011 in Illinois, the average rate from 2pm-7pm was 10.42¢/kWh, while the average rate from 1am-6am was 2.36¢/kWh. The highest rate of 11.9¢/kWh occurs at 4pm, and is over 5X larger than the lowest rate of 2.3¢/kWh from 2am-5am. In this case, with January 3rd’s consumption pattern and battery/inverter inefficiencies, SmartCharge is still capable of reducing costs by 59%, or \$1.78 (from \$3.02 to \$1.24). However, Figure 4.4 demonstrates that the actual savings also depend on the on-site storage capacity. In this case, with 12kWh

of usable energy storage, SmartCharge is only able to shift five hours of consumption during the highest rate daytime periods to the lowest rate nighttime periods. In particular, there is not enough capacity to store low-cost nighttime energy for use during the mid-price periods. As a result, consumption in the late morning and early evening remains unchanged. With 12kWh of storage capacity, the cost reduction falls to 32%, or \$0.96 (from \$3.02 to \$2.06) for the day.

Of course, home consumption patterns and hourly rates vary each day, which may decrease (or increase) a home's actual yearly savings. To understand why home consumption patterns are important, consider the following scenario using the Ontario TOU pricing plan. In Ontario, while SmartCharge may fully charge its battery array during the lowest rate period (7pm-7am), it may also consume that stored energy during the day's first high rate period (7am-11am). If the home expects to consume at least the battery array's entire usable capacity during the day's second high rate period (5pm-9pm), it is cost-effective, assuming ideal batteries, to fully charge the batteries during the mid-rate period (11am-5pm) when electricity costs are 17% less than in the high rate period. However, if the home only expects to use 20% of the battery's capacity during the subsequent high rate period, it is only cost-effective to charge the battery 20% during the mid-rate period, since there will be an opportunity to charge the battery further (for 33% less cost) during the next low-rate period. In this case, charging the battery more than 20% wastes money. Introducing more price tiers, as in real-time markets, complicates the problem further. As a result, we frame the problem of minimizing the daily electricity bill as a linear optimization problem.

4.4.2 Problem Formulation

While batteries exhibit numerous limitations (e.g., charging rate, capacity), inefficiencies (e.g., energy conversion efficiency, self-discharge), and non-linear relationships (e.g., between capacity, lifetime, depth of discharge, discharge rate, ambient temper-

ature, etc.), SmartCharge’s normal operation places it at the efficient end of these relationships. The system mostly charges the battery once a day during the night, which prevents stratification and extends battery lifetime by limiting the number of charge-discharge cycles. The self-discharge rate of valve-regulated absorbed glass mat (VRLA/AGM) lead-acid batteries (commonly called sealed lead-acid batteries), estimated at 1-3% per month, is insignificant, amounting to no more than \$13 per year for a 12kWh battery array with an average electricity price of 10¢/kWh. Sealed lead-acid batteries are generally 85-95% efficient, while inverters are 90-95% efficient. For SmartCharge’s battery array and inverter, we assume an energy conversion efficiency of 80%, which mirrors the efficiency rating for VRLA/AGM lead-acid batteries in a recent Department of Energy report on energy storage technologies [93]. Thus, the batteries waste 1W for every 4W they are able to store and re-use. Additionally, depth of discharge (DOD) for sealed lead-acid batteries impacts their lifetime, i.e., the number of charge-discharge cycles, due to the crystallization of lead sulfate on the battery’s metal plates. In our evaluation, we find that a DOD of 45% minimizes battery costs by balancing lifetime with usable storage capacity for a typical battery designed for home photovoltaic (PV) installations, e.g., the Sun Xtender PVX-2580L [105].

The ambient temperature and rate of discharge also have an impact on usable capacity, according to Peukert’s law. To maximize lifetime, we expect SmartCharge installations to reside in a climate-controlled room with a temperature near 25C. Rated capacity is typically based on a C/20 discharge rate, i.e., the rate of discharge necessary to deplete the battery’s capacity in 20 hours. A discharge rate higher or lower than C/20 results in less or more usable capacity, respectively. The home in our case study has averaged near 1kW per hour over the last two years, so a 20kWh battery capacity approaches this rating. As we show in §4.6, reasonable battery capacities for SmartCharge with a 45% DOD are near or above 20kWh. Finally, sealed lead-acid batteries are capable of fast charging up to a C/3 rate, i.e., charges

to full capacity in three hours [75]. In §4.6, we use a maximum charge rate of $C/4$ for the usable storage capacity, which translates to a $C/8$ rate for a battery used at 45% DOD. As we show, faster charging rates are not beneficial, since market-based pricing plans generally offer long low-rate periods for charging at night.

Given the constraints above, we frame SmartCharge’s linear optimization problem as follows. The objective is to minimize a home’s electricity bill using a battery array with a usable capacity (after accounting for its DOD) of C kWh. We divide each day into T discrete intervals of length l from 1 to T . We then denote the power charged to the battery during interval i as s_i , the power discharged from the battery as d_i , and the power consumed from the grid as p_i . We combine both the battery array and inverter inefficiency into a single inefficiency parameter e . Finally, we specify the cost per kWh over the i th interval as c_i , and the amount billed as m_i . Formally, our objective is to minimize $\sum_{i=1}^T m_i$ each day, given the following constraints.

$$s_i \geq 0, \forall i \in [1, T] \quad (4.1)$$

$$d_i \geq 0, \forall i \in [1, T] \quad (4.2)$$

$$s_i \leq C/4, \forall i \in [1, T] \quad (4.3)$$

$$\sum_{t=1}^i d_t \leq e * \sum_{t=1}^i s_t, \forall i \in [1, T] \quad (4.4)$$

$$\left(\sum_{t=1}^i s_t - \sum_{t=1}^i d_t/e \right) * I \leq C, \forall i \in [1, T] \quad (4.5)$$

$$m_i = (p_i + s_i - d_i) * I * c_i, \forall i \in [1, T] \quad (4.6)$$

The first and second constraint ensure the energy charged to, or discharged from, the battery is non-negative. The third constraint limits the battery’s maximum charging rate. The fourth constraint specifies that the power discharged from the battery

<i>Model</i>	<i>12am</i> to 7am	<i>7am</i> to 11am	<i>11am</i> to 5pm	<i>5pm</i> to 7pm	<i>7pm</i> to 12am	Average (%)
SVM Linear	14.77	27.32	46.72	18.49	47.03	29.5
SVM RBF	22.44	63.77	71.93	17.84	35.01	42.51
SVM Polynomial	4.74	4.62	6.48	7.99	5.14	5.75

Table 4.1. Average prediction error (%) over 40 day sample period for SVM with different kernel functions.

is never greater than the power charged to the battery multiplied by the inefficiency parameter. The fifth constraint states that the energy stored in the battery array, which is the difference between the energy charged to or discharged from the battery over the previous time intervals, cannot be greater than its capacity. Finally, the sixth constraint defines the price the home pays for energy during the i th interval. The objective and constraints define a linearly constrained optimization problem that is solvable using standard linear programming techniques. SmartCharge solves the problem at the beginning of each day to determine when to switch between grid and battery power. Since the approach uses knowledge of next-day consumption patterns, we next detail statistical machine learning techniques for predicting next-day consumption and quantify their accuracy for our case study home.

4.5 ML-based Demand Prediction

As discussed in §4.4, solving SmartCharge’s linear optimization problem requires *a priori* knowledge of next day consumption patterns. A simple approach to predicting consumption is to use past-predicts-future models that assume an interval’s consumption will closely match either that interval’s consumption from the previous day or the prior interval’s consumption. As we show, the approach does not work well for the multi-hour intervals in Ontario’s TOU pricing plan. Instead, we develop

statistical machine learning (ML) techniques to accurately predict consumption each interval. While our techniques have numerous applications, e.g., dispatch scheduling in microgrids, we focus solely on their application to SmartCharge in this chapter.

We experimented with a variety of prediction techniques, including Exponentially Weighted Moving Averages (EWMA), Linear Regression (LR), and Support Vector Machines (SVMs) with various kernel functions, including Linear, Polynomial, and Radial Basis Function (RBF) kernels. EWMA is a classic past-predicts-future model that predicts consumption in the next interval as a weighted sum of the previous interval’s consumption and an average of all previous intervals’ consumption. More formally, EWMA predicts the energy consumption for each interval on day k as $\hat{E}_C(k+1) = \alpha E_C(k) + (1 - \alpha)\hat{E}_C(k)$, where α is a configurable parameter that alters the weight applied to the most recent interval versus the past. Note that since each interval’s power consumption is different, we apply EWMA to each interval independently on a daily basis. As might be expected, since home consumption patterns vary largely around mealtimes, we found that predicting consumption based on the preceding interval to be highly inaccurate.

Both LR and SVM are regression techniques that combine and correlate numerous indicators (or features) of future power consumption to predict next-day usage. We experimented with a total of nine features: outdoor temperature and humidity, month, day of week, previous day power, previous interval power, as well as whether or not it is a weekend day or a holiday. We also included the EWMA prediction as an additional feature. To predict next-day temperature and humidity, we used weather forecasts from the National Weather Service available from the National Digital Forecast Database (<http://www.nws.noaa.gov/ndfd/>). To evaluate our techniques we used power data collected every second from our case study home over a period of four months from June to September 2011. For the LR and SVM models, we used the first 70 days of the data set for model training, and the last 40 days for




Figure 4.5. Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.

evaluating the model’s accuracy. We use the LibSVM library [39] to implement our LR and SVM models. Our SVM models use the *nu*-SVR regression algorithm, which we found always performed better than the ϵ -SVR algorithm [39]. For simplicity, we only predict consumption for the Ontario TOU rate periods in Figure 4.3.

Before training our model, we employed Correlation-based Feature Subset Selection (CFSS) to refine the number of input features [61]. CFSS evaluates the predictive ability of each individual feature along with the degree of redundancy between features. We apply CFSS separately for each of the five intervals, since the pattern of power consumption varies each interval. CFSS reduces the number of features in prediction model from nine to: four for 12am-7am, seven for 7am-11am, seven for 11am-5pm, six for 5pm-9pm, and five for 9pm-12am. In general, we find that more features are useful during periods with high, variable consumption.

We then experimented with multiple variations of LR models, including least squares and different regularized models (LASSO, ElasticNet, and Ridge Regression), since we found that temperature, humidity, and past data were approximately linear with respect to power consumption. However, our best performing LR model

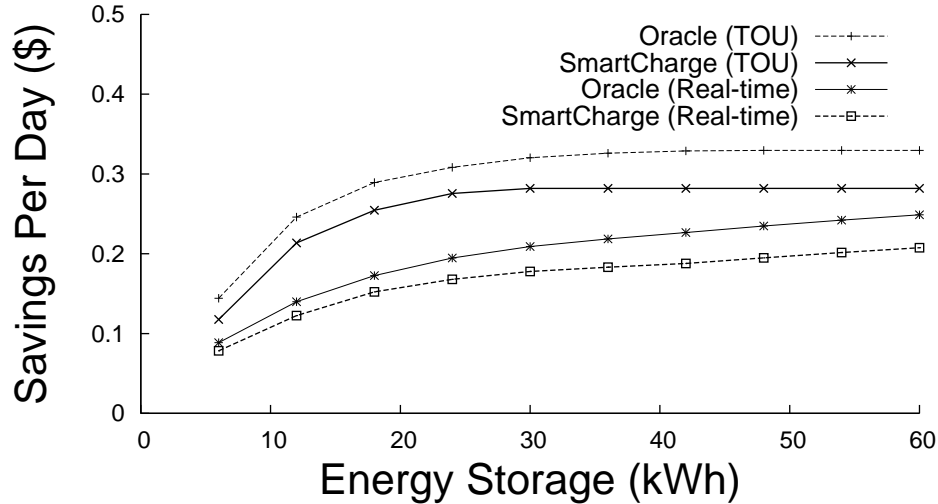


Figure 4.6. Average dollar savings per day for both real-time and TOU prices in our case study home.

(ElasticNet) had an average error of 37%. EWMA performed much better, although Figure 4.5 demonstrates its limitations in predicting future consumption. The figure shows actual power consumption each day during the first interval (12am-7am), as well as EWMA ($\alpha = 0.35$) and the SVM-Polynomial model. EWMA is unable to predict large spikes or dips in consumption before they occur. Instead, EWMA’s predictions never vary too far from the mean usage. In contrast to EWMA, the SVM approach is able to partially predict many of the spikes and dips in consumption. Over our 40 day testing period, we found that SVM-Polynomial had an average error of only 5.75%. The SVM model with the Linear and RBF kernel performed worse than EWMA, as Table 5.1 shows, with a 29.5% and 42.5% average error, respectively. As a result, in §4.6 we use SVM-Polynomial to evaluate SmartCharge.

4.6 Experimental Evaluation

To illustrate SmartCharge’s potential for savings, we use the home described in §4.4 to evaluate the savings using real hourly real-time and TOU rate plans in simu-




Figure 4.7. Average percentage savings for both real-time and TOU prices in our case study home.

lation. We also implement a small-scale SmartCharge prototype using a home UPS system and a few household appliances. For real-time prices, we use rates from June to September 2011 in the hourly day-ahead market run by the New England Independent System Operator (ISO), which operates the electricity market in our home’s region. We use historical market data publicly available that ISOs are required to publish [86]. Since we use day-ahead market prices, we have perfect knowledge of next-day prices. For TOU pricing, we use the Ontario rate plan from Figure 4.3. While our home is not located in Ontario, it lies at the same latitude and experiences a similar climate. Thus, the prices are not entirely mismatched to our home’s consumption profile. In our experiments, we vary the pricing plans and battery characteristics to see how future price trends and battery technology impact savings. To predict next-day usage, we use the SVM-Polynomial model described §4.5. Finally, to quantify the optimal savings, we compare with an oracle that has perfect knowledge of next-day consumption.

Unless otherwise noted, our experiments use home power data from the same 40 day period in late summer as the previous section. We use CPLEX, a popular inte-

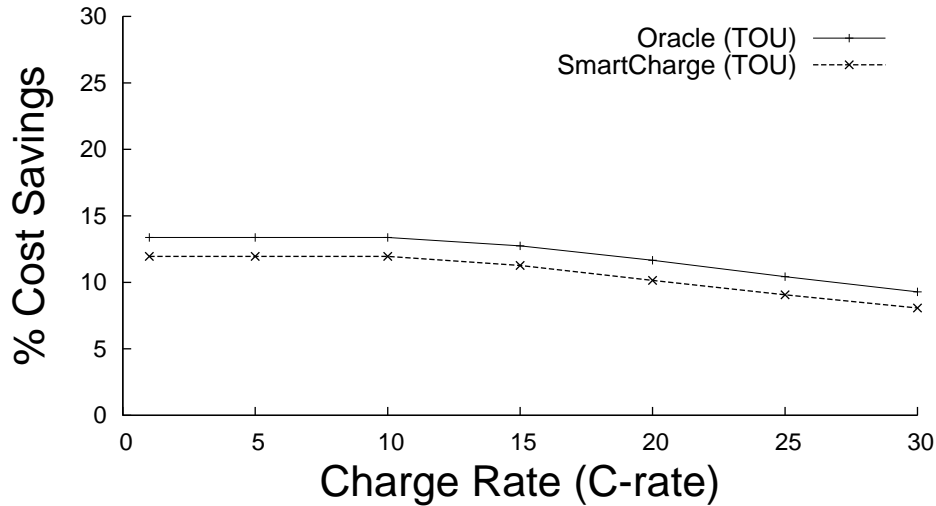



Figure 4.8. SmartCharge’s savings as a function of the charging rate for a 12kWh storage capacity.


ger and linear programming solver, to encode and solve SmartCharge’s optimization problem, given next-day prices and expected consumption levels. Note that we consider only usable storage capacity in kWh in this section, which is distinct from (and typically much less than) battery capacity. In the next section, we discuss the battery capacity necessary to attain a given storage capacity. As mentioned in §4.4, we use an energy conversion efficiency of 80% for the battery and a C/4 charging rate for the usable storage capacity.

4.6.1 Household Savings

Figure 4.6 shows the average savings per day in USD for both the real-time and TOU rate plans over the 40 day period, as a function of storage capacity, while Figure 4.7 shows the savings as a percentage of the total electricity bill. The graphs show that a storage capacity beyond 30kWh does not significantly increase savings. Further, smaller storage capacities, such as 12kWh, are also capable of reducing costs, near 10% for SmartCharge. If we extrapolate the savings over an entire year, we estimate that SmartCharge with 24kWh of storage is capable of saving \$101.59.



(a)



(b)

Figure 4.9. Varying the average electricity price (a) and the peak-to-off-peak price ratio (b) impacts savings.

Finally, the graphs show that SmartCharge’s performance is close to that of an oracle with perfect knowledge of future consumption: mispredictions only cost an estimated \$12.09 each year with 24kWh storage capacity, or near 12% of the total savings. Due to different price levels, the TOU plan saves slightly more dollars per day, while the real-time plan saves a larger percentage of the bill. As we show next, both the pricing plan and battery characteristics impact the savings. Since the savings for both the real-time and TOU rate plan are similar, for clarity we focus our remaining results on the TOU rate plan, which is more widely used today.

The experiments above assume that we use today’s battery characteristics and price levels. Of course, a more efficient battery and inverter would increase the usable storage capacity in a battery array. As the experiments above indicate, increasing storage capacity increases the savings up to a 30kWh capacity. Figure 4.8 demonstrates that the maximum charging rate has a minimal effect on savings, since the TOU rate plan (as well as the real-time plan) offer a long period of relatively low rates during the night for charging. The charging rate need only be high enough, e.g., a C/10 rate, to charge the battery over these periods. Figures 4.9(a) and (b) show how the savings change if we vary either the average price (while keeping price ratios constant) or the peak-to-off-peak price ratio (while keeping the average price constant) for a 12kWh capacity. The graphs demonstrate that, as expected, rising prices or ratios significantly impact the savings. In the former case, the relationship is linear, with a doubling of today’s average price resulting in a doubling of the savings for SmartCharge. Thus, if average electricity prices continue to rise 5% per year, as in the past, SmartCharge’s expected savings should also increase at 5% per year. In the latter case, while the savings rate decreases slowly as the ratio increases, the savings nearly doubles (up 88%) if the current ratio increases slightly from 1.6 to 2.

Finally, Figure 4.10 shows the additional savings homes are able to realize by sharing battery capacity with neighbors. Sharing is beneficial when homes exhibit peaks at different times by allowing them to share the available storage capacity. For the experiment, we use power data for a single day from a pool of 353 additional homes we monitor (described below), such that each point is an average of twenty runs with a set of k randomly chosen homes. We report both the additional dollar and percentage savings per home. We include 90% confidence intervals for the dollar savings. The experiment shows that sharing a battery array between homes results in additional savings as we increase the number of homes. As expected, more homes require more storage capacity to reap additional benefits. With 10 homes sharing

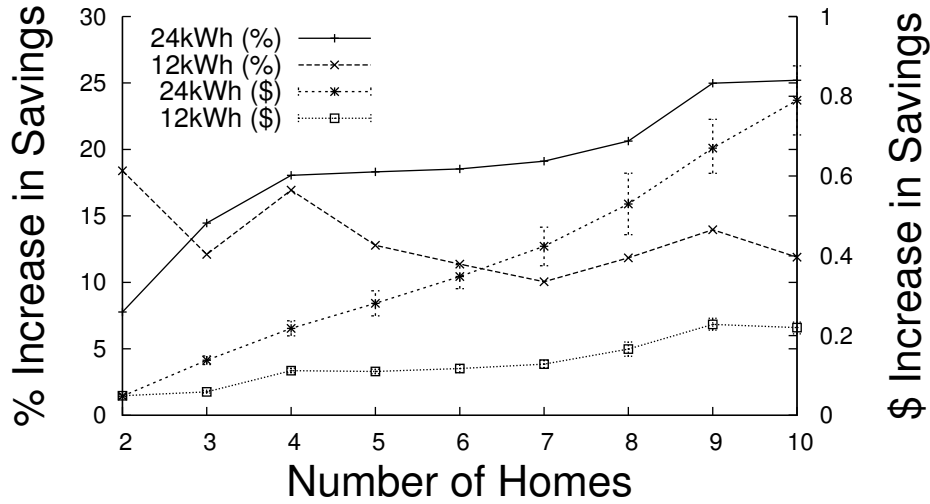



Figure 4.10. Additional savings (in % and \$) from sharing 12kWh and 24 kWh between homes.

24kWh per home, the additional savings is 25%. However, with 12kWh per home the percentage savings does not increase beyond 15% when sharing with more than four homes.


4.6.2 Grid Peak Reduction

The purpose of real-time and TOU rate plans is to lower peak electricity usage across the entire grid. We evaluate the potential grid-scale effect of SmartCharge using power data from a large sampling of homes. We gather power data at scale from thousands of in-panel energy meters that anonymously publish their data to the web. Since we do not know if the meters are installed in commercial, industrial, or residential buildings, we filter out sources that do not have typical household power levels and profiles, i.e., peak power less than 10kW and average power less than 3kW. We also filter out sources with large gaps in their data. After filtering, we select 435 homes from the available sources.

Figure 4.11(a) plots the peak power over all the homes as a function of the fraction of homes using SmartCharge with 12kWh of energy storage. The figure shows that



(a)



(b)

Figure 4.11. With 22% of homes using SmartCharge, the peak demand decreases by 20% (a) and demand flattens significantly (b).

SmartCharge is capable of reducing peak power by 20% when 22% of homes use the system, as long as the homes randomize when they begin overnight charging. If everyone begins charging at the same time, e.g., at 12am at night, the peak reduction decreases to a maximum of only 8%. Even using randomized charging, if more than 22% of consumers install SmartCharge, then the peak reduction benefits begin to decrease, due to a nighttime “rebound peak”. Once 45% of consumers use the system the evening rebound peak actually becomes larger than the original peak without SmartCharge. The same point occurs when only 24% of homes use the system without

randomized charging. Of course, the experiments assume that prices do not change in response to homes installing SmartCharge, i.e., a large fraction of homes install the system simultaneously. A more plausible and realistic scenario is that the rate of adoption slowly rises with the differential between the peak and off-peak prices. In this scenario, SmartCharge’s load shifting would alter prices in each rate period. At some point, as Vytelingum et al.[110] formally show, the price changes would make the system increasingly less attractive for new users, as the difference between peak and off-peak prices would approach zero.

We discuss SmartCharge’s economics at scale further in §4.7. Figure 4.11(b) shows grid power usage over time, with 0% and 22% of the homes using SmartCharge with randomized charging, and demonstrates how SmartCharge causes demand to “flatten” significantly. Such a peak reduction would have a profound effect on generation costs, likely lowering them by more than 20% [83]. Finally, with 22% of homes using SmartCharge, the increase in total energy usage is only 2%. The result demonstrates that the benefits of flattening likely outweigh the increased energy consumption due to battery/inverter inefficiencies.

4.6.3 Lab Prototype Results

We constructed a small-scale proof-of-concept prototype using a home UPS connected to a few common household appliances. While not typically designed for entire homes, today’s UPSs include the inverters, transfer switches, charge controllers, battery enclosure, cabling, and battery sensors necessary for a SmartCharge system in a single appliance. We chose the APC Smart-UPS 2200VA XL as our UPS, which includes software to monitor its capacity and charge/discharge state. The UPS has a usable capacity of 450Wh, but is expandable to 16kWh, at a discharge rate of 100W/s. The UPS switches to battery in roughly 25ms, which is less than the holdup time, i.e., the duration a device is able to sustain operation without power, in modern

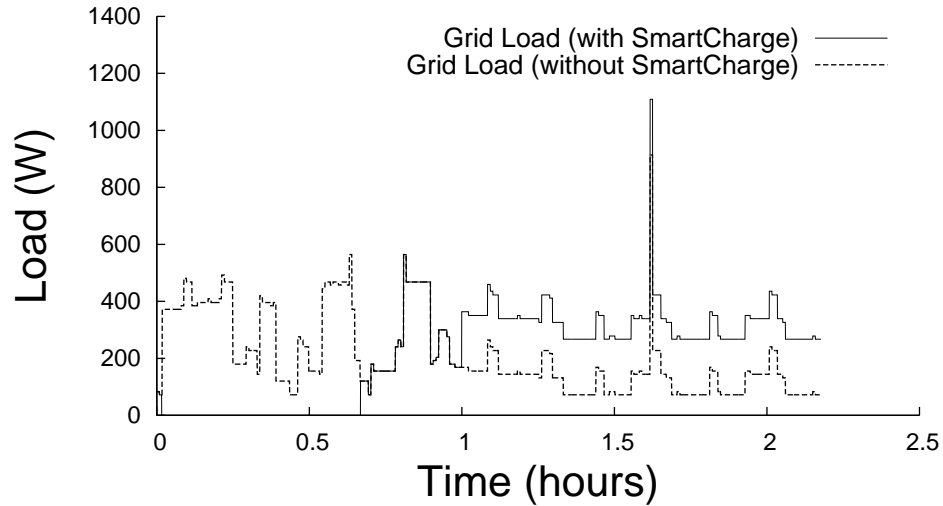


Figure 4.12. Our UPS-based prototype reduces peak usage by 69% when using a few common appliances.

power supplies. We experimented with both charging and discharging the UPS. The unit charges from 45% to 100% capacity in 80 minutes at a linear rate, and discharges in 35 minutes with an average load of 384W. We connect a refrigerator, freezer, dehumidifier, and two laptops to the UPS system. We then emulate a TOU rate plan over a two hour period, where the first hour corresponds to a peak period and the second corresponds to an off-peak period. Figure 4.12 shows that in this simple case SmartCharge uses battery power during the peak period and then switches to grid power during the off-peak period. Without SmartCharge, during the peak period the grid load was on average 298W and during the off-peak period it was 128W. With SmartCharge, the peak period has an average grid load of only 91W while the off-peak period has an average load of 324W, resulting in a 69% reduction in peak electricity consumption.

4.7 Cost-Benefit Analysis

The previous section shows that SmartCharge cuts an electric bill by 10-15% with today's market-based pricing plans. In this section, we first discuss SmartCharge's return on investment (ROI), including its installation and maintenance costs, and then discuss its advantages over centralized energy storage. We ground our discussion using price quotes, primarily from the altE store (<http://www.altestore.com>), for widely-available commercial products.

4.7.1 Return-on-Investment

In many instances, homes already have the necessary infrastructure to implement SmartCharge. For example, many homes in developing countries already utilize UPSs because of instability in the power grid. As we discuss below, in the future, homes with photovoltaic (PV) systems may require on-site energy storage to balance an intermittent supply with demand without the aid of net metering. Batteries in electric vehicles (EVs) could also serve as energy storage. In each case, the homes already include the required infrastructure and battery capacity to implement SmartCharge. Since the homes would not need new infrastructure, the ROI is positive in these cases. Below, we discuss the ROI for homes that do not already have the necessary infrastructure.

Table 5.2 shows cost estimates for purchasing and installing SmartCharge's components. For the inverter, we assume Apollo Solar's True Sinewave Inverter, which combines an inverter, battery charger, and transfer switch into a single appliance. To read battery state and control the appliance, we attach an additional communications gateway available for the inverter. Numerous home energy meters are available: The Energy Detective (TED) is a popular choice and costs \$200. Nearly any server is adequate to support SmartCharge's software. We use an embedded DreamPlug server at a cost of \$159 as the gateway in the homes we now monitor. To hold the battery

array, we assume two MNEBE-C 12-battery modular enclosures. Finally, we estimate \$200 for cabling and a day's labor at \$500 for installation. The total estimated cost, excluding batteries, is \$4871.

Of course, SmartCharge's largest expense is its battery array. Sealed VRLA/AGM lead-acid batteries are the dominant battery technology for stationary home UPSs and PV installations, due to their combination of low price, high efficiency, and low self-discharge rate. By contrast, lithium ion batteries, while lighter and more appropriate for EVs, are much more expensive. We use, as an example, the Sun Xtender PVX-2580L with a 3kWh rated capacity (at a C/20 discharge rate), which costs \$570 [105] and is designed for deep-cycle use in home PV systems. The battery's manual specifies its lifetime as a function of its number of charge-discharge cycles and the DOD each cycle. We use the data to estimate the yearly cost of batteries—in \$/kWh of *usable* storage capacity—as a function of the depth of discharge (Figure 4.13) amortized over their lifetime, assuming SmartCharge's typical single charge-discharge cycle per day. The usable storage capacity takes DOD into account: a battery rated for 10kWh operated at 50% depth of discharge has a usable capacity of only 5kWh. Figure 4.13 demonstrates that cost begins to increase rapidly after a 45% DOD, with an estimated cost of \$118/kWh of usable capacity.

In the U.S., SmartCharge likely qualifies for a Residential Renewable Energy Tax Credit, reducing its cost by 30%. Additionally, U.S. state and local governments offer an assortment of tax incentives for energy-efficiency improvements [44], which we estimate lower costs by 20%. Despite the advantages, today's lead-acid batteries are still too expensive to produce a positive ROI at current electricity prices. For instance, while 24kWh of usable storage capacity saves \$91.25 per year using the Ontario TOU rate plan, batteries alone would cost \$1416 per year assuming the take breaks above. However, recent advancements in battery technology promise to dramatically reduce battery costs in the near future. Lead-carbon batteries have

<i>Component</i>	<i>Total</i>
Inverter	\$2099.00
Battery Charger	-
Transfer Switch	-
Inverter Gateway	\$287.00
Energy Monitor	\$200.00
Server	\$159.00
Battery Enclosure	\$1426.00
Cabling	\$200.00
Labor	\$500.00
Total	\$4871.00

Table 4.2. Estimated cost breakdown for installing SmartCharge’s supporting infrastructure.

an expected lifetime 10x longer than today’s sealed lead-acid batteries at roughly the same cost [49, 56, 93]. Figure 4.14 shows the extended lifetime using data from recent tests conducted at Sandia National Labs comparing today’s sealed lead-acid battery and a new lead-carbon battery (the UltraBattery) [93].

Lead-carbon batteries combined with modest and expected price increases (25%) and peak-to-off-peak ratios (25%) would produce a positive ROI for SmartCharge in a few years. Assuming this scenario, Figure 4.15 plots SmartCharge’s yearly expense, including battery and infrastructure costs (amortized over 20 years), along with its estimated yearly savings for our case study home, as a function of usable storage capacity. Note that our ROI estimates do not include the savings from lowering generation costs for *all* homes by reducing peak demands. As Figure 4.11 shows, enabling only 22% of homes with SmartCharge would dramatically reduce peak demands, and, hence, generation costs for *all* homes, even those that have not invested in the system. Since all homes benefit from lower prices, utilities may consider subsidies that spread costs across all consumers, which for 22% of homes would lower costs by nearly 5X.

Alternatively, utilities might consider modifying their pricing plans to incentivize SmartCharge in all homes by increasing the fraction of the bill based on peak usage. While many utilities charge large consumers based on their peak usage over

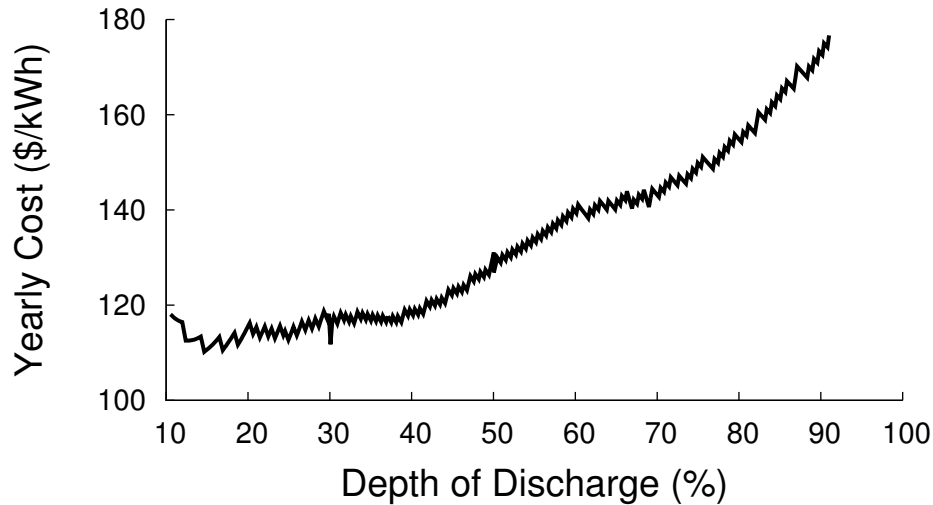


Figure 4.13. Amortized cost per kWh as a function of depth of discharge.

a day or month [30], residential bills typically do not include such a charge. Incorporating a substantial peak usage charge in electric bills would prevent the large rebound peaks in Figure 4.11 by directly incentivizing homes to flatten demand, rather than shift as much demand as possible to low-cost periods (causing the rebound peak). With market-based plans that only charge per-kWh, as more consumers install SmartCharge and shift their demand to low-cost periods, the price difference between the low-cost and high-cost periods would lessen to reflect the new demand distribution, thus lowering the ROI and discouraging additional homes from installing the system. A substantial peak-usage charge would maintain the financial incentives and continue to flatten demand (and prevent rebound peaks) as the fraction of SmartCharge-enabled homes approaches 100%.

A full discussion of SmartCharge’s impact on the economics of electricity generation is outside the scope of this chapter. However, it is clear that today’s market-based pricing plans assume that the price elasticity of electricity demand is low, i.e., changes in price do not have a significant impact on demand. SmartCharge fundamentally changes this fact by making demand nearly fully elastic with price.




Figure 4.14. Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [93].

4.7.2 Distributed vs. Centralized

Utilities have already begun to deploy large, centralized battery arrays to reduce peak usage and integrate more wind and solar farms, which require substantial energy storage to match an intermittent supply with variable demand. However, distributing battery storage throughout the grid has a number of inherent advantages over a centralized approach. In particular, home energy storage may serve as backup power during extended blackouts, lessening the economic impact of power outages and promoting a more stable grid. A centralized system also introduces a single point of failure. Further, substantial home energy storage may be a catalyst for implementing microgrids, where matching supply and demand is difficult without an energy buffer. Storing energy at its point-of-use also reduces transmission losses by eliminating losses incurred from generator to battery array.

Finally, perhaps the most important argument for installing many distributed battery arrays in homes, rather than large centralized arrays, is to encourage distributed generation without relying on net metering. While today’s PV installations typically use net metering to offset costs by selling energy back to the grid, it is not a scalable




Figure 4.15. SmartCharge’s projected yearly expense and savings assuming recent battery advancements.

long-term solution. Injecting significant quantities of power into the grid from unpredictable and intermittent renewables has the potential to destabilize the grid by making it difficult to balance supply and demand. SmartCharge provides an alternative to net metering to offset costs in home PV systems that use batteries instead of net metering. We are currently studying how to include renewables in SmartCharge’s algorithm. Our initial results suggest that homes with PV installations also benefit from SmartCharge.

4.8 Conclusion

In this chapter, we explore how to lower electric bills using SmartCharge by storing low-cost energy for use during high-cost periods. We show that typical savings today are 10-15% per home with the potential for significant grid peak reduction (20% with our data). Finally, we analyze SmartCharge’s costs, and show that recent battery advancements combined with an expected rise in electricity prices may make

SmartCharge's return on investment positive for the average home within the next few years.

CHAPTER 5

GREENCHARGE: MANAGING RENEWABLE ENERGY IN SMART BUILDINGS

Renewable energy integration can further boost savings from energy storage under variable pricing. Therefore, we now extend SmartCharge’s architecture and algorithm to include on-site renewables with energy storage and grid energy to minimize electricity bills.

5.1 Introduction and Motivation

Buildings today consume more energy (41%) than either of society’s other broad sectors of energy consumption—industry (30%) and transportation (29%) [7]. As a result, even small improvements in building energy efficiency, if widely adopted, hold the potential for significant impact. The vast majority (70%) of building energy usage is in the form of electricity, which, due to environmental concerns, is generated at “dirty” power plants far from population centers. As a result, nearly half (47%) of energy use in residential buildings is lost in electricity transmission and distribution (T&D) from far-away power plants to distant homes [7]. An important way to decrease both T&D losses and carbon emissions is through distributed generation (DG) from many small on-site renewable energy sources deployed at individual buildings and homes. Unfortunately, in practice, DG has significant drawbacks that have, thus far, prevented its widespread adoption. In particular, DG primarily relies on solar panels and wind turbines that generate electricity intermittently based on uncontrollable and changing environmental conditions. Since the energy consumption density,

in kilowatt-hours (kWh) per square foot, is higher than the energy generation density of solar and wind deployments at most locations, buildings must still rely heavily on the electric grid for power.

Another major drawback of DG is that large centralized power plants benefit from economies-of-scale that cause their generation costs, even accounting for T&D losses, to be significantly lower than DG. As a result, today's DG deployments rely heavily on net metering—where buildings sell the unused energy they produce back to the utility company—to offset their cost relative to grid energy. DG is a much less financially attractive where net metering is not available. Net metering laws and regulations vary widely across states—it is not available in four states and the regulations are weak in many others [96]. Further, even where available, states typically place low caps on both the total number of participating consumers and the total amount of energy contributed per customer [96]. After exceeding these caps, utilities are no longer required to accept excess power from DG deployments. As one example, the state of Washington caps the total number of participating consumers at 0.25% of all customers. One reason for the strict laws limiting DG's contribution is that injecting significant quantities of power into the grid from unpredictable renewables at large scales has the potential to destabilize the grid by making it difficult, or impossible, for utilities to balance supply and demand. Large baseload power plants that produce the majority of grid energy are simply not agile enough to scale their own generation up and down to offset significant fractions of renewable generation.

Thus far, current laws have not been an issue, since today's energy prices do not make DG financially attractive enough to reach even these low state caps. However, more widespread adoption of DG is critical to meeting existing goals for increasing the fraction of environmentally-friendly renewable energy sources. For example, the Renewables Portfolio Standard targets 25% of electricity generation from intermittent renewables [44], while California's Executive Order S-21-09 in California calls for

33% of generation from renewables by 2020 [104]. Given current laws, if and when DG becomes more widespread, buildings will have to look beyond net metering to balance on-site energy generation and consumption, while also reducing DG’s costs. We envision consumers using a combination of on-site renewables, on-site battery-based energy storage, and the electric grid to satisfy their energy requirements, while also balancing local supply and demand.

In parallel, we envision the adoption of market-based electricity pricing providing a new opportunity to recoup the loss of net metering revenue, while also introducing new financial incentives for DG where net metering is not available. Many utilities are transitioning from conventional fixed-rate pricing models, which charge a flat fee per kilowatt-hour (kWh), to new market-based schemes, e.g., real-time or time-of-use pricing, which more accurately reflect electricity’s cost by raising and lowering prices during peak and off-peak periods, respectively. Satisfying peak demands is significantly more expensive ($\sim 10x$) than off-peak demands, since peak demands drive both capital expenses—by dictating the number of power plants, transmission lines, and substations—and operational expenses—“peaking” generators are generally dirtier and costlier to operate than baseload generators [68]. For instance, Illinois already requires utilities to provide residential customers the option of using hourly electricity prices based directly on wholesale prices [101], while Ontario charges residential customers based on a time-of-use scheme with three different price tiers (off-, mid-, and on-peak) each day [87].

The primary contribution of this chapter is a new system architecture and control algorithm, called GreenCharge for managing on-site renewables, on-site energy storage, and grid energy in buildings to minimize grid energy costs for market-based electricity prices. Our system determines both the fraction of power to consume from the grid versus on-site battery-based energy storage, as well as when and how much to charge battery-based storage using grid energy. The primary inputs to our

control algorithm are 1) the battery’s current energy level, 2) a prediction of future solar/wind energy generation, 3) a prediction of future energy consumption patterns, and 4) market-based electricity prices. The output is the amount of power to consume from the grid, as well as the power to discharge or charge the battery from renewables or the grid, over each rate period. We evaluate our system using a collection of real data sets, including power consumption data from a real home, energy harvesting data from a solar and wind deployment, National Weather Service (NWS) forecast data, and TOU pricing data from Ontario, Canada.

We compare GreenCharge with two other approaches: i) an approach from initial work, called SmartCharge [78], that only uses energy storage without renewables to reduce prices and ii) an oracle with perfect knowledge of future energy consumption and generation. GreenCharge extends our initial work on SmartCharge in multiple ways. First, SmartCharge only optimized prices by determining when and how much to charge a battery at off-peak hours. GreenCharge extends this idea to account for intermittent renewable generation, e.g., by using forecast-based models to predict future energy harvesting—a major enhancement to SmartCharge. In addition, this chapter includes new material describing our use of communication protocols in implementing a GreenCharge prototype, as well as a revised linear programming formulation and algorithm that accounts for renewable generation. Finally, our work includes substantial experiments to understand the impact of adding renewables to SmartCharge. Our results show that GreenCharge saves an additional 10-15% on electric bills beyond SmartCharge, which only uses a battery, and is near the performance of an oracle with perfect future knowledge.

5.2 Related Work

Daryanian et al. [43] first identified the opportunity to exploit energy storage in real-time electricity markets using a linear programming formulation similar to

ours. However, their problem formulation ignores many of the battery inefficiencies that influence the realizable savings. Further, the work does not address stochastic demand in residential settings, whereas we develop machine learning techniques to accurately predict next-day consumption. In addition, we also conduct experiments to analyze the peak reduction effects of energy storage in the grid using real data, as well as analyze the ROI for installing and maintaining the system. Finally, we include renewables into the system, as well as use a model for predicting renewable generation, which has not been considered in prior work to the best of our knowledge.

More recent work explores a similar problem as ours, but from different perspectives and without renewable generation. For example, van de ven et al. [46] model the problem as a Markov Decision Process and claim that there is a threshold-based stationary cost-minimizing policy. The policy is optimal assuming that consumption is independent and identically distributed (i.i.d.). A preliminary evaluation with simulated demands following an i.i.d. distribution shows cost savings up to 40%. In contrast, we take a more experimental approach using traces of real home power usage, solar panel generation, and market-based rate plans. For the home in our case study, which has an aggregate power usage close to the average U.S. home, we show that the optimal savings is never more than 20% with realistic energy storage capacities ($< 60\text{kWh}$). Rather than solving the problem with respect to a particular demand distribution, we distill the problem to a linear program that uses our prediction model of future consumption levels

Vytelingum et al. [110] and Carpenter et al. [37] both focus on the economics of storage at scale, which we also discuss. Vytelingum et al. show that for sufficiently low adoption rates, the difference between the peak and off-peak prices approaches zero, reducing the financial incentives for installing energy storage. Similarly, in parallel with our work, Carpenter et al. also show that today's pricing schemes may increase the grid's peak demand at scale if prices do not adjust to demand. The work studies

the profitability of a variety of different pricing schemes, and their effectiveness in decreasing grid demand peaks at scale. Koutsopoulos et al. [72] explore the problem from the perspective of a utility operator. In this case, the utility controls when to charge and discharge battery-based storage to minimize generation costs, assuming the marginal cost to dispatch generators increases super-linearly as utilities move up the dispatch stack to satisfy increasing demand. In contrast to our problem, the approach is more applicable to large centralized energy storage facilities. We discuss the trade-offs between distributed and centralized energy storage in §5.7.2.

5.3 GreenCharge Architecture

Figure 5.1 depicts GreenCharge’s architecture, which utilizes a power transfer switch that is able to toggle the power source for the home’s electrical panel between the grid and a DC→AC inverter connected to a battery array. On-site solar panels or wind turbines connect to, and charge, the battery array. A smart gateway server continuously monitors 1) electricity prices via the Internet, 2) household consumption via an in-panel energy monitor, 3) renewable generation via current transducers, and 4) the battery’s state of charge via voltage sensors. Our SmartCharge system, which we compare against in this work, utilizes the same architecture, but does not use renewables [78].

Before the start of each day, the server solves an optimization problem based on the next day’s expected electricity prices, the home’s expected consumption and generation pattern, and the battery array’s capacity and current state of charge, to determine when to switch the home’s power source between the grid and the battery array. The server also determines when to charge the battery array when the home uses grid power. In §5.7, we provide a detailed estimate of GreenCharge’s installation and maintenance costs based on price quotes for widely-available commercial products.




Figure 5.1. A depiction of GreenCharge’s architecture, including its battery array and charger, DC→AC inverter, solar and/or wind energy sources, power transfer switch, energy/voltage sensors, and gateway server.

5.3.1 Network Communication and Sensing

One challenge with instantiating GreenCharge’s architecture is transmitting sensor data about energy consumption, energy generation, and battery status to GreenCharge’s smart gateway server in real time. The simplest way to measure energy consumption and generation is to wrap current transducers (CT) around wires in the building’s electrical panel. In this case, two CTs are necessary to cover both legs of a building’s split leg input power from the grid, as well as a CT for each connection to a renewable source. Note that CTs use the Hall Effect [60] for measuring voltage

and current, and only require wrapping a sensor around a wire without cutting any wires. CTs must be installed in the panel, since this is the only place in the building that has the incoming grid lines exposed for sensors. Since electrical panels are often in remote corners of a building, transmitting readings wirelessly is difficult. While wired Ethernet is an attractive option, it requires running an Ethernet cable from GreenCharge’s gateway server to the electrical panel. Instead, to overcome wireless interference and prevent running an Ethernet cable into the panel, GreenCharge uses a powerline-based communication protocol to transmit readings to the server.

Multiple types of powerline-based communication protocols exist. The most common are X10, Insteon, and HomePlug. X10 is by far the oldest protocol, having been developed in 1975; it is primarily used for controlling applications, which only requires sending brief, short control messages. Unfortunately, X10 has severe bandwidth limitations (a maximum of 20bps) and reliability problems, which make it undesirable for continuous real-time sensing. The bandwidth limitations alone prevent X10 from being used to continuously sense multiple data sources. Since powerline is a broadcast network, the 20bps bandwidth is across *all* devices. In addition to the bandwidth limitations, the protocol has no acknowledgements, so it is impossible to detect packet losses and retransmit. Further, powerline noise caused by switched mode power supplies results in substantial losses with X10 in most buildings. In our own prototype, we initially used the Energy Detective (TED) power meter for monitoring electricity consumption and generation at the electrical panel. However, we discovered that the meter uses an unreliable X10-like protocol that experiences communication problems while sending data over the powerline due to sensitivity to noise. While the display blinks orange when the problems occur, the data masks the problem by always recording the last power reading as the current power reading.

Insteon is an improvement to X10 that includes acknowledgements, retransmissions, and optimizations to overcome powerline noise. However, Insteon still has band-

width limitations that, in practice, reduce its maximum rate to near 180bps [118]. While useful for controlling devices via the powerline, it is still insufficient for continuous real-time sensing of multiple data sources. Thus, in our own prototype we chose a power meter that uses the HomePlug Ethernet-over-powerline protocol. Unlike Insteon and X10, Homeplug was initially designed to stream high definition audio and video data from the Internet to televisions. As a result, it was designed from the outset to support high-bandwidth applications. HomePlug modems exist that are capable of transmitting up to 200Mbps. Since HomePlug simply implements Ethernet over the powerline, it can support a standard TCP stack to ensure reliable communication. Our prototype uses an eGauge power meter [50], and uses HomePlug to continuously transmit power consumption and generation readings over a building's powerline to GreenCharge's gateway server. Below, we discuss how the server gets current market prices for electricity.

5.3.2 Market-based Electricity Pricing

Most utilities still use fixed-rate plans for residential customers that charge a flat fee per kilowatt-hour (kWh) at all times. In the past, market-based pricing plans were not possible, since the simple electromechanical meters installed at homes had to be read manually, e.g., once per month, and were unable to record when homes consumed power. However, utilities are in the process of replacing these old meters with smart meters that enable them to monitor electricity consumption in real time at fine granularities, e.g., every hour or less. As a result, utilities are increasingly experimenting with market-based pricing plans for their residential customers. To cut electricity bills, GreenCharge relies on residential market-based pricing that varies the price of electricity within each day to more accurately reflect its cost. We expect many utilities to offer such plans in the future.




Figure 5.2. Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.

There are multiple variants of market-based pricing. Figure 5.2 shows rates over a single day for both a time-of-use (TOU) pricing plan used in Ontario, and a real-time pricing plan used in Illinois. TOU plans divide the day into a small number of periods with different rates. The price within each period is known in advance and reset rarely, typically every month or season. For example, the Ontario Electric Board divides the day into four periods (7pm-7am, 7am-11am, 11am-5pm, and 5pm-7pm) and charges either a off-peak-, mid-peak, or on-peak rate (6.2¢/kWh, 9.2¢/kWh, or 10.8 ¢/kWh) each period [87]. The long multi-hour periods and well-known rates enable consumers to plan their usage across reasonable time-scales and adopt low-cost daily routines, e.g., running the dishwasher after 7pm each day. However, while TOU pricing more accurately reflects costs than fixed-rate pricing, it is not truly market-based since actual prices vary continuously based on supply and demand.

TOU pricing is a compromise between fixed-rate pricing and real-time pricing, where prices vary each hour (or less) and reflect the true market price of electricity. Unfortunately, real-time pricing complicates planning. Since prices may change significantly each hour, consumers must continuously monitor prices and adjust their

daily routines, which may now have different costs on different days. Illinois was the first U.S. state to require utilities to offer residential consumers the option of using real-time pricing plans. While some utilities use real-time prices not known in advance, most utilities use day-ahead market prices, which are set one day in advance. Since utilities purchase most of their electricity in day-ahead markets, e.g., 98% in New York [85], next-day prices are well-known.

There are many possible ways for GreenCharge’s gateway server to monitor prices in real time. In the simplest case, utilities can provide simple web pages with current prices. For example, Illinois utilities are already required to do this, e.g., www.powersmartpricing.org/chart posts next-day prices each evening. Utilities may also use explicit protocols to “push” prices to GreenCharge’s gateway server whenever they change. For example, utilities could run publish/subscribe protocols that interact with smart meters to broadcast price changes. In this case GreenCharge’s gateway server could interact with a building’s local smart meter to discover prices. Authors in [67], [57] propose to combine IP multicast and publish-subscribe technologies to scale real-time price broadcast to millions of users for Ecogrid [48]. When using smart meters, utilities could disseminate prices using the smart meter’s communication protocol, e.g., often cellular wireless or wired powerline, rather than the public Internet.

Transactive control system, presented in [62], proposes another way of price dissemination in smart grids. In transactive control, responsive demand assets are controlled by a single, shared, price-like value signal. It defines a hierarchical node structure and the signal path through these nodes, and includes the predicted day-ahead price values. Alternatively, IEC 61850 ([11]), which has been used between DER (Distributed Energy Resources) plants for energy and price information exchange, can be extended for price exchange in smart grids. [77] presents a survey of

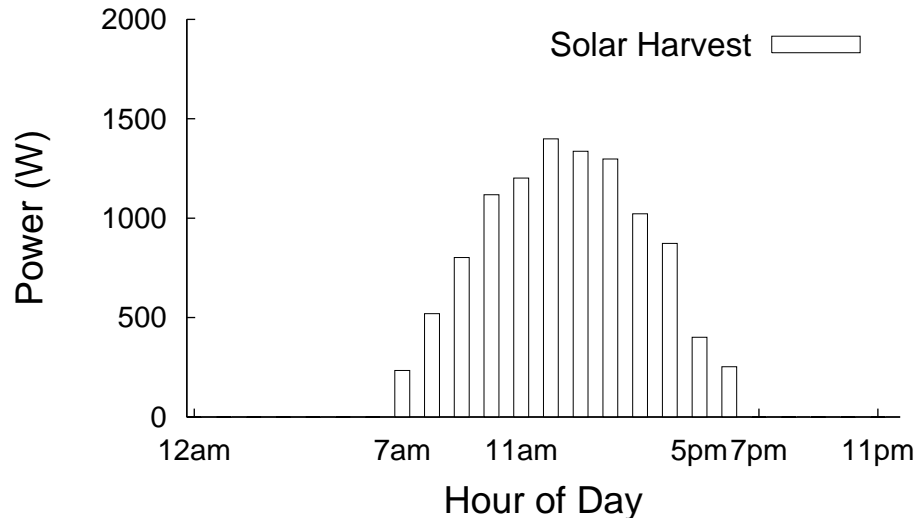


Figure 5.3. Example solar harvest data from a day in August.

a set of existing communication protocols. The report also analyzes suitability of the surveyed protocols for their application in real-time price exchange.

GreenCharge is compatible with any method above for retrieving real-time prices, and works well with both TOU and real-time pricing plans. In either case, GreenCharge solves the optimization problem detailed in the next section at the end of each day to determine when to switch between grid and battery power to minimize costs, based on next-day prices and expected next-day consumption. The number of periods each day—four in Ontario or twenty-four in Illinois—simply changes a parameter in the optimization’s constraints.

5.4 GreenCharge Algorithm

GreenCharge cuts electricity bills by combining on-site renewable generation with energy storage that stores energy during low-cost periods for use during high-cost periods. As discussed earlier, GreenCharge extends our SmartCharge system that only uses energy storage to cut electricity bills without renewables. The total possible savings each day is a function of both the home’s rate plan and its pattern

of generation and consumption. Throughout the chapter, we use power data from a real home we have monitored for the past two years as a case study to illustrate GreenCharge’s potential benefits. The home is an average 3 bedroom, 2 bath house in Massachusetts with 1700 square feet. To measure electricity, we instrument the home with an eGauge energy meter [50], which installs in the electrical panel by wrapping two 100A current transducers around each leg of the home’s split-leg incoming power. We have monitored the home’s power consumption every second for the past two years. In 2010, the home consumed 8240kWh at a cost of \$1203.53 (or 22.6 kWh/day), while in 2011 it consumed 9732kWh at a cost of \$1339.51 (or 26.7 kWh/day). The costs are near the \$1419 average U.S. home electric bill. Separately, we have deployed solar panels to study variation in solar power generation. Figure 5.3 depicts power generation from a sunny day.

5.4.1 Potential Benefits

To better understand GreenCharge’s potential for savings, it is useful to consider a worst-case scenario where 100% of the home’s consumption occurs during the day’s highest rate period. Figure 5.4 then compares GreenCharge using renewable production from Figure 5.3 with a home has only energy storage but not renewables (labeled SmartCharge), and home with no energy storage or renewables. Now consider our home’s hourly electricity use on January 3rd, 2012, as depicted in Figure 5.4 in red. On this day, the home consumed 43.7 kWh, primarily due to the occupants running multiple laundry loads after returning from a holiday trip. With Ontario’s TOU plan, if the home had consumed 100% of the day’s power during the 10.8¢/kWh on-peak period, and all consumption was shifted to the 6.2¢/kWh off-peak period, then the maximum savings is 43%, or \$2.01 (from \$4.72 to \$2.71) for the day. Since the home did not consume 100% of its power during the on-peak period, the maximum realizable savings (if we shift all of the on-peak and mid-peak consumption to the

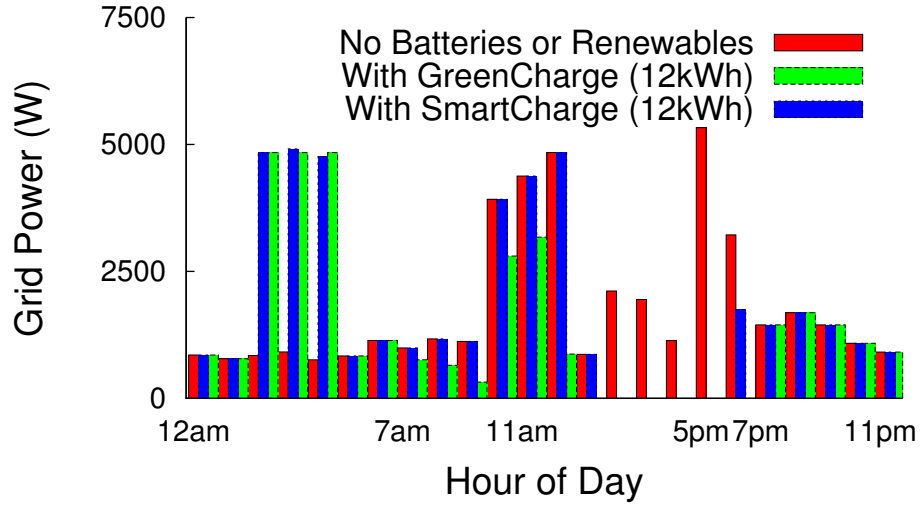


Figure 5.4. Example from January 3rd with and without GreenCharge using Illinois prices from Figure 5.2.

off-peak period) is only 30%, a decrease of \$1.14 for the day (from \$3.85 to \$2.71). In practice, battery and inverter inefficiencies, which combined are ~80% efficient, reduce the savings further, to \$0.99 for the day. Finally, if we then add in the 10.5kW generated by renewables the savings increases by \$0.93 to \$1.92. This per-day savings rate translates to a yearly savings of \$702, if the system achieves it every day.

Real-time pricing plans, as in Illinois, offer even more potential for savings, since the difference between the highest and lowest rate is significantly larger than a typical TOU plan. Of course, energy consumption and generation patterns, as well as hourly rates vary each day, which may decrease (or increase) a building’s actual yearly savings. To understand why energy consumption and generation patterns are important, consider the following scenario using the Ontario TOU pricing plan. In Ontario, while GreenCharge may fully charge its battery array during the lowest rate period (7pm-7am), it may also consume that stored energy during the day’s first high rate period (7am-11am). If the home expects to consume at least the battery array’s entire usable capacity, even when accounting for renewable generation, during the day’s second high rate period (5pm-9pm), it is cost-effective, assuming ideal batteries, to

fully charge the batteries during the mid-rate period (11am-5pm) when electricity costs are 17% less than in the high rate period. However, if the home only expects to use 20% of the battery’s capacity during the subsequent high rate period, e.g., because renewables will generate some power during this time, it is only cost-effective to charge the battery 20% during the mid-rate period, since there will be an opportunity to charge the battery further (for 33% less cost) during the next low-rate period. In this case, charging the battery more than 20% wastes money. Introducing more price tiers, as in real-time markets, complicates the problem further. As a result, we frame the problem of minimizing the daily electricity bill as a linear optimization problem.

5.4.2 Problem Formulation

While batteries exhibit numerous limitations (e.g., charging rate, capacity), inefficiencies (e.g., energy conversion efficiency, self-discharge), and non-linear relationships (e.g., between capacity, lifetime, depth of discharge, discharge rate, ambient temperature, etc.), GreenCharge’s normal operation places it at the efficient end of these relationships. The system mostly charges the battery once a day during the night, which prevents stratification and extends battery lifetime by limiting the number of charge-discharge cycles. The self-discharge rate of valve-regulated absorbed glass mat (VRLA/AGM) lead-acid batteries (commonly called sealed lead-acid batteries), estimated at 1-3% per month, is insignificant, amounting to no more than \$13 per year for a 12kWh battery array with an average electricity price of 10¢/kWh. Sealed lead-acid batteries are generally 85-95% efficient, while inverters are 90-95% efficient. For GreenCharge’s battery array and inverter, we assume an energy conversion efficiency of 80%, which mirrors the efficiency rating for VRLA/AGM lead-acid batteries in a recent Department of Energy report on energy storage technologies [93]. Thus, the batteries waste 1W for every 4W they are able to store and re-use. Additionally, depth of discharge (DOD) for sealed lead-acid batteries impacts their lifetime, i.e., the

number of charge-discharge cycles, due to the crystallization of lead sulfate on the battery’s metal plates. In our evaluation, we find that a DOD of 45% minimizes battery costs by balancing lifetime with usable storage capacity for a typical battery designed for home photovoltaic (PV) installations, e.g., the Sun Xtender PVX-2580L [105].

The ambient temperature and rate of discharge also have an impact on usable capacity, according to Peukert’s law. To maximize lifetime, we expect GreenCharge installations to reside in a climate-controlled room with a temperature near 25C. Rated capacity is typically based on a C/20 discharge rate, i.e., the rate of discharge necessary to deplete the battery’s capacity in 20 hours. A discharge rate higher or lower than C/20 results in less or more usable capacity, respectively. The home in our case study has averaged near 1kW per hour over the last two years, so a 20kWh battery capacity approaches this rating. As we show in §5.6, reasonable battery capacities for GreenCharge with a 45% DOD are near or above 20kWh. Finally, sealed lead-acid batteries are capable of fast charging up to a C/3 rate, i.e., charges to full capacity in three hours [75]. In §5.6, we use a maximum charge rate of C/4 for the usable storage capacity, which translates to a C/8 rate for a battery used at 45% DOD. As we show, faster charging rates are not beneficial, since market-based pricing plans generally offer long low-rate periods for charging at night.

Given the constraints above, we frame GreenCharge’s linear optimization problem as follows. The objective is to minimize a home’s electricity bill using a battery array with a usable capacity (after accounting for its DOD) of C kWh. We divide each day into T discrete intervals of length I from 1 to T . We then denote the power charged to the battery from the grid during interval i as s_i , the renewable power charged to the battery as g_i , average renewable power available to the home as r_i , the power discharged from the battery as d_i , and the power consumed from the grid as p_i . We combine both the battery array and inverter inefficiency into a single inefficiency parameter e . Finally, we specify the cost per kWh over the i th interval as c_i , and the

amount billed as m_i . Formally, our objective is to minimize $\sum_{i=1}^T m_i$ each day, given the following constraints.

$$s_i \geq 0, \forall i \in [1, T] \quad (5.1)$$

$$d_i \geq 0, \forall i \in [1, T] \quad (5.2)$$

$$g_i \geq 0, \forall i \in [1, T] \quad (5.3)$$

$$g_i \leq r_i, \forall i \in [1, T] \quad (5.4)$$

$$s_i \leq C/4, \forall i \in [1, T] \quad (5.5)$$

$$g_i \leq C/4, \forall i \in [1, T] \quad (5.6)$$

$$\sum_{t=1}^i d_t \leq e * \sum_{t=1}^i s_t + e * \sum_{t=1}^i g_t, \forall i \in [1, T] \quad (5.7)$$

$$\left(\sum_{t=1}^i s_t + \sum_{t=1}^i g_t - \sum_{t=1}^i d_t/e \right) * I \leq C, \forall i \in [1, T] \quad (5.8)$$

$$m_i = (p_i + s_i - d_i) * I * c_i, \forall i \in [1, T] \quad (5.9)$$

The first second and third constraint ensure the energy charged to, or discharged from, the battery is non-negative. The fourth constraint ensures that total renewable energy charged to the battery is less than or equal to the available renewable energy. The fifth and sixth constraint limits the battery's maximum charging rate. The seventh constraint specifies that the power discharged from the battery is never greater

<i>Model</i>	<i>12am-7am</i>	<i>7am-11am</i>	<i>11am-5pm</i>	<i>5pm-7pm</i>	<i>7pm-12am</i>	Average (%)
SVM-Linear	14.77	27.32	46.72	18.49	47.03	29.5
SVM-RBF	22.44	63.77	71.93	17.84	35.01	42.51
SVM-Polynomial	4.74	4.62	6.48	7.99	5.14	5.75

Table 5.1. Average prediction error (%) over 40 day sample period for SVM with different kernel functions.

than the total power charged to the battery multiplied by the inefficiency parameter. The eighth constraint states that the energy stored in the battery array, which is the difference between the energy charged to or discharged from the battery over the previous time intervals, cannot be greater than its capacity. Finally, the ninth constraint defines the price the home pays for energy during the i th interval. The objective and constraints define a linearly constrained optimization problem that is solvable using standard linear programming techniques. GreenCharge solves the problem at the beginning of each day to determine when to switch between grid and battery power, and when to charge the battery from grid vs renewables. SmartCharge uses a similar linear programming formulation without the constraints specific to renewable energy. Since the approach uses knowledge of next-day consumption and generation patterns, we next detail techniques for predicting next-day consumption and generation, and quantify their accuracy for our case study home.

5.5 Predicting Consumption and Generation

As discussed in §5.4, solving GreenCharge’s linear optimization problem requires *a priori* knowledge of next day consumption and generation patterns. We develop a machine learning based approach to predicting demand, and use an approach developed in prior work [97] to predict next day energy harvesting based on weather forecasts. We discuss each mode in turn.

5.5.1 ML-based Demand Prediction

A simple approach to predicting consumption is to use past-predicts-future models that assume an interval’s consumption will closely match either that interval’s consumption from the previous day or the prior interval’s consumption. As we show, the approach does not work well for the multi-hour intervals in Ontario’s TOU pricing plan. Instead, we develop statistical machine learning (ML) techniques to accurately predict consumption each interval. While our techniques have numerous applications, e.g., dispatch scheduling in microgrids, we focus solely on their application to GreenCharge in this chapter.

We experimented with a variety of prediction techniques, including Exponentially Weighted Moving Averages (EWMA), Linear Regression (LR), and Support Vector Machines (SVMs) with various kernel functions, including Linear, Polynomial, and Radial Basis Function (RBF) kernels. EWMA is a classic past-predicts-future model that predicts consumption in the next interval as a weighted sum of the previous interval’s consumption and an average of all previous intervals’ consumption. More formally, EWMA predicts the energy consumption for each interval on day k as $\hat{E}_C(k + 1) = \alpha E_C(k) + (1 - \alpha)\hat{E}_C(k)$, where α is a configurable parameter that alters the weight applied to the most recent interval versus the past. Note that since each interval’s power consumption is different, we apply EWMA to each interval independently on a daily basis. As might be expected, since home consumption patterns vary largely around mealtimes, we found that predicting consumption based on the preceding interval to be highly inaccurate.

Both LR and SVM are regression techniques that combine and correlate numerous indicators (or features) of future power consumption to predict next-day usage. We experimented with a total of nine features: outdoor temperature and humidity, month, day of week, previous day power, previous interval power, as well as whether or not it is a weekend day or a holiday. We also included the EWMA prediction

as an additional feature. To predict next-day temperature and humidity, we used weather forecasts from the National Weather Service available from the National Digital Forecast Database (<http://www.nws.noaa.gov/ndfd/>). To evaluate our techniques we used power data collected every second from our case study home over a period of four months from June to September 2011. For the LR and SVM models, we used the first 70 days of the data set for model training, and the last 40 days for evaluating the model’s accuracy. We use the LibSVM library [39] to implement our LR and SVM models. Our SVM models use the *nu*-SVR regression algorithm, which we found always performed better than the ϵ -SVR algorithm [39]. For simplicity, we only predict consumption for the Ontario TOU rate periods in Figure 5.2.

Before training our model, we employed Correlation-based Feature Subset Selection (CFSS) to refine the number of input features [61]. CFSS evaluates the predictive ability of each individual feature along with the degree of redundancy between features. We apply CFSS separately for each of the five intervals, since the pattern of power consumption varies each interval. CFSS reduces the number of features in prediction model from nine to: four for 12am-7am, seven for 7am-11am, seven for 11am-5pm, six for 5pm-9pm, and five for 9pm-12am. In general, we find that more features are useful during periods with high, variable consumption.

We then experimented with multiple variations of LR models, including least squares and different regularized models (LASSO, ElasticNet, and Ridge Regression), since we found that temperature, humidity, and past data were approximately linear with respect to power consumption. However, our best performing LR model (ElasticNet) had an average error of 37%. EWMA performed much better, although Figure 5.5 demonstrates its limitations in predicting future consumption. The figure shows actual power consumption each day during the first interval (12am-7am), as well as EWMA ($\alpha = 0.35$) and the SVM-Polynomial model. EWMA is unable to predict large spikes or dips in consumption before they occur. Instead, EWMA’s

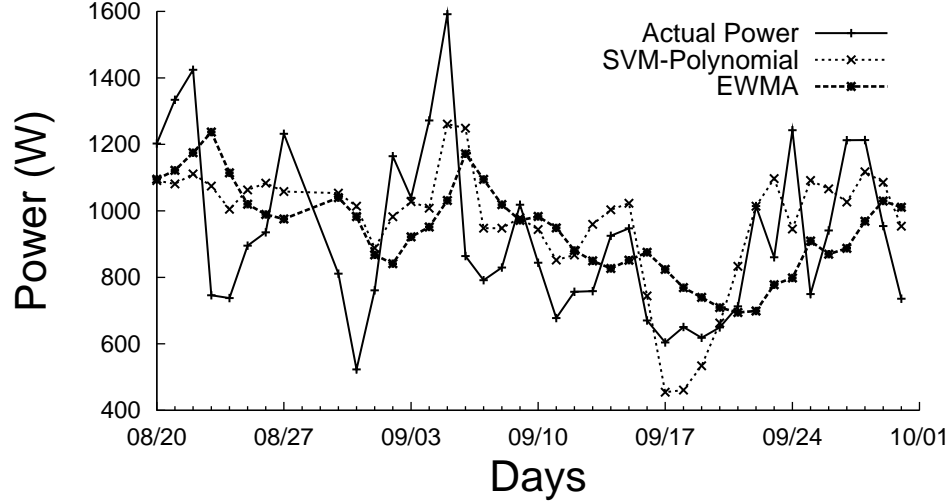


Figure 5.5. Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.

predictions never vary too far from the mean usage. In contrast to EWMA, the SVM approach is able to partially predict many of the spikes and dips in consumption. Over our 40 day testing period, we found that SVM-Polynomial had an average error of only 5.75%. The SVM model with the Linear and RBF kernel performed worse than EWMA, as Table 5.1 shows, with a 29.5% and 42.5% average error, respectively. As a result, in §5.6 we use SVM-Polynomial to evaluate SmartCharge.

5.5.2 Predicting Energy Harvesting from Weather Forecasts

For predicting the harvested solar energy we use the prediction model presented in [97]. For a given solar panel deployment this model translates the forecasted sky cover, by National Weather Service (NWS), into solar energy harvesting prediction. The NWS publishes weather forecast including sky condition forecast, every hour. The forecast contains predicted sky condition for next 24 hours. The model computes predicted solar harvesting power for every hour as:

$$Power = MaxPower * (1 - SkyCondition) \tag{5.10}$$




Figure 5.6. Average dollar savings per day for both SmartCharge and GreenCharge in our case study home.

$Power$ in (5.10) is the predicted solar harvesting power, $MaxPower$ is the maximum possible solar power that can be harvested from the given solar panel in a given hour of day assuming perfectly sunny day, and $SkyCondition$ is the fraction of sky that is covered with clouds.

5.6 Experimental Evaluation

To illustrate GreenCharge’s potential for savings, we use the home described in §5.4 to evaluate the savings using Ontario’s TOU rate plans in simulation from Figure 5.2. While our home is not located in Ontario, it lies at the same latitude and experiences a similar climate. Thus, the prices are not entirely mismatched to our home’s consumption and generation profile. In our experiments, we vary the pricing plans and battery characteristics to see how future price trends and battery technology impact savings. To predict next-day usage, we use the SVM-Polynomial model described in §5.5. Similarly, to predict next-day generation, we use the forecast-based




Figure 5.7. Average percentage savings for both SmartCharge and GreenCharge in our case study home.

model from §5.5. Finally, to quantify the optimal savings, we compare with an oracle that has perfect knowledge of next-day consumption and generation.

Unless otherwise noted, our experiments use home power data from the same 40 day period in late summer as the previous section, and generation data from our own solar panel installation scaled up to generate equal to the home’s average power consumption. We use CPLEX, a popular integer and linear programming solver, to encode and solve GreenCharge’s (and SmartCharge’s) optimization problem, given next-day prices and expected consumption levels. Note that we consider only usable storage capacity in kWh in this section, which is distinct from (and typically much less than) battery capacity. In the next section, we discuss the battery capacity necessary to attain a given storage capacity. As mentioned in §5.4, we use an energy conversion efficiency of 80% for the battery and a $C/4$ charging rate for the usable storage capacity.




Figure 5.8. SmartCharge’s and GreenCharge’s savings as a function of the charging rate for a 24kWh storage capacity.

5.6.1 Household Savings

Figure 5.6 shows the average savings per day in USD for the TOU rate plan over the 40 day period, as a function of storage capacity, while Figure 5.7 shows the savings as a percentage of the total electricity bill. The graphs show that a storage capacity beyond 30kWh does not significantly increase savings. Further, smaller storage capacities, such as 12kWh, are also capable of reducing costs, near 10% for SmartCharge and 20% for GreenCharge. If we extrapolate the savings over an entire year, we estimate that GreenCharge with 24kWh of storage is capable of saving \$200, while SmartCharge is capable of saving \$100. Finally, the graphs show that GreenCharge’s performance is close to that of an oracle with perfect knowledge of future consumption and generation: mispredictions only cost a few dollars each year with 24kWh storage capacity, or under 10% of the total savings.

The experiments above assume that we use today’s battery characteristics and price levels. Of course, a more efficient battery and inverter would increase the usable storage capacity in a battery array. As the experiments above indicate, increasing storage capacity increases the savings up to a 30kWh capacity. We evaluate the effect




Figure 5.9. Varying the average electricity price (a) and the peak-to-off-peak price ratio (b) impacts savings.

of maximum battery charging rate on home savings using TOU pricing plan over 40 day traces in presence of 24kWh battery capacity. Figure 5.8 demonstrates that the maximum charging rate has a minimal effect on savings, since the TOU rate plan offers a long period of relatively low rates during the night for charging. The charging rate need only be high enough, e.g., a C/10 rate, to charge the battery over these periods. Figures 5.9(a) and (b) show how the savings change if we vary either the average price (while keeping price ratios constant) or the peak-to-off-peak price ratio (while keeping the average price constant) for a 24kWh capacity, assuming C/4 charging rate for the usable storage capacity, for both GreenCharge and SmartCharge. The graphs demonstrate that, as expected, rising prices or ratios significantly impact the savings. In the former case, the relationship is linear, with a doubling of today's average price resulting in a doubling of the savings for both GreenCharge and SmartCharge. Thus, if average electricity prices continue to rise 5% per year, as in the past, the expected savings for both systems should also increase at 5% per year. In the latter case, while the savings rate decreases slowly as the ratio increases, the savings nearly doubles (up 88%) for both GreenCharge and SmartCharge if the current ratio increases slightly from 1.6 to 2.




Figure 5.10. Additional savings (in % and \$) from sharing 12kWh and 24 kWh between homes.

Finally, Figure 5.10 shows the additional savings homes are able to realize by sharing battery capacity with neighbors. Sharing is beneficial when homes exhibit peaks at different times by allowing them to share the available storage capacity. For the experiment, we use power data for a single day from a pool of 353 additional homes we monitor (described below), such that each point is an average of twenty runs with a set of k randomly chosen homes. We report both the additional dollar and percentage savings per home. We include 90% confidence intervals for the dollar savings. The experiment shows that sharing a battery array between homes results in additional savings as we increase the number of homes. As expected, more homes require more storage capacity to reap additional benefits. With 10 homes sharing 24kWh per home, the additional savings is 25%. However, with 12kWh per home the percentage savings does not increase beyond 15% when sharing with more than four homes.

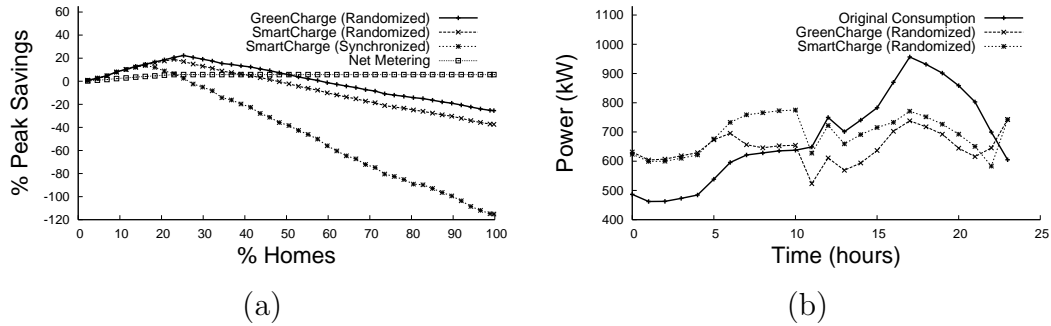


Figure 5.11. With 25% of homes using GreenCharge, the peak demand decreases by 22.5% (a) and demand flattens significantly (b).

5.6.2 Grid Peak Reduction

The purpose of market-based rate plans is to lower peak electricity usage across the entire grid. We evaluate the potential grid-scale effect of GreenCharge using power data from a large sampling of homes. We gather power data at scale from thousands of in-panel energy meters that anonymously publish their data to the web. Power consumption trace for each home is at the granularity of one hour. Since we do not know if the meters are installed in commercial, industrial, or residential buildings, we filter out sources that do not have typical household power levels and profiles, i.e., peak power less than 10kW and average power less than 3kW. We also filter out sources with large gaps in their data. After filtering, we select 435 homes from the available sources.

Figure 5.11(a) plots the peak power over all the homes as a function of the fraction of homes using GreenCharge and SmartCharge with energy storage. For these experiments we assume that each home has an available energy storage equal to half the home’s average daily consumption. Charging rate of $C/4$ for the usable storage capacity is assumed. The figure shows that GreenCharge and SmartCharge are capable of reducing peak power by roughly 20% when little more than 20% of homes use the system, as long as the homes randomize when they begin overnight charg-




Figure 5.12. Demand flattening with Net Metering.

ing. If everyone begins charging at the same time, e.g., at 12am at night, the peak reduction decreases to a maximum of only 8%. Even using randomized charging, if more than 22% of consumers install GreenCharge or SmartCharge, then the peak reduction benefits begin to decrease, due to a nighttime “rebound peak”. Once 45% of consumers use the system the evening rebound peak actually becomes larger than the original peak. The same point occurs when only 25% of homes use the system without randomized charging. ‘Net Metering’ represents those homes which have on-site renewable deployments, however, they don’t have on-site battery installations for storing this energy. Hence, the renewable energy is consumed as soon as it is generated. In contrast to GreenCharge and SmartCharge the peak savings from ‘Net Metering’ increase from 0% to 5.75% and then flattens out. The reason being, net metering does not use any on-site battery storage, it simply uses the renewable energy whenever it is available else the power is drawn from the grid. Also, as can be seen from figure 5.12 net metering effectively flattens out the mid day peaks between 11am and 2pm, however, it does poorly to shave the evening peak which occurs after 5pm. This is because solar energy harvest reduces significantly towards sunset.

Clearly, battery storage is required to shave the evening peaks. Another important observation from figure 5.12 is that net metering increases the difference between the minimum and maximum power drawn from the grid during day time, i.e., between 7am to 7pm, hence making load on the grid less predictable and sporadic.

All our experiments assume that prices do not change in response to homes installing battery-based energy storage, i.e., a large fraction of homes install the system simultaneously. A more plausible and realistic scenario is that the rate of adoption slowly rises with the differential between the peak and off-peak prices. In this scenario, the gradual load shifting would alter prices in each rate period. At some point, as Vytelingum et al.[110] formally show, the price changes would make the system increasingly less attractive for new users, as the difference between peak and off-peak prices would approach zero.

We discuss GreenCharge’s and SmartCharge’s economics at scale further in §5.7. Figure 5.11(b) shows grid power usage over time, with 0% and 22% of the homes using GreenCharge and SmartCharge with randomized charging, and demonstrates how both approaches cause demand to “flatten” significantly. Such a peak reduction would have a profound effect on generation costs, likely lowering them by more than 20% [83]. Finally, with 20% of homes using GreenCharge or SmartCharge, the increase in total energy usage is only 2%. The result demonstrates that the benefits of flattening likely outweigh the increased energy consumption due to battery/inverter inefficiencies.

5.7 Cost-Benefit Analysis

The previous section shows that GreenCharge cuts an electric bill by 20% with today’s market-based pricing plans, compared to around a 10% decrease with SmartCharge. In this section, we first discuss GreenCharge’s return on investment (ROI), including its installation and maintenance costs. We ground our discussion using price quotes,

primarily from the altE store (<http://www.altestore.com>), for widely-available commercial products.

5.7.1 Return-on-Investment

In many instances, homes already have the necessary infrastructure to implement GreenCharge. For example, many homes in developing countries already utilize UPSs because of instability in the power grid. In addition, homes with photovoltaic (PV) systems require on-site energy storage to balance an intermittent supply with demand without the aid of net metering. Batteries in electric vehicles (EVs) could also serve as energy storage. In each case, the homes already include the required infrastructure and battery capacity to implement GreenCharge. Since the homes would not need new infrastructure, the ROI is positive in these cases. Below, we discuss the ROI for homes that do not already have the necessary infrastructure.

Table 5.2 shows cost estimates for purchasing and installing GreenCharge's components. For the inverter, we assume Apollo Solar's True Sinewave Inverter, which combines an inverter, battery charger, and transfer switch into a single appliance. To read battery state and control the appliance, we attach an additional communications gateway available for the inverter. Numerous home energy meters are available: The Energy Detective (TED) is a popular choice and costs \$200. Nearly any server is adequate to support GreenCharge's software. We use an embedded DreamPlug server at a cost of \$159 as the gateway in the homes we now monitor. To hold the battery array, we assume two MNEBE-C 12-battery modular enclosures. Finally, we estimate \$200 for cabling and a day's labor at \$500 for installation. The total estimated cost, excluding batteries, is \$4871. Of course, GreenCharge's largest expenses are its battery array and solar panel installation. We discuss each below.

Sealed VRLA/AGM lead-acid batteries are the dominant battery technology for stationary home UPSs and PV installations, due to their combination of low price,

<i>Component</i>	<i>Total</i>
Inverter	\$2099.00
Battery Charger	-
Transfer Switch	-
Inverter Gateway	\$287.00
Energy Monitor	\$200.00
Server	\$159.00
Battery Enclosure	\$1426.00
Cabling	\$200.00
Labor	\$500.00
Total	\$4871.00

Table 5.2. Estimated cost breakdown for installing SmartCharge’s supporting infrastructure.

high efficiency, and low self-discharge rate. By contrast, lithium ion batteries, while lighter and more appropriate for EVs, are much more expensive. We use, as an example, the Sun Xtender PVX-2580L with a 3kWh rated capacity (at a C/20 discharge rate), which costs \$570 [105] and is designed for deep-cycle use in home PV systems. The battery’s manual specifies its lifetime as a function of its number of charge-discharge cycles and the DOD each cycle. We use the data to estimate the yearly cost of batteries—in \$/kWh of *usable* storage capacity—as a function of the depth of discharge (Figure 5.13) amortized over their lifetime, assuming GreenCharge’s typical single charge-discharge cycle per day. The usable storage capacity takes DOD into account: a battery rated for 10kWh operated at 50% depth of discharge has a usable capacity of only 5kWh. Figure 5.13 demonstrates that cost begins to increase rapidly after a 45% DOD, with an estimated cost of \$118/kWh of usable capacity.

While solar panel prices are dropping dramatically, current prices are \$7-\$9 per watt for installing solar generation. Since both the average consumption in our example home (and the average across the U.S.) is 1kW, it would cost \$4000 for a system capable of producing half the home’s electricity. Of course, since a solar installation does not produce its maximum power all the time, our home would likely need a

installation with at least a 4x larger capacity than our desired output. As a result, to generate half the home's electricity from solar panels would cost \$16,000-\$20,000.

In the U.S., GreenCharge likely qualifies for a Residential Renewable Energy Tax Credit, reducing its cost by 30%. Additionally, U.S. state and local governments offer an assortment of tax incentives for energy-efficiency improvements [44], which we estimate lower costs by 20%. Despite the advantages, today's lead-acid batteries and solar panels are still too expensive to produce a positive ROI at current electricity prices. For instance, while 24kWh of usable storage capacity saves \$91.25 per year using the Ontario TOU rate plan, batteries alone would cost \$1416 per year assuming the take breaks above. However, recent advancements in battery technology promise to dramatically reduce battery costs in the near future. Lead-carbon batteries have an expected lifetime 10x longer than today's sealed lead-acid batteries at roughly the same cost [49, 56, 93]. Figure 5.14 shows the extended lifetime using data from recent tests conducted at Sandia National Labs comparing today's sealed lead-acid battery and a new lead-carbon battery (the UltraBattery) [93]. In addition, solar panel prices per installed watt are predicted to drop to \$1 per watt over the next decade.

Lead-carbon batteries combined with modest and expected price increases (25%) and peak-to-off-peak ratios (25%), as well as a decrease in solar panel prices, would produce a positive ROI for GreenCharge in a few years. As Figure 5.11 shows, enabling only 20% of homes with GreenCharge would dramatically reduce peak demands, and, hence, generation costs for *all* homes, even those that have not invested in the system. Since all homes benefit from lower prices, utilities may consider subsidies that spread costs across all consumers, which for 20% of homes would lower costs by nearly 5X.

Alternatively, utilities might consider modifying their pricing plans to incentivize GreenCharge (and SmartCharge) in all homes by increasing the fraction of the bill based on peak usage. While many utilities charge large consumers based on their




Figure 5.13. Amortized cost per kWh as a function of depth of discharge.

peak usage over a day or month [30], residential bills typically do not include such a charge. Incorporating a substantial peak usage charge in electric bills would prevent the large rebound peaks in Figure 5.11 by directly incentivizing homes to flatten demand, rather than shift as much demand as possible to low-cost periods (causing the rebound peak). With market-based plans that only charge per-kWh, as more consumers install the system and shift their demand to low-cost periods, the price difference between the low-cost and high-cost periods would lessen to reflect the new demand distribution, thus lowering the ROI and discouraging additional homes from installing the system. A substantial peak-usage charge would maintain the financial incentives and continue to flatten demand (and prevent rebound peaks) as the fraction of GreenCharge-enabled homes approaches 100%.

A full discussion of GreenCharge’s impact on the economics of electricity generation is outside the scope of this chapter. However, it is clear that today’s market-based pricing plans assume that the price elasticity of electricity demand is low, i.e., changes in price do not have a significant impact on demand. GreenCharge fundamentally changes this fact by making demand nearly fully elastic with price.




Figure 5.14. Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [93].

5.7.2 Distributed vs. Centralized

Utilities have already begun to deploy large, centralized battery arrays to reduce peak usage and integrate more wind and solar farms, which require substantial energy storage to match an intermittent supply with variable demand. However, distributing battery storage and energy harvesting throughout the grid has a number of inherent advantages over a centralized approach. In particular, local energy storage and generation serves as backup power during extended blackouts, lessening the economic impact of power outages and promoting a more stable grid. A centralized system also introduces a single point of failure. Further, substantial home energy storage and generation may be a catalyst for implementing microgrids, where matching supply and demand is difficult without an energy buffer. Storing and generating energy at its point-of-use also reduces transmission losses by eliminating losses incurred from generator to battery array.

Finally, perhaps the most important argument for installing many distributed battery arrays and energy harvesting deployments in homes, rather than large centralized arrays, is to encourage distributed generation without relying on net metering. While

today's PV installations typically use net metering to offset costs by selling energy back to the grid, it is not a scalable long-term solution. Injecting significant quantities of power into the grid from unpredictable and intermittent renewables has the potential to destabilize the grid by making it difficult to balance supply and demand. GreenCharge provides an alternative to net metering to offset costs in home PV systems that use batteries instead of net metering.

5.8 Conclusion

In this chapter, we explore how to lower electric bills using GreenCharge by storing low-cost energy for use during high-cost periods. We show that typical savings today are near 20% per home with the potential for significant grid peak reduction (20% with our data). Finally, we analyze GreenCharge's costs, and show that recent battery advancements combined with an expected rise in electricity prices and decrease in solar panel prices may make GreenCharge's return on investment positive for the average home within the next few years.

CHAPTER 6

SCALING DISTRIBUTED ENERGY STORAGE FOR GRID PEAK REDUCTION

Although optimizing building energy footprints (such as in Chapters 4 and 5) reduces electricity bills, it does not necessarily make the aggregate grid-wide demand profile sustainable—as indicated by results in Section 4.6.2. Utilities need to shave the peak demands on their grids so as to make generation more sustainable, and optimize grid’s operational and capital costs. Hence, they are transitioning to variable pricing plans. However, even though homes can cut their bills using storage with variable pricing, energy storage adoption at scale can worsen the aggregate peak on the grid: Simultaneous battery charging across several homes during low price periods can lead to the formation of tall “rebound peaks.”

In this chapter we propose a simple solution to address the energy storage scaling problem by augmenting variable electricity pricing plans with a peak demand surcharge. We also present PeakCharge, an online peak-aware charging algorithm to optimize the use of energy storage in the presence of a peak demand surcharge.

6.1 Introduction and Motivation

As is now well-known, a significant fraction of the electric grid’s capital and operational expenses (CapEx and OpEx) result from satisfying its peak power demands. For example, recent work estimates that 10%-18% of North American CapEx, in terms of energy generation capacity, is idle and wasted over 99% of the year [54]. Similarly, peak demand also influences OpEx, by i) requiring utilities to operate high

cost and inefficient “peaking” generators to meet demand [10], ii) contributing to higher transmission charges, which are set based on peak demand, and iii) forcing utilities to offset supply shortages by purchasing electricity in the wholesale market at inopportune times, i.e., when it is most expensive. Thus, reducing peak demand and its impact on CapEx and OpEx is an important part of ongoing smart grid research efforts. One way to reduce peak demand that has received significant attention in the research community is leveraging energy storage to shift some demand from peak to off-peak periods. To shift demand, prior work proposes to store energy during off-peak periods, which increases off-peak demand, and use it during peak periods, which then decreases peak demand [37, 43, 46, 59, 72, 109, 110].

To implement the approach, utilities may either i) install large-scale *centralized* energy storage systems at strategic points in the grid, such as at power plants and substations [72], and directly control when they store and release energy, or ii) incentivize consumers to install and control their own small-scale energy storage systems *distributed* at buildings throughout the grid. Prior research has focused largely on the latter case, since the increasing adoption of variable rate pricing plans by utilities [41, 87, 101] provides an incentive [37, 43, 46, 59, 109, 110], and endowing buildings with energy storage has additional value-added benefits, e.g., providing power during outages and conditioning power to increase its quality. Since variable rate pricing plans charge higher rates during periods of peak demand, consumers that store energy during off-peak periods—when prices are low—and use it during peak periods—when prices are high—are able to lower their electricity bill. While many energy storage technologies exist, including pumped water storage, flywheels, and compressed air, batteries are currently the most viable option for storing energy at building-scale.

Prior research analyzes the potential savings for residential [37, 43, 46, 110] and industrial [59, 109] consumers to install batteries. The focus is largely on cost-benefit analyses using existing pricing plans, which vary electricity’s price per kilowatt-hour

(kWh). Unfortunately, for the reasons below, these plans *provide only a weak incentive for distributed energy storage and do not promote its adoption at large scales.*

Large Upfront Capital Costs. Since today’s pricing plans typically exhibit low prices during off-peak nighttime periods and high prices during peak daytime periods, they incentivize consumers to shift *all of their demand* to the off-peak period. Of course, the cost of batteries limits the amount of storage capacity available to shift demand. In our prior work on SmartCharge, we show that for a residential home with near the average U.S. electricity usage, $\sim 24\text{kWh}$ of capacity¹ maximizes the return-on-investment (ROI) when taking into account battery costs [78]. Given typical battery lifetimes, we estimate the *annual* amortized cost to maintain 24kWh of energy storage to be \$1416 [78]. Since the annual electricity bill for an average U.S. home is \$1419 [38], battery costs effectively prevent (at current price levels) a positive ROI using this much energy storage.

Rebound Peaks and Grid Instability. Current pricing plans incentivize *all consumers* to charge their batteries during off-peak, low-price periods. Thus, at large scales, simultaneous battery charging during off-peak periods will trigger rebound peaks if prices do not change to reflect the resulting increases in off-peak demand. Our prior work shows that if prices do not change and 100% of consumers install 24kWh of energy storage, then the peak demand period will migrate to the (previously) off-peak period and actually *increase*, rather than decrease, peak demand by nearly 120% [78]. Note that most variable rate pricing plans in use are Time-of-Use (TOU) plans with rates that do not react quickly to changes in demand, but instead are manually reset by utilities on an infrequent basis, e.g., monthly or seasonally [87].

Uncertain Return-on-Investment. One way to prevent rebound peaks is to alter electricity rates in real-time as peak and off-peak demand changes. Although

¹Operated at a maximum of 45% depth-of-discharge.

not widespread, some utilities are experimenting with real-time pricing (RTP) plans for residential consumers, where rates vary dynamically each hour based on demand [41, 101]. Unfortunately, consumers only benefit from energy storage by exploiting the difference between peak and off-peak prices. With RTP plans, as peak demand declines and off-peak demand rises due to the increasing use of energy storage, the difference between the peak and off-peak price narrows, reducing energy storage’s benefits [110]. In the extreme, if grid demand is near flat then the price of electricity will be similar at all times [37, 78, 110]. Once the peak/off-peak price differential is not large enough to compensate for the conversion losses from storing energy in batteries, there is no benefit to using energy storage. Our prior work estimates that grid demand would be near flat once just 22% of consumers install 24kWh of energy storage [78], which is consistent with related work [37, 110]. Consumers are unlikely to invest in energy storage with such uncertain future long-term benefits.

Socialized Benefits and Free Riders. For residential consumers, the annual cost to install and maintain battery-based energy storage is much higher—around 10X for average consumers in the U.S.—than the annual savings on an electric bill using current battery costs, electricity rates, and pricing plans [78]. However, prior work does not consider the grid-wide reductions in generation costs from lowering the grid’s aggregate peak demand. Unfortunately, with existing pricing plans, these cost savings are distributed (or socialized) across all consumers, since they manifest themselves as cheaper electricity rates. Thus, variable rate pricing plans provide a weak, non-optimal incentive for energy storage. Strengthening the incentive requires eliminating free riders to ensure that the consumers that invest in energy storage reap its full benefits, especially given the large capital costs.

The problems above arise from the interaction between current pricing plans and battery charging algorithms that minimize cost. We argue that solving these problems requires re-designing both pricing plans and charging algorithms to explicitly

encourage energy storage adoption. In particular, any charging algorithm should prevent grid instability regardless of the pricing plan, similar to how TCP prevents Internet congestion even though it does not maximize end-user bandwidth. Likewise, pricing plans should sustain, not eliminate, the incentive to use energy storage as capacity scales. Finally, the charging algorithm and pricing plan should work together to ensure a stable grid, while also maximizing energy storage’s ROI at scale.

6.1.1 Contributions

Ideally, energy storage distributed at buildings throughout the grid would behave like centrally-controlled energy storage of equal capacity. That is, the “right” fraction of buildings would i) charge their batteries whenever grid demand is below average and ii) discharge their batteries whenever grid demand is above average, such that aggregate grid demand remains flat and constant at the average. Of course, ensuring the behavior of any self-organizing distributed system emulates that of an equivalently-sized centralized system is challenging. In this case, determining when and how many batteries should charge requires explicit feedback from the grid and coordination among all buildings, which does not scale. This chapter targets an alternative approach: designing a charging algorithm and pricing plan where individual consumers (partially) flatten their own demand. As we discuss, our distributed approach does not require global coordination between consumers and the utility, and addresses each of the issues with scaling distributed energy storage.

The main drawback to incentivizing consumers to flatten their own demand is that it may require more aggregate energy storage capacity to flatten grid demand than the minimum required using a centralized approach. Since batteries are expensive, minimizing overall storage capacity and distributing it as widely as possible among consumers is critical to reducing per-consumer capital costs and increasing ROI. Our hypothesis is that, when consumers’ peak demand is well-aligned, a charging algo-

rithm and pricing plan that flattens each consumer’s demand uses aggregate storage capacity near the optimal centralized approach. In evaluating our hypothesis we make the following contributions.

Incentive-compatible Design. We describe the *storage adoption dilemma* that arises as energy storage scales. We show that existing charging algorithms and pricing plans cannot simultaneously minimize an electric bill and ensure grid stability at scale. In particular, preventing rebound peaks requires some (explicit or implicit) feedback from the grid to signal algorithms to rate-limit charging as demand rises. To resolve the dilemma, we propose augmenting variable rate plans with a peak demand surcharge, and then modifying charging algorithms to account for it. Our system, called PeakCharge, is a complete redesign of our SmartCharge system [78] that optimizes a consumer’s electricity costs in the presence of a peak demand surcharge.

Closed-loop Experimentation. We implement a closed-loop simulator, which replays traces of real household demand, using a representative generator dispatch stack, which specifies the cost to generate electricity as demand rises, to dynamically compute electricity rates based on demand. Our simulator is closed-loop since our charging algorithm reacts to the rates, which in-turn alters demand and then changes the rates. In contrast, prior work has evaluated energy storage using only open-loop simulations, where consumer behavior does not affect prices. Using our simulator, we experimentally verify the undesirable behavior of existing charging algorithms and pricing plans at scale.

Grid- and Consumer-scale Evaluation. We evaluate both the grid- and consumer-scale effects of PeakCharge, comparing it with prior “greedy” approaches that store as much energy as possible during low-price periods. Our analysis shows that, when compared with these systems, PeakCharge i) reduces upfront capital costs since it requires significantly less storage capacity per consumer and ii) increases ROI, since a

peak surcharge mitigates free riding and maintains energy storage’s incentive at large scales, while requiring aggregate storage capacity within 18% of optimal.

6.2 Related Work

Numerous researchers have studied the use of energy storage at homes and buildings to shift demand and cut electricity bills under emerging variable rate electricity pricing plans. Daryanian et al. [43] was the first to propose this form of energy arbitrage. This work, as well as work by van de ven et al. [46], study the problem from a theoretical standpoint, e.g., assuming certain demand distributions, without evaluating their solutions on real data. More recently, our own work on SmartCharge [78], as well as work by Carpenter et al. [37], study a similar problem in a realistic setting taking into account battery inefficiencies, stochastic demand in residential settings, and existing variable rate pricing plans in Ontario and Illinois. Both papers mention the problems with scaling distributed energy storage to many consumers, but neither i) explores the full implications of large scale adoption, including the decreasing ROI for consumers with storage as adoption scales nor ii) proposes or evaluates a solution to the problem.

While the data sets in these papers are different, they both show that $\sim 20\%$ of homes using energy storage maximizes the grid’s peak reduction. After this point, rebound peaks and simultaneous battery charging begin to reduce energy storage’s benefits, ultimately leading to a higher peak usage than without energy storage if prices do not react to demand. In earlier work, Vytelingum et al. [110] shows formally that under variable electricity rate pricing plans there is a Nash equilibrium that maximizes social welfare, e.g., cost savings, once only 38% of U.K. households use energy storage (based on a U.K. data set). Although slightly higher than the $\sim 20\%$ of homes found above, the paper’s trend is the same: beyond a certain point with existing variable rate electricity prices the benefits of consumers installing energy

storage begin to decrease. We argue that, due to the high cost of batteries, when designing incentives for distributed energy storage, the goal should be to encourage the distribution of aggregate capacity as widely as possible among consumers.

While the work above focuses on residential settings, prior work has also looked at similar problems from the perspective of industrial consumers, particularly data centers [59, 109], but has not examined the impact of storage at scale. Prior work also highlights the effect of variable rate pricing on grid stability [91, 92], showing that real-time pricing has the potential to create an unstable closed feedback loop. We show this experimentally in Figure 6.5 in the presence of large-scale energy storage. Finally, we know of no work that proposes and evaluates using a peak demand surcharge to maintain a stable grid and prevent rebound peaks by incentivizing consumers to flatten their own demand.

6.3 Overview and Approach

Our work leverages the use of battery-based energy storage systems to reduce electricity costs. We assume an intelligent battery-based energy storage system that is capable of determining when, and how much, to charge and discharge batteries based on variable electricity rates over time to minimize electricity costs. To be cost-effective, these systems must i) limit energy storage capacity due to battery costs, which, amortized over their lifetime, are currently \$100-\$200 per year per kWh of usable capacity for the VRLA/AGM lead acid variety widely used in stationary energy storage systems, and ii) account for the $\sim 20\%$ conversion loss from storing energy in batteries. Note that, since a lead-acid battery's lifetime is a function of its depth-of-discharge (DOD), a 24kWh battery operated at 50% DOD has only 12kWh of usable capacity. As in past work, we consider both the savings from batteries and their cost (20% energy loss and capital cost) when considering a system's ROI.




Figure 6.1. Prior switch-based architectures do not significantly lower an individual building’s peak demand. Figure from [78].

6.3.1 PeakCharge Architecture

Previously proposed architectures for leveraging energy storage [78] use a programmatic power transfer switch, which allows them to toggle a building’s power supply between the grid and a battery. Thus, in addition to a charging algorithm that decides when and how much to charge batteries, the system also decides when to toggle the building’s power supply between the grid and the battery, based on expectations of future prices and demand. Of course, when batteries supply power, the building’s load dictates the rate of discharge due to Kirchhoff’s laws. Although not programmatic, such switches are common in commercial standby UPS systems, which automatically switch to battery power when grid voltage falls below a preset threshold. The coarse switching architecture works well in previous systems, since they connect to the grid and charge batteries during lengthy low-price periods at night before switching to battery power during lengthy high-price periods during the day.

In contrast, we assume a system architecture that is capable of controlling a battery’s rate of discharge independent of the building’s load. For example, if a building

is consuming 1kW of power, the system is able to control the fraction of the 1kW the battery supplies, with the grid supplying the remainder. Thus, the system may choose to satisfy 1kW of demand using 500W via the battery and 500W via the grid, or using 200W via the battery and 800W via the grid. Controlling the rate of discharge is necessary for PeakCharge’s approach, which encourages buildings to flatten their demand rather than simply shift large amounts of demand from daytime to nighttime. As Figure 6.1 demonstrates, for individual buildings, the simple switching architecture does not significantly reduce (or flatten) an individual building’s peak demand. The figure (from [78]) illustrates how, due to off-peak battery charging, our prior switch-based SmartCharge system simply shifts the original peak demand to the off-peak period to minimize electricity costs.

There are two primary ways to control a battery’s rate of discharge. A simple approach is to install multiple switches capable of switching separate fractions of a building’s load between grid and battery power. For example, the system may be able to individually switch each circuit. In this case, the system controls the rate of discharge by monitoring the load on each circuit and switching some subset of circuits to the battery to achieve a specific rate of discharge. An alternative, cleaner approach depicted in Figure 6.2 is to connect the battery in parallel to the grid and use a discharge controller to programmatically limit the rate of discharge. These controllers use pulse-width modulation (PWM) to control the charge or discharge rate by connecting and disconnecting the battery at rapid frequencies. Unfortunately, controllers capable of programmatically setting the rate of discharge are not widely available, since their primary purpose today is in testing equipment [119]. However, programmatic control may become more widespread in the future, since recent work beyond our own also requires this capability [76, 117]. We assume this latter method is available to control the discharge rate in PeakCharge.




Figure 6.2. PeakCharge architecture, which includes a battery array capable of programmatically controlling the rate of discharge wired in parallel with the grid.

Finally, both our work and prior work derives from the fact that the marginal cost for a utility to generate each additional watt of power increases non-linearly as utilities activate additional generators to satisfy increasing demand. Utilities maintain a dispatch stack of generators: as grid demand rises utilities activate, or “dispatch,” additional generators in ascending order of their marginal cost. Figure 6.3 shows the demand-cost function we use to compute generation costs based on demand in our closed-loop simulator, and demonstrates the non-linear relationship between cost and demand. To derive our function, we scaled real demand-cost data from the Southeastern U.S. from a 2008 report [55] by the Federal Energy Regulatory Commission

(FERC) to match the peak demand in our traces, discussed in Section 6.6, while also ensuring a median electricity cost of 10¢/kWh, which is near the average cost of electricity in the U.S. We then fitted an exponential function to this scaled data for use in our simulations.

6.3.2 The Storage Adoption Dilemma

Figure 6.4 depicts the *storage adoption dilemma*, a variant of the classic prisoner’s dilemma, that arises from the use of distributed energy storage at large scales to minimize electricity costs in the presence of variable rate electricity prices. At the top of the figure, variable demand for power first causes the price of electricity over time to change based on the demand-cost function from Figure 6.3. Variable pricing, in turn, incentivizes consumers to adopt energy storage to reduce their costs by shifting demand to low price periods. However, as more consumers shift demand using energy storage, the difference between the grid’s peak and off-peak demand narrows resulting in a flatter grid demand profile. As a result, prices also flatten to reflect the new demand distribution. Unfortunately, flat prices eliminate the incentive to use energy storage, which causes demand to vary again and the cycle to repeat. Of course, our depiction is idealistic. In practice, the cycle may stall due to weak incentives for energy storage and completing each step would take a long time, potentially requiring significant regulatory changes and large capital investments.

The storage adoption dilemma may also cause grid instability if prices do not react fast enough to changes in demand. To demonstrate the potential for instability, we ran a simple experiment using our trace-driven closed-loop simulator to show how grid power demand could experience significant oscillations even if utilities alter prices each day based on the previous day’s demand. Day-ahead planning is common, since consumers require some pricing feedback to adjust their behavior and utilities require some advance notice to activate generators. The simulator, which we discuss




Figure 6.3. The model we use in our simulator of the marginal cost to generate electricity as demand increases. The fitted function we use is based on scaled data from a recent FERC report [55].

in Section 6.6, takes traces of demand as input; in this case, we use demand traces we collected of 194. For our experiment, 49 of the 194 homes have usable energy storage capacity that is 50% of their average daily demand, where each home uses a “greedy” charging algorithm similar to prior work [78], which minimizes electricity costs by charging as much as possible during the lowest price periods. Our simulator is closed-loop, since it computes the next day’s prices each hour using Figure 6.3’s demand-cost function and the previous day’s demand.

Figure 6.5 shows how the peak demand periods change dramatically each day. On the first day, everyone charges during the low-price period at night (12am-6am), which increases demand during that period and, hence, also increases the price of electricity during that time on the second day. As a result, on the second day the lowest-price period shifts to the morning (6am-12pm), which is the low-demand period from the previous day, and causes peak demand to shift dramatically from the nighttime 12am-6am period to the morning 6am-12pm period. Since generators require lead time to activate, utilities carefully plan generator dispatch schedules each day based on the




Figure 6.4. An idealized depiction of the cycle that causes the storage adoption dilemma.

previous day's demand. If demand were to change dramatically each day, as in this scenario, the grid would be incapable of balancing supply and demand. This simple example highlights how existing pricing plans and battery charging algorithms may cause grid instability under certain conditions. Ensuring grid stability should be a priority of any battery charging algorithm, regardless of the pricing plan.

6.3.3 An Optimal Approach

Before describing PeakCharge’s charging algorithm, we first define and consider an optimal centralized battery charging scheme. Ideally, to minimize generation costs based on the demand-cost function from Figure 6.3, the optimal approach would shift aggregate grid demand such that it was the same—equal to average demand—all the time. If we assume a centrally controlled battery array, then an optimal algorithm simply charges and discharges batteries whenever grid demand is below or above average, respectively, such that demand is always equal to the average. With this algorithm, the minimum energy capacity necessary to flatten demand is equal to the maximum capacity ever required to charge or discharge the batteries to sustain the average. Of course, a centralized system has drawbacks compared to the distributed approach, which provides value-added benefits to consumers.

As an example, Figure 6.6 depicts a grid demand profile from the first day of our trace of 194 homes, as well as the average demand for the day. In this case, the maximum capacity required to charge or discharge the battery occurs between hour 16 and 23 and equals 392kWh (equivalent to the area between the instantaneous demand and the average demand from hour 16 to 23). With the optimal approach this 392kWh of storage would reduce generation costs by 23% based on our demand-cost function in Figure 6.3. If this storage capacity were distributed evenly among all 194 homes, then each home would need only 2.02kWh of usable energy storage (or 4.5kWh of rated capacity used at 45% depth-of-discharge to maximize lifetime). This capacity is over 5X less than the 24kWh of rated capacity each home requires to maximize energy storage’s ROI based on our previous SmartCharge work [78], which uses existing variable rate pricing plans and a “greedy” charging algorithm. Since battery costs scale linearly with capacity, maintaining 5X less capacity decreases costs by 5X (from \$1416 amortized per year to maintain 24kWh to \$266 per year to maintain 4.78kWh). The example demonstrates how minimizing capacity, and




Figure 6.5. Load oscillations in our simulated microgrid, in presence of day-ahead real time pricing.

distributing it as widely as possible among consumers reduces the ROI per consumer of energy storage.

6.4 Scalable Design

The storage adoption dilemma discourages distributed energy storage from scaling to a large fraction ($> 20\%$) of consumers. Unfortunately, variable rate electricity prices incentivize consumers with energy storage to use greedy battery charging algorithms, which charge batteries as much as possible during the lowest price periods to minimize electricity costs. At large scales, the use of greedy charging algorithms results in either i) large rebound peaks (if prices do not react to changing demand), ii) grid instability (if prices react slowly to changing demand as in Figure 6.5), or iii) no benefit to the consumer (if prices react quickly to changing demand by flattening). None of these outcomes is desirable. Variable rate pricing is effective at reducing peak demand today only because electricity’s price elasticity of demand is typically low, i.e., consumers do not react strongly to changes in electricity’s price. As a result,

only a small fraction of consumer demand shifts to low price periods. In contrast, large-scale distributed energy storage makes electricity’s price completely elastic with demand, causing a large fraction of demand to shift to the lowest price period.

6.4.1 The Effect of a Peak Demand Surcharge

Properly incentivizing distributed energy storage at scale requires rethinking electricity pricing plans. Our premise is that augmenting existing variable rate pricing plans with a peak demand surcharge (or penalty) is a simple and effective way of addressing the storage adoption dilemma. A peak demand surcharge bills consumers $\$X/\text{kWh}$ based on their peak demand over an N minute interval within some billing period M . Typical values are $N = 1$ hour and $M = 1$ day; for example, in this case, the consumer in Figure 6.1 would incur an additional charge for using $\sim 5\text{kWh}$ during their peak hour of that day. Utilities already use such a peak demand surcharge for large, primarily industrial, consumers. Put simply, a large peak demand surcharge incentivizes consumers to flatten their own demand to minimize their peak, rather than simply shift as much demand to the lowest price period. Of course, if consumers flatten their own demand, then grid demand will also flatten. As we discuss below, penalizing peak usage addresses the problems mentioned earlier in the beginning of this chapter.

6.4.1.1 Benefits

First, flattening a consumer’s demand takes significantly less energy storage capacity than shifting all of it to the lowest price period. For example, Figure 6.7 shows that, while 12kWh of usable energy storage is only capable of shifting a fraction of demand to the low price period, it is more than enough to completely flatten the original demand from Figure 6.1. As a result, the approach encourages distributing aggregate storage capacity widely across consumers, requiring less storage capacity per consumer, and resulting in lower upfront capital costs and higher per-consumer




Figure 6.6. Instantaneous and average grid demand for 194 homes in our trace.

ROI. In effect, to flatten grid demand, the approach incentivizes a large number of consumers to install a small amount of energy storage (and make a small investment), rather than incentivizing a small number of consumers to install a large amount of energy storage (and make a large investment). Second, the approach prevents rebound peaks and grid instability, since consumers are (partially) flattening, rather than shifting, their demand. Third, the approach maintains the incentive to use energy storage as capacity scales, since consumers always benefit from not paying an additional peak demand surcharge, regardless of other consumers' behavior. Finally, utilities can mitigate free riding by altering the peak demand surcharge, in addition to the electricity rate, as generation costs change, since only the set of consumers with energy storage are able to automatically optimize for peak demand. Thus, a higher peak demand surcharge and lower rates will penalize consumers with energy storage less than consumers without it.

Of course, utilities must use a peak demand surcharge in conjunction with existing variable rate schemes, since only charging based on peak demand would encourage *more* energy use. For example, with only a peak demand surcharge, if a residential




Figure 6.7. While 12kWh of energy storage is capable of shifting only a fraction of demand to the low price period, it is more than enough to completely flatten the demand from Figure 6.1.

consumer runs a dryer that causes their peak load to be 7kW over an hour, the consumer has *no incentive to ever reduce their demand below 7kW*. Since the home in Figures 6.1 and 6.7 has an average electricity usage of $\sim 1\text{kW}$, billing solely based on peak demand would allow the consumer to use 7X more electricity at no extra cost. Thus, utilities must balance the size of the peak demand surcharge with the electricity rates to encourage flattening without incentivizing consumers to use significantly more electricity. We examine how the size of the peak demand surcharge affects our battery charging algorithm in Section 6.6.

While this chapter focuses primarily on how a peak demand surcharge addresses the storage adoption dilemma, it also has other benefits. For instance, homes without energy storage could reduce their electricity costs using automated load scheduling techniques that flatten demand, e.g., via SmartCap [34] or nPlug [58]. Consumers have little monetary incentive to use these techniques today, since most deferrable loads, e.g., refrigerators, air conditioners, heaters, dehumidifiers, are unable to defer their usage (by up to 12 hours) to low-price nighttime periods without causing signifi-

cant harm, e.g., spoiled food or an uncomfortable environment. In addition, as recent work shows, flattening demand using a battery preserves privacy [76, 117], since it removes power variations that Non-Intrusive Load Monitoring (NILM) algorithms use to identify appliance usage and behavioral patterns. Unfortunately, with existing variable rate plans consumers with a battery must choose to either use it to reduce their electricity bill or preserve privacy, but not both. A peak demand surcharge could enable consumers to minimize their electricity bill *and* preserve their privacy.

6.4.1.2 Drawback

The primary drawback to encouraging consumers to flatten their own demand is that, in aggregate, it may require consumers to install more energy storage capacity than necessary to flatten grid demand. To understand why, consider a simple grid with only two homes, where each day the first home uses 1kW from 12am-12pm and 2kW from 12pm-12am, while the second home uses 2kW from 12am-12pm and 1kW from 12pm-12am. In this case, to flatten their own demand, each home requires 6kWh of energy storage for a total of 12kWh of capacity, which the homes would charge at a rate of 500W/hour when usage is 1kW and discharge at a rate of 500W/hour when usage is 2kW. However, in aggregate, the two homes' demand is already flat—using exactly 3kW all the time—without any energy storage. Thus, in this case, energy storage is not necessary. The waste occurs because the peak periods of the two homes are not aligned with each other; if their peak periods were exactly aligned then they would each require 6kWh to flatten demand (and 12kWh would be the minimum capacity necessary to flatten aggregate demand). In general, the aggregate energy storage capacity necessary to flatten grid demand by flattening each home's demand will diverge more from the optimal amount the more the peak and off-peak periods of the homes become less aligned. Of course, in practice, homes in the grid exhibit peak demand at similar times, which naturally reduces the divergence from optimal.

We quantify this divergence using our closed-loop simulator and demand traces in Section 6.6.

As we discuss, augmenting existing variable rate plans with a peak demand surcharge requires rethinking the greedy charging algorithms used in prior work. Below, we present PeakCharge’s peak-aware charging algorithm, which minimizes a consumer’s electricity bill in the presence of a peak-demand surcharge.

6.5 Peak-aware Charging

Our initial approach to designing PeakCharge’s battery charging algorithm was to simply modify the algorithm from our prior work on SmartCharge [78]. SmartCharge uses a linear program (LP) that executes at the beginning of each day and takes as input next-day electricity prices, which are typically well-known, and expected demand each hour to determine how to charge and discharge a battery throughout the day. The LP is optimal, i.e., minimizes costs, if future demand is known. For completeness, we include a description of this LP and its constraints in the Appendix. Since future demand is not known, we use machine learning, specifically a type of support vector machine, to predict next-day demand over the five multi-hour periods each day in Ontario’s TOU pricing plan [87]. On average, our predictions were within 6% of demand for a representative home, and SmartCharge’s LP achieved cost savings within 10-15% of an oracle with perfect future knowledge. As a result, for PeakCharge, we initially added constraints to SmartCharge’s LP to account for the new peak demand surcharge. As with SmartCharge, given perfect future knowledge, the LP is optimal at minimizing costs. Unfortunately, our experimental results were far from optimal: the PeakCharge variant *did not* result in flatter consumer demand and *did not* minimize electricity costs. In this case, rates were based on the Ontario TOU scheme, and the peak demand surcharge was applied to the peak hour of each day.

We found the reason for the poor results to be that, with a peak-demand surcharge, the LP is highly sensitive to the prediction of the peak demand hour each day. Unfortunately, predicting next-day demand at the granularity of an hour is much less accurate than over the multi-hour periods used in SmartCharge. Further, ensuring high prediction accuracy for the peak hour is more difficult than ensuring a high average accuracy. In contrast, SmartCharge, which only optimized for variable electricity prices, is much less sensitive to prediction accuracy. While there are corner cases the LP is able to optimize for with accurate predictions, in general, it will always charge a battery as much as possible during the lowest-rate nighttime periods. Thus, simply charging the battery at its maximum rate overnight, regardless of the predictions of next-day demand, accounts for the vast majority of the savings in SmartCharge and other systems. One implication of a sensitivity to demand prediction accuracy is that optimizing for a peak demand surcharge becomes more difficult the longer the time interval M the peak is evaluated over, since predictions are generally less accurate over longer time horizons. For instance, many utilities charge industrial consumers a surcharge for their peak demand hour within an entire month, requiring them to accurately predict demand (including the peak) each hour of the month to determine when to charge and discharge the battery using our LP above.

6.5.1 Optimizing for the Peak

Based on our experiences with the LP, rather than retrofit a greedy algorithm originally designed to minimize costs for variable electricity rates, to account for a peak demand surcharge, we instead began by designing an algorithm to minimize costs for a peak demand surcharge in the absence of variable electricity rates. Our starting point is the same algorithm we outlined in Section 6.3.3 for flattening grid demand, but applied to individual buildings. The system selects a target average

power and then simply charges and discharges batteries whenever demand is below or above the target, respectively, such that demand is always equal to the average. Note that rather than run our algorithm once per day (at the beginning of the day) using predictions of next-day demand, as with SmartCharge, this algorithm naturally operates in an online manner, adjusting the charging and discharging of the battery in real time based on changing demand. This peak-centric algorithm works well as long as i) the target average is near the actual average power, and ii) the storage capacity is large enough to flatten demand.

If the target average is too small, then the approach will not store enough energy to reduce the peak by its maximum amount; if the target is too large, then it will store more energy than necessary throughout the day. However, importantly, while the algorithm is sensitive to a prediction of average power, it does not require shorter time-scale predictions of future demand, e.g., hourly day-ahead predictions, as in the LP approach. Average power predictions over long time periods tend to be much more accurate than demand predictions over short time-scales far into the future. In fact, when predicting average power, the longer the time-scale, generally the more accurate the prediction [97], e.g., the average power of a home each year tends to vary less than each day. In addition, accurate predictions of average power over long periods, e.g., a day or month, do not require sophisticated methods [97, 98]. In this chapter, to predict average demand over an interval, we simply use the average demand over the previous interval.

If the available storage capacity is too small, then the approach may discharge batteries when demand is only slightly above average, causing there to be little energy left for the highest peaks. In this case, short time-scale predictions of future demand are necessary to optimize use of the available storage capacity, i.e., save stored energy for the highest peaks each day or month. Thus, with a peak demand surcharge, the less storage capacity a consumer has, the more fine-grained and accurate the predic-

tions required to minimize cost. However, importantly, as discussed in Section 6.4.1, flattening consumer demand requires much less energy storage capacity than shifting demand to take advantage of variable rates.

6.5.2 Optimizing for Peaks and Variable Rates

Our peak-centric algorithm focuses only on flattening demand. As a result, it minimizes a consumer’s electricity costs when their bill is based solely on a peak demand surcharge in absence of variable rates charged per kWh of energy use. Given our basic algorithm, we must modify it to optimize for cost in the presence of both a peak demand surcharge and variable rates. With a high peak demand surcharge the algorithm should behave like the peak-centric algorithm, and with a low peak demand surcharge the algorithm should behave greedily, i.e., by charging at its maximum rate during low-price periods. To understand the decision of whether to behave greedily or peak-centric, consider the inequality below, which compares the benefit of greedily taking advantage of variable rates versus the cost of raising peak demand. In this case, we consider only two rate periods: high and low, where C_{high} is the cost per kWh during the high rate period and C_{low} is the cost per kWh during the low rate period. In addition, T is the length of the low-price period, P is the cost per kWh of usage during the peak hour each day, e_{loss} is the energy conversion loss as a percentage stored energy (typically 80% in practice), and X_{max} is the maximum charging rate of the battery.

$$X_{max}e_{loss}C_{high}T - X_{max}C_{low}T > X_{max}P \quad (6.1)$$

The left side of the inequality is the maximum monetary benefit of greedily charging the battery at its maximum rate during the low-price period and then discharging it during the high-price period, while the right side is the cost of the peak demand surcharge from charging the battery at its maximum rate. If the inequality holds then the benefit of charging greedily during the lowest-price period is greater than the cost,

signaling that a consumer should act greedily. If not, then a consumer may benefit from acting peak-centric by charging (or discharging) at less than the maximum rate during low-price periods. Unfortunately, determining exactly how much less to charge (or discharge) than the maximum is challenging, requiring the same accurate short time-scale, e.g., hourly, predictions of future demand that SmartCharge’s LP requires. As a result, we adopt a heuristic approach using four simple cases, as outlined below, based on whether the electricity rate is high or low and the demand is above or below average.

- If the electricity rate is **low** and demand is **below average**, then greedily charge at the maximum rate if (6.1) holds, else charge at a rate to sustain the target average demand.
- If the electricity rate is **low** and demand is **above average**, then greedily charge at the maximum rate if (6.1) holds, else discharge at a rate to sustain the target average demand.
- If the electricity rate is **high** and demand is **below average**, then greedily discharge at the full rate (bounded by the building’s demand) if (6.1) holds, else do nothing.
- If the electricity rate is **high** and demand is **above average**, then greedily discharge at the full rate (bounded by the building’s demand) if (6.1) holds, else discharge at a rate to sustain the target average demand.

Rather than add more cases, to extend the approach to multiple rate periods, we simply divide each period into two bins, based on whether its price is higher and lower than average, and compute C_{high} and C_{low} by taking the average of the cost per period (weighted by the length of the period) in each respective bin. Based on the cases, if the inequality holds then the algorithm simply acts greedily by charging at the maximum rate when the electricity price is low, while discharging at the maximum rate (bounded


by the building’s demand) when the electricity price is high. In contrast, if the inequality does not hold, then the algorithm simply toggles to using the peak-centric algorithm, with one exception. If the electricity rate is high and demand is below average, it balances the objective of the greedy algorithm, i.e., to discharge, and the peak-centric algorithm, i.e., to charge, by doing nothing. Note that, in the extreme, since variable rates are based on the grid’s demand, as grid demand flattens the rates will equal each other and the algorithm will become entirely peak-centric.

Using our peak-aware algorithm above, when the peak demand surcharge is high relative to the electricity rates, the algorithm above charges and discharges the battery to hit the expected average demand; in contrast, when it is low, the algorithm devolves to a greedily charges the battery at the maximum rate during the lowest price periods.


6.5.3 Summary

Our *peak-aware* algorithm above optimizes for a peak demand surcharge by using inequality (6.1) to determine when to act greedily and when to optimize for the peak. In the next section, we compare the peak-aware algorithm with an online *greedy* algorithm that is conceptually similar to our previous LP-based approach, but operates in an online manner by charging at the maximum rate during the lowest-price periods at night and discharging during the highest price periods during the day. Since battery capacity is typically much lower than each day’s energy usage, this simple variant performs similarly to our previous LP-based approach.


As an additional point of comparison, we also experiment with a variant of the greedy algorithm with an additional congestion parameter, which limits the maximum charging rate of the battery by a factor P_{limit} , which is between 0 and 1. Enforcing a limit on the battery charging rate is a simple way to ensure grid stability and prevent rebound peaks, even using greedy charging. Of course, in practice, this parameter requires feedback and enforcement from a utility, which could either directly dissem-



(a)



(b)



(c)

Figure 6.8. Generation cost savings compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans. Zoom-in of generation cost savings for peak-aware algorithm (c).

inate P_{limit} to consumers or allow them to indirectly infer it, e.g., by using subtle changes in line voltage as a signal of grid demand as in nPlug [58]. We show that while our *congestion-aware greedy variant* prevents rebound peaks and grid instability, without a peak demand surcharge, it reduces the savings (and ROI) of energy storage for consumers.




Figure 6.9. Peak reduction as a percentage compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans.




Figure 6.10. Time series of aggregate grid demand for TOU (a) and DART (b) pricing for both without energy storage and using our peak-aware algorithm with each home having energy storage.

6.6 Evaluation

To evaluate the charging algorithms from the previous section, we built a closed-loop simulator that takes as input traces of building energy usage. The simulator is closed-loop, since it determines the price of electricity each hour of each day based on the demand from i) the same hour on the previous day and ii) the demand-cost function in Figure 6.3; we call this Day-Ahead Real-Time (or DART) pricing, since each day’s prices are known at the beginning of the day. In addition to DART, our simulator also supports open-loop TOU pricing, where prices do not change based on




Figure 6.11. Generation cost savings (a) and grid peak reduction (b) as we vary the size of each home’s energy storage capacity.

demand. As in our prior work [78], we use TOU prices based on Ontario’s rates [87]; specifically, 6.3¢ per kWh from 11pm to 6am (off-peak period), 11.8¢ per kWh between 6am to 10am and 4pm to 11pm (peak periods), and 9.9¢ per kWh from 10am to 4pm (mid-peak period).

In addition to the rate plans above, the simulator also supports a peak demand surcharge in \$/kW of peak usage. The surcharge applies to the highest average demand over a 30 minute sliding window across each day. Our default surcharge in the experiments below, unless otherwise noted, is \$3/kW. This surcharge is high relative to the rates, i.e., inequality (6.1) does not hold, although with DART pricing rates may rise (since they vary every day based on demand). As in our prior work [78], we use a maximum charge rate of $C/4$ for the usable storage capacity, i.e., the battery charges to full capacity in 4 hours, which translates to a $C/8$ rate for a battery used at 45% DOD. We use power demand traces from 194 homes, which have an in-panel energy meter to record usage each minute, for ten consecutive days. While our traces are not at utility scale, i.e., with tens of thousands of residential homes as well as commercial and industrial buildings, they are sufficient to verify the trends in using energy storage at scale and to explore the behavior of our algorithm. However, the

benefits of storage at scale will certainly vary based on the characteristics of each grid’s (and building’s) demand profile.

As with our previous work [33], we plan to make our traces available for download from our Smart* data repository located at <http://smart.cs.umass.edu>. Finally, we experiment with the algorithms from the previous section—greedy, peak-aware, and congestion-aware greedy—using DART (closed-loop) and TOU (open-loop) rate plans, examining both the grid-scale and consumer-scale effects.

6.6.1 Grid-scale Effects

We first examine the effect of rebound peaks when consumers use energy storage at large scales. Figure 6.8 shows the savings in generation costs across the entire grid compared to no energy storage, as the percentage of homes using energy storage scales up with both DART (a) and TOU (b) pricing. In this case, each home has usable energy storage capacity that is 50% of their average daily demand. The graph shows that as the number of homes using energy storage scales up, the greedy algorithm, akin to our SmartCharge algorithm, increases generation costs, i.e., the savings are negative, due to simultaneous battery charging and large rebound peaks after 20% of homes use energy storage (with DART).

TOU pricing scales slightly better than DART because the closed-loop pricing results in very low rates when consumers are not charging, which reflect in the next day’s prices. Since TOU rates do not change, it does not suffer from oscillations in peak demand or prices. The congestion-aware greedy variant (with a limit on the charging rate of 60% the maximum rate) in both cases is more scalable, but still results in rebound peaks once enough homes adopt energy storage (35% for DART and 45% for TOU). In contrast, the peak-aware algorithm steadily decreases the grid’s generation costs as more homes use energy storage, signaling that homes and the grid are successfully flattening demand. The congestion-aware greedy variant

demonstrates an important point: consumers could prevent rebound peaks by rate-limiting their charging, but they would reduce their cost savings.

Figure 6.8(c), which zooms in on the results for the peak-aware algorithm, demonstrates the steadily increasing, rather than decreasing, cost savings as more homes use energy storage. Notice maximum savings from the three variants is similar: this reflects that in each case the aggregate energy storage is sufficient to flatten demand. The difference with the peak-aware algorithm is that it distributes this capacity across 100% of consumers, while the other algorithms distribute it across a much smaller set of consumers. Figure 6.9 shows the corresponding reduction in peak demand for the same experiment, which, as expected, shows similar trends as the generation cost savings. Namely, the grid’s peak demand steadily decreases, rather than increases, as more homes use energy storage with the peak-aware algorithm. As an example of this decrease in peak demand, Figure 6.10(a) and (b) show the time-series of power usage of an example day in our trace both with and without energy storage (for DART and TOU pricing).

We also run a similar experiment, but rather than vary the percentage of homes using energy storage we vary the amount of energy storage each home has as a fraction of its demand. In this case, 100% of homes have energy storage. Figures 6.11(a) and (b) show the generation cost savings and peak reduction, respectively. In both cases, the results show that each home only needs energy storage capacity that is a small fraction of its average demand. In Figure 6.11(a), generation cost savings stop increasing once homes have energy storage capacity that is 20% of their average demand. Similarly, Figure 6.11(b) shows the grid’s peak not decreasing further at the same 20% threshold. At the 20% threshold, the aggregate storage capacity across homes is near the optimal storage capacity required to flatten demand (from Section 6.3.3). In contrast, using our previous work on SmartCharge without a peak

demand surcharge, a representative home required 50% of their average demand in storage capacity to maximize their ROI [78].

The previous experiments used the same, relatively high, peak demand surcharge of \$3/kW, such that inequality (6.1) does not hold and our peak-aware algorithm focuses on flattening demand. Figure 6.12 demonstrates the percentage peak reduction across the 194 homes in our traces as we vary the peak demand surcharge for both TOU and DART pricing plans. In this case, all homes use energy storage with usable capacity that is 50% of their average demand. As expected, for low surcharge values the homes are greedy resulting in large rebound peaks—larger than the original peaks—from simultaneous battery charging during low-price periods. However, once the peak demand surcharge passes the threshold defined by inequality (6.1) the homes switch to flattening their demand. In this case, for high peak demand surcharges the algorithm reduces the peak 10-15%. The threshold near \$0.60/kW represents the tipping point where the benefit of charging the battery at its maximum rate during the lowest-price hour is not worth the cost of the peak demand surcharge.

Lastly, we revisit the analysis of the optimal minimum amount of energy storage capacity necessary to flatten grid demand in a centrally controlled system (from Section 6.3.3). For that trace, the minimum capacity required was 393kWh. For the same trace, our peak-aware algorithm requires 481kWh of aggregate capacity across the grid to maximize its peak reduction and generation savings from flattening demand. As Figure 6.10 demonstrates, the peak-aware algorithm is not able to perfectly flatten demand due to inaccuracy in choosing the target average, which may result in high peaks if a home’s battery is empty during a period of high demand. While predicting average demand over a long period is more accurate than predicting hourly demand over the same period, it is not perfect due to changing consumer behavior. Despite the inaccuracy, the approach comes within 18% of the optimal centralized approach.

6.6.2 Consumer-scale Effects

With distributed energy storage in use large scales, the peak reductions and cost savings for individual consumers mirror the reductions and savings in the grid. To show this, we use the TOU pricing scheme, which, as shown above, performs similarly to DART. We take a representative home with near the average demand for a home in our trace and look at its individual peak reduction and cost savings (as both a percentage and in dollars) as function of the home’s usable energy storage capacity for our peak-aware algorithm. In this case, 100% of homes have energy storage. Figure 6.13 shows the results. As in the grid, the peak-aware algorithm reaches its maximum peak reduction (Figure 6.13(a)) when usable storage capacity is only 20% of average demand (rather than 50% with a greedy approach like SmartCharge [78]). Further, for less than 2X the energy storage capacity, the percentage cost savings in the electric bill (Figure 6.13(b)) is in the same 10-15% range as we found in SmartCharge. Thus, the ROI for the consumer is much higher, since consumers achieve similar savings using much less energy storage capacity, which overwhelmingly dominates the cost of the system. Finally, Figure 6.13 shows the average per day dollar savings, which mirrors the trend in the overall percentage savings.

The results above clearly demonstrate the benefits, as the use of energy storage scales, of incentivizing consumers to flatten their own demand using a peak demand surcharge. For the similar cost savings per consumer using energy storage, the approach requires much less storage capacity, resulting in a higher ROI. Further, this ROI does not diminish as more consumers use energy storage, since the utility is able to control each consumer’s incentivize to flatten demand using the peak demand surcharge independent of the electricity rates. This also mitigates free riding, since consumers without energy storage are less able to control their peak demand and benefit from reducing their peak. By incentivizing all consumers to use energy storage, the approach encourages distributing the aggregate energy storage necessary to





Figure 6.12. Increasing the peak demand surcharge prevents rebound peaks in the grid by incentivizing consumers to flatten their demand.


flatten grid demand across a wide set of consumers, which all share in the savings. In prior work, with only variable electricity rates, the consumers not using energy storage also benefit from lower overall generation costs (and electricity rates) as storage capacity increases, which, in turn, diminishes the savings for the consumers that use energy storage.

6.7 Conclusion


This chapter examines the effects of using energy storage distributed at buildings and homes throughout the grid to flatten grid demand. In particular, we show that as more consumers adopt energy storage, a number of problems arise that impact grid stability and generation costs. As a result, we propose to augment traditional variable rate electricity pricing plans with a substantial peak demand surcharge, which incentivizes consumers to flatten their demand rather than shift it all to a low-price period. Utilities already use peak demand surcharges for large industrial consumers;



(a)



(b)



(c)

Figure 6.13. Percentage peak reduction (a), percentage cost savings (b), and dollar cost savings (c) for an individual home using our peak-aware algorithm as the home’s energy storage capacity varies using our peak-aware algorithm under a peak-demand surcharge.

we argue that, to incentivize distributed energy storage, they may want to broaden their use to other consumers.

We then design PeakCharge, which includes an online algorithm to minimize electricity costs in the presence of variable rates and the peak demand surcharge. Using a closed-loop simulator, we show that our algorithm is effective at both i) maintaining the incentives for consumers to use energy storage at large scales and ii) ensuring grid stability. Further, our results indicate the aggregate energy storage capacity to flatten

<i>Parameter</i>	<i>Definition</i>
T	Time in T discrete intervals 1 to T
I	Length of each T
s_i	Power charged in interval i
d_i	Power discharged in interval i
e	Battery efficiency, ≤ 1
p_i	Grid power demand in interval i
c_i	Power cost per kWh in interval i
m_i	Charge for electricity in interval i
C	Battery capacity in kWh
l_i	Aggregate grid demand in interval i .
L	Peak grid demand over all intervals
c'	Peak demand surcharge

Table 6.1. Parameter definitions for linear program.

grid demand by incentivizing consumers to flatten their own demand is within 18% of the minimum, optimal capacity to flatten grid demand in a centralized system. Since flattening a home’s demand requires over 2X less storage capacity per consumer to maximize consumer savings, it significantly reduces the ROI of energy storage, which is dominated by battery costs.

6.A Appendix

Below is the modification of the LP from SmartCharge to minimize an electricity bill using energy storage in the presence of a peak demand surcharge, given future knowledge of next-day prices and next-day demand each hour. Table 6.1 defines the optimization’s parameters. The formal objective is to minimize $\sum_{i=1}^T m_i$ each day, given constraints below. The first five constraints are present in SmartCharge’s original LP: constraints (1) and (2) ensure positive energy is charged and discharged from the battery, constraint (3) bounds the battery’s charging rate, constraint (4) preserves conservation of energy (including energy conversion efficiency), and constraint (5) bounds the battery’s capacity. The final three constraints (in bold) are

necessary to optimize for a peak demand surcharge: constraint (6) computes the bill based on variable rate prices and the peak demand surcharge, constraint (7) represents grid's demand in the i th interval, and constraint (8) is the size of the peak demand surcharge.

$$s_i \geq 0, \forall i \in [1, T] \quad (6.2)$$

$$s_i \leq C/4, \forall i \in [1, T] \quad (6.3)$$

$$d_i \geq 0, \forall i \in [1, T] \quad (6.4)$$

$$\sum_{t=0}^i d_t \leq e * \sum_{t=0}^i s_t, \forall i \in [1, T] \quad (6.5)$$

$$\left(\sum_{t=0}^i s_t - \sum_{t=0}^i d_t/e \right) * I \leq C, \forall i \in [1, T] \quad (6.6)$$

$$\mathbf{m}_i = (\mathbf{p}_i + \mathbf{s}_i - \mathbf{d}_i) * \mathbf{I} * \mathbf{c}_i + \mathbf{L} * \mathbf{c}', \forall i \in [1, T] \quad (6.7)$$

$$\mathbf{l}_i = \mathbf{p}_i + \mathbf{s}_i - \mathbf{d}_i, \forall i \in [1, T] \quad (6.8)$$

$$\mathbf{l}_i \leq \mathbf{L}, \forall i \in [1, T] \quad (6.9)$$

CHAPTER 7

INTEGRATING ENERGY STORAGE IN ELECTRICITY DISTRIBUTION NETWORKS

As we have seen in the previous chapters, energy storage can effectively optimize the grid’s aggregate consumption profile, and thereby cut the electricity generation, transmission and distribution costs (by shaving the peaks). However, so far we have only considered employing lead-acid batteries at customer premises. With several competing energy storage technologies in market, employing hybrid combinations of storage technologies may provide more cost effective and profitable solutions. Besides, energy storage can possibly be placed at multiple levels of the grid hierarchy.

In this chapter, we examine the efficacy of employing different combinations of storage technologies at different levels of the grids distribution hierarchy. We present an optimization framework for modeling the primary characteristics that dictate the lifetime cost of many prominent energy storage technologies. Our framework captures the important tradeoffs in placing different technologies at different levels of the distribution hierarchy with the goal of minimizing a utility’s operating costs.

7.1 Introduction and Motivation

Nearly 40% of energy in the U.S. is consumed in the form of electricity [73]. Increasing the percentage of electrical energy is an important part of creating a clean and sustainable energy supply, as “green” energy, e.g., from solar and wind, is generally consumed in the form of electricity. In addition, transmitting and distributing electricity is significantly more efficient than transmitting and distributing other captive energy sources, e.g., via oil and gas pipelines or trucks. However, electricity

transmission and distribution (T&D) costs are non-trivial, and, in some cases, such as New York and southern California, now dominate generation costs [89]. The cost and carbon footprint to generate electricity is a complex function of the electricity demand patterns, mix of generators and fuel sources, penetration of renewable energy, and T&D efficiency.

A significant fraction of these costs are determined by the electric grid's peak power demand. The peak demand influences capital costs by dictating the capacity (and number) of transmission lines, substations, transformers, etc., since utilities must size these to service the peak. In addition, since the "peaking" generators utilities activate to satisfy demand peaks are significantly less efficient and more expensive to operate than baseload generators that are continuously active, peak power demands also influence operational costs. Thus, satisfying even brief peak demand periods has a disproportionate affect on capital and operational expenses. For example, recent estimates attribute as much as 20% of the grid's generation costs in the U.S. to servicing only the top 100 hours of peak demand each year [83]. Finally, since energy lost in transmission and distribution is a function of the square of current, rising peak demand results in quadratically higher transmission losses.

The importance of reducing peak demand is one of the primary motivations for Demand Response (DR) programs, which attempt to coerce consumers into actively shifting their load from peak to off-peak periods. Since requiring consumers to actively change their behavior to shift load is often not effective, recent work has explored the use of energy storage to automatically shift load in the background, i.e., by storing energy during off-peak periods and using it during peak periods [37, 78, 79]. Prior work in this area has generally examined deploying energy storage devices (ESDs) in individual homes, where the approach can potentially reduce a consumer's electricity bill if electricity prices vary over time, e.g., such that peak prices are higher than

off-peak prices. In fact, such energy arbitrage is an explicit use-case cited by Tesla for its new PowerWall battery, which is designed for deployment in homes [25].

While prior research, and now commercial products, target energy storage for homes, such storage can be deployed at any level of the grid’s hierarchy from the lowest level (at homes) to the medium level (at distribution transformers) to the top level (at distribution and bulk power substations). The choice of where to deploy energy storage presents interesting tradeoffs. For example, using energy storage in individual homes to reduce the home’s peak demand requires more aggregate storage capacity than employing storage at a higher level of the grid hierarchy, since each home’s peak demand does not occur at the same time yielding some smoothing from statistical multiplexing at higher levels. Since prior research largely focuses on deploying energy storage in homes, it also generally focuses on only a single type of storage technology: in particular, batteries [37, 79]. However, while batteries are the only small-scale energy storage appropriate for homes at current price points, other ESDs become more feasible at higher levels of the grid hierarchy. Energy storage technologies differ in their cost, lifetime, energy-efficiency, etc. For example, flywheels exhibit a high energy-efficiency and lifetime, but have a high self-discharge rates and cost, while lead-acid batteries exhibit a low self-discharge rate and cost, but have a shorter lifetime and lower energy-efficiency.

Thus, our hypothesis is that intelligently employing hybrid combinations of different energy storage technologies at multiple levels of the grid’s hierarchy has the potential to reduce costs relative to deploying only a single storage technology at a single level of the hierarchy. In evaluating our hypothesis, we make the following contributions.

- **ESD and Grid Modeling.** We extensively model important ESD operational characteristics, including energy density, self-discharge rate, cycle lifetime, power ramp time, etc., to capture their tradeoffs. We examine the deployment of dif-

ferent ESDs using a simple model of the grid’s electricity distribution hierarchy, which includes the various costs associated with generating, transmitting, and distributing electricity.

- **Optimization Framework.** Using our above models, we develop an optimization framework that enables us to examine the benefit of using different combinations of ESDs at different levels of the hierarchy. The goal of our optimization framework is to minimize generation costs, including the capital, operational, and storage costs, for different configurations of ESDs.
- **Implementation and Evaluation.** We implement our optimization framework and then use it to evaluate in simulation the cost and benefit of different storage configurations using smart meter data from 5000 customers of a local utility. In doing so, we identify key insights into the benefits of different ESD technologies at different levels of the grid. We find that deploying hybrid ESDs at an individual level typically improves savings over any single-technology ESD deployment, while deploying multi-level hybrid ESDs typically provides the best savings. Overall, we find that ESDs can reduce distribution-related capital and operational costs by up to 12%.

7.2 Related Work

Much of the work on DR in the grid using ESDs has focused on cutting costs for customers with storage in presence of variable electricity pricing. For example, [43] presents an optimization approach to cut costs using ESDs in presence of spot electricity prices. Similarly, in [78] authors propose the use of energy storage in homes to cut their electricity bills under a variable prices [46], which they model as a Markov Decision Process. However, none of these approaches specifically look at cutting the costs for the utility. In fact, as noted in [37], such approaches can increase the peak demand on the grid and thereby increase its op-ex and cap-ex.

In contrast to the work above, the authors in [79] propose the use of ESDs for cutting peak demand on the grid and reducing its generation costs. Prior work has also proposed renewable energy integration to reduce consumption from the grid, e.g., [106, 118]. However, all of these consider ESD/renewable deployment only at customer premises (homes), and they evaluate their solutions only for a specific ESD technology. Finally, there has been considerable work in cost-aware provisioning and DR for datacenters, e.g., [88, 111]. The closest to our work is that done by Wang et al. [111]; here, the authors present a framework for modeling different ESDs, and the tradeoffs of placing them at different levels of datacenter power hierarchy. The authors evaluated the proposed framework using traces from real datacenters. As opposed to this, we have formulated and evaluated the solution for an electricity distribution network. We model several distribution network features which are absent in a datacenter, e.g., power losses in distribution.

7.3 Background

7.3.1 Electric Grid

The electric grid is an interconnected network for delivering electricity from suppliers to consumers. Electricity is generated at power plants, often far from population centers, using different types of generators and fuels with different operational characteristics. Generated electricity exits the power plant and is stepped up to high voltages for long-distance transmission, since high voltages reduce transmission losses. At a substation near the final destination, a step-down transformer reduces the transmission voltage for distribution to both industrial and residential customers. At this point, distribution lines deliver electricity from the substation to end-consumers. In this work, we focus primarily on the large number of small-scale residential consumers in the grid, since they represent the vast majority of end-points in the distribution network.

7.3.1.1 Distribution Network

Figure 7.1 highlights the basic structure of electricity distribution in the grid. Electricity is fed into a bulk power substation, or a subtransmission station, which service a few “load areas” of customer demand. The bulk power substation routes the electricity to distribution substations. A distribution substation may then route the power to thousands of homes [1, 3]. Before being delivered to a building, distribution transformers near the building steps down the voltage of electricity. The number of consumers fed by a single distribution transformer varies: several homes may be fed off a single transformer in urban areas, or rural distribution may require one transformer per consumer [15].

In general, multiple distribution transformers may be connected in parallel. However, due to a lack of access to the distribution graph of an existing network and, for simplicity, in this chapter, we assume the topology of the distribution network as shown in Figure 7.1. We base this simple model on information that is available in public domain [1, 3, 15], and use it in our experimental evaluation. Here, we assume each distribution transformer supports five homes, each distribution substation serves 500 transformers, and two distribution sub-stations are served by one bulk power substation. While our absolute results are specific to this simple model of a distribution network, we believe that many of our key insights are applicable to a range of real topologies, since we base our topology on publicly-available information. Importantly, our methodology and analyses extends to other types of distribution networks.

7.3.1.2 T&D Losses in the Grid

A fraction of electricity is lost in transmission and distribution. In the US, nationally, roughly 6% to 6.5% of the total electricity is lost each year [51]. Losses are generally divided equally between transmission and distribution. For example, in




Figure 7.1. Typical electricity distribution hierarchy.

New York, transmission losses accounted for a total of 3.18% loss, while distribution losses accounted for the loss of 3.3% of the total annual electricity [9]. We use these loss values in our evaluation.

7.3.2 Electric Utility's Generation Costs

An electric utility generates, transmits, and distributes electricity for sale in the electricity market [16]. A consumer's electric bill is generally divided into three categories related to electricity's generation, transmission, and distribution, as listed below [13, 17, 24].

Energy Charge. Consumers are charged based on the total amount of energy, in kilowatt-hours (kWh), they consume over a billing period. This charge incorporates the cost for a utility to generate the energy or buy the energy on the open market.

Distribution Charge. Consumers are charged a fee to enable utilities to recover the cost of operating and maintaining the distribution system. This charge typically has two components: an energy component, based on the amount of kWh consumed over the billing period, and a peak power component, based on the highest peak power demand in kilowatts (kW) over the billing period [24].

Transmission Charge. Consumers are charged a fee to enable utilities to recover the costs related to the delivery of electricity over high-voltage transmission lines. This energy is generally purchased from a third-party and not generated by the local utility. As with the distribution charge, this charge has an energy component and a peak power component [24].

In some cases, consumers are not charged for energy, distribution, and transmission individually, but rather, the charges are included as part of the electricity rate. In addition to these costs, utilities also have expenditures related to the cost of materials and supplies and capital (including depreciation).

Expenses that are dependent on the total energy consumption are dictated by the pattern of end-user consumption, which cannot be controlled by a utility. However, the generation, transmission, and distribution costs incurred as a result of demand peaks can be reduced by curtailing the peaks. In addition, reducing demand peaks enables utilities to gain savings from avoided electricity costs, which include the marginal cost to produce and deliver one more unit of electrical energy. The avoided cost consists of two components—avoided energy costs (\$/MWh) and avoided capacity costs (\$/kW-month) [5]—which represent lower generation costs and the need for less peak capacity from lower peak demands.

7.3.3 Energy Storage to Lower Utility’s Costs

Energy storage devices can be used to store energy during low demand periods, which can then be used later to satisfy customer demands during peak demand pe-

riods, thereby reducing the net peak on the higher levels of the grid. As capital and operational expenses of the grid are largely determined by the peaks, energy storage can cut these expenses for the grid.

7.3.3.1 Energy Storage Technologies

We examine the potential for the following energy storage technologies to reduce an electric utility's distribution costs.

Compressed Air Energy Storage (CAES): With Compressed Air Energy Storage, off-peak grid power is used to compress air underground. Later, when the energy is needed, this compressed air is released to power an electric generation and produce electricity. These systems are typically large, often requiring significant real-estate for storing compressed air [40], e.g., underground tanks.

Ultra-capacitors (UC): Ultra-capacitors operate similarly to electrostatic capacitors, except that they can hold significantly more energy in a size similar to that of conventional capacitors [35]. UCs are now often being used for large-scale uninterruptible power supplies (UPS) in data centers, hospitals, industrial buildings, etc. [23].

Flywheels (FW): Flywheels store kinetic energy in rotating discs. These discs are made to turn a generator for producing electricity. Flywheels can be very efficient in storing energy over short durations; however, they have high self-discharge rates due to losses from friction [18, 35]. One example of a flywheel energy storage plant is the Beacon Power plant in New York [14].

Lead Acid batteries (LA): Lead-acid batteries are one of the most widely used energy storage devices. They have long been the primary technology for stationary energy storage at both grid-scales and in off-grid homes [21].

Lithium-Ion batteries (LI): Lithium Ion batteries are the most popular type of rechargeable batteries; they are known for their relatively high efficiency and energy

ESD	CAES	UC	FW	LA	LI
Efficiency (%)	68	95	95	80	85
Discharge:Charge Rate	4	1	1	10	5
Self-discharge (%per day)	low	20	100	0.2	0.1
Energy Density (Wh/L)	6	30	80	80	150
Power Density (W/L)	0.5	30000	1600	125	450
Ramp Time (sec)	600	0.001	0.001	0.001	0.001
Max DoD (%)	100	100	100	80	80
Energy Cost (\$/kWh)	50	500	1000	200	525
Cycle Lifetime	15000	100000	12000	2000	5000
Expected Lifetime (Years)	20	20	15	4	10

Table 7.1. ESD Parameters [40] [35] [111] [94]

density [22]. Lithium Ion batteries are the primary technology in mobile systems, e.g., electric vehicles, due to their light weight in comparison with lead acid batteries. However, these batteries are now being deployed in conjunction with renewable energy to provide energy storage for homes, as evidenced by Tesla’s recent introduction of the PowerWall home battery based on lithium-ion technology [25].

While diesel generators, and other captive sources can also be considered energy storage devices, we do not consider them separately here. Pumped hydroelectric is another widely used storage technology in the grid; however, since it requires significant infrastructure, it is not readily deployable in the distribution networks.

7.3.3.2 Energy Storage Characteristics

Below we list key characteristics of the energy storage devices that are relevant to our optimizations. Table 7.1 list the default values of the parameters used in our study, which we derive from various scientific studies. Note that these parameters are inputs to our framework and while they vary significantly across technologies, we do not further consider the impact of varying them for a particular storage technology We are specifically interested in how these characteristics yield different trade-offs when

placing various storage technologies at different levels of the distribution hierarchy to minimize a utility's distribution costs. We model the following characteristics:

Energy Storage Capacity: The energy storage capacity represents the total amount of energy that a device can store. Generally, the capacity is expressed in kilowatt-hours (kWh).

Maximum Charge and Discharge Rates: Usually expressed as E-rate, the maximum charge and discharge rates are a measure of the rate at which a battery can be charged or discharged relative to its total capacity [4]. For example, a 2E discharge rate is the discharge rate necessary to fully discharge the battery in half an hour.

Efficiency: Use of energy storage results in energy loss due to energy conversion. We employ a constant efficiency factor for each energy storage technology to capture these losses; e.g., typical lead-acid batteries are 80% efficient.

Self-Discharge Rate: The self-discharge rate is a phenomenon in energy storage by which the ESD loses stored energy merely with passage of time. The self-discharge rate can be significant for some technologies, such as flywheels. For energy storage technology k we model its self-discharge per unit time as a constant factor μ_k .

Cycle Lifetime: A ESD's lifetime is usually expressed in terms of number of charge-discharge cycles. Typically, ESD lifetime is measured based on the number of cycles as a function of the depth of discharge (DoD). For a given energy storage technology, we limit its DoD and the number of charge-discharge cycles at the given DoD over a given time horizon; thereby, we control the lifetime of an energy storage device, and capture its amortized per unit energy storage cost over its lifetime in our model.

Energy Density: The energy density is the nominal battery energy per unit volume (Wh/L). The energy density determines the battery size required to achieve a given energy output [4].

Power Density: Power density is defined as the maximum available power per unit volume (W/L). The power density determines the battery size required to achieve a given power output [4].

Power Ramp Time: The Power Ramp Time is the start-up latency associated with a given energy storage technology before it can start delivering its maximum power. This ramp-up is similar to the start-up acceleration in vehicles. Ramp times of most storage devices are very low, however for compressed air storage the ramp time may be several minutes.

7.4 Problem Statement

Although energy storage can reduce peaks and cut costs, the problem of storage deployment presents several interesting tradeoffs. Peak reduction at a given level of the grid's hierarchy enables provisioning the infrastructure at that level, as well as higher levels, for a lower peak. Therefore, storage deployment at lower levels of the hierarchy appears more beneficial than reducing the peak only at higher levels. However, in general, peaks at higher levels of the hierarchy are smaller than the sum of individual peaks at the lower levels, such as at homes; this occurs because individual homes peak at different times. The statistical multiplexing gains due to the spreading of individual peaks over time makes aggregate peaks at higher levels smaller. Therefore, deploying energy storage at the higher levels would require much less energy storage capacity, and hence lower aggregate energy storage costs.

In addition to deployment choices, the choice of storage technologies also presents tradeoffs in their cost, lifetime, efficiency, energy density, etc. For example, compressed air energy storage has a low energy cost and long expected lifetime, but a low energy efficiency and requires significant space for deployment. In contrast, lead-acid batteries have a higher energy-efficiency and a smaller form-factor, but also much higher energy costs and much shorter expected lifetime. Furthermore, differ-

ent storage technologies are suitable for shaving different types of peaks: compressed air storage is suitable for wide peaks, lead-acid batteries work well for less frequent medium-width peaks, and ultra-capacitors are best for very narrow peaks (up to a minute). Since the nature of the peak demand depends on the level of the grid hierarchy—medium-width peaks are more likely at homes, whereas wide peaks are frequent higher in the hierarchy—the best choice may differ at each level.

In this work, we address the problem of deploying energy storage across a distribution grid hierarchy to cut a utility’s distribution costs. As there are a number of variables involved, such as large distribution hierarchies, time varying demand profiles, different types of storage technologies, and a range of electricity pricing plans, it is not easy to formulate a heuristic solution for this problem. Therefore, we frame it as an optimization problem. We define the problem as follows: given an electricity distribution network, an estimate of power demands, and a set of available storage technologies, the problem is to find an optimal sizing and placement of energy storage devices across the distribution hierarchy so as to minimize a utility’s expenses (Section 7.3.2) for distributing electricity. Our framework is general enough to provide storage provisioning solutions for a range of consumption profiles, electricity pricing plans, storage technologies, and distribution networks.

7.5 Energy Storage Provisioning and Control Framework

We now present our optimization framework for energy storage provisioning. We intend the framework to provide storage deployment solutions for a distribution hierarchy with the goal of optimizing a utility’s cap-ex and op-ex. Inputs to the framework include power demand, the distribution network topology, and cap-ex and op-ex costs. The framework solution then outputs the optimal choice, placement, and size of energy storage devices across the hierarchy, along with optimal energy storage control patterns.




Figure 7.2. Illustrative graph depicting distribution hierarchy

Throughout the formulation we assume the grid distribution network is a directed graph $G = (V, E)$ —as shown in Figure 7.2—where V is the set of nodes (i.e., homes, distribution transformers, substations, etc.) and E represents the directed edges between these nodes. In Figure 7.2, circles and squares represent the nodes, and arrows represent the edges. Two types of nodes are shown in the figure: squares are the leaf nodes and circles are the non-leaf nodes. In the distribution network, leaf nodes typically represent homes. In our formulation, all leaf nodes are represented by the set labeled LeafNodes. Node r (Figure 7.2) represents the root node. Finally, across the formulation, k represents the k -th energy storage technology (out of K) and t is the time interval (between 1 to T).

7.5.1 Inputs

Power Consumption (Demand): Broadly, there are two types of problems associated with energy storage deployments in the grid’s distribution network. The first problem is determining the proper energy storage capacity and where in the hierarchy to deploy it. The second problem is determining how to charge-discharge the energy storage device to clip the peaks and realize cost savings. In this chapter, we solve the first problem—the energy storage sizing and deployment problem, i.e.,

figuring out how much energy storage should be deployed and where. Therefore, we assume that historical power consumption traces are available, and we use these for future provisioning of storage in the distribution network. We assume that prior work can be employed to derive an accurate power demand time-series at each home (e.g., [28] [107]). Further, utilities have extensive power consumption logs over time for their customers, which can also be used as an input to our optimization framework. We divide time into T slots, each of length I . For home u we assume its power demand to be a time series $UsrDmnd_{u,t}$, where $t \in [1, T]$. We present results for both real and synthetic consumption power time series.

Capital Expenditure (Infrastructure Cost): We model two types of infrastructure costs: first, maintenance and upgrade cost, second, avoided (or marginal) capacity costs. These costs vary significantly between utilities and between locations within utilities: ranging from \$2.51/kW/month to \$46.34/kW/month [12, 65, 71]. In our experiments, we consider a cap-ex saving in the range of \$6/kW/month to \$30/kW/month resulting from peak reduction. These savings are obtained by reducing a watt of consumption from the maximum power draw $Peak_u^{max-orig}$. At any time t , power draw at node u is given by the sum of power draws of all the nodes in the sub-tree with root at the node u at time t and net energy drawn by the energy storage devices at its vicinity; the sum is denoted by $Demand_{u,t}$. The corresponding size of the tallest shaved peak at u is denoted by $Peak_u^{shvd}$.

Operational Expenditure (Tariffs): Electricity tariffs are good indicators of the operational costs incurred in distribution of electricity. Most prevalent electricity tariff models charge customers for their total energy consumption, i.e., customers have to pay a flat $\$C_e^{unit}/kWh$ of their consumption. Recently, to shave the peak demand on their grids, utilities have introduced a peak penalty on the tallest consumption peaks across the billing cycle. Typically, the peaks are computed as a sliding window of 30 minutes over the billing period. End users then pay a penalty of $\$/kW$ based

on the tallest peak. For a utility, this peak penalty translates to the energy charged in generating the peak power and the cost incurred in routing the electricity to the distribution network (as in [20]). Note that our model for the value of peak reduction derives directly from the way electric utility companies charge for the peaks, therefore the marginal value of peak reduction is constant.

Our model for capital and operational expenditure for the utilities is derived from the information available in electricity bills and reports published by the utility companies, as reported in [5, 12, 13, 17, 24]. As utilities pass on their costs to the customers, we believe utility bills closely model the actual expenses of electric utility companies. In addition, several real-world factors, such as resource availability and market price, affecting utility expenses are accounted for while computing the avoided costs. For example, among other factors, [12] accounts for wholesale electric energy price, projections of natural gas prices, generation capacity costs, cost of controlling CO₂ emissions, and the effect of implementation of anticipated federal regulations.

7.5.2 Optimization Problem Formulation

Decision Variables: All notations used in the framework are summarized in Table 7.4. Our decision variables capture both the sizing and placement of energy storage and how to operate it to minimize the peak demand. The energy storage capacity of a storage technology of type k deployed at node u is denoted by $C_{k,u}$. The average power fed into and drawn out of an energy storage device at u during time slot t is denoted by $S_{k,u,t}$ and $D_{k,u,t}$.

Optimization Objective: Our optimization objective is:

$$\text{Minimize}(\text{CapEx} + \text{OpEx} + \text{StorageCost}) \quad (7.1)$$

The objective function has three components: capital expenditure ($CapEx$), operational expenditure ($OpEx$), and $StorageCost$. Each component is normalized to our experiment’s time horizon.

CapEx includes the capital expenses due to infrastructure deployment for electricity distribution. Assuming α_u is the maintenance, upgrade, and marginal capacity costs for each watt of infrastructure provisioning at node u , *CapEx* is given by equation (7.2).

$$\text{CapEx} = \sum_{u \in V} \text{Peak}_u^{\text{shvd}} * \alpha_u \quad (7.2)$$

OpEx is the expected utility operational costs in electricity distribution and can be represented as in (7.3). The OpEx has three components, respectively: peak surcharge paid on the tallest demand peak served by the utility, electric energy cost paid on the total electricity served to the customers, and additional avoided costs of electricity incurred as a result of inefficiencies in energy storage devices.

$$\begin{aligned} \text{OpEx} = & \text{Peak}_{\text{root}}^{\text{shvd}} * b + \sum_t \text{Demand}_{\text{root},t} * I * a_t \\ & + \sum_{k,u,t} (S_{k,u,t} - D_{k,u,t}) * I * \gamma \end{aligned} \quad (7.3)$$

In the above, b is the per unit surcharge (\$/kW) on the tallest peak and a_t is the unit cost of electricity at time t . $\text{Peak}_{\text{root}}^{\text{shvd}}$ is the tallest peak seen at the root node, which incurs the peak surcharge. γ is the avoided electric energy cost (AEEC in \$/MWh), as some energy is lost in the energy storage conversion process, this lost energy incurs extra avoided costs which is added to the utility operational costs. Note that $\text{Demand}_{\text{root},t}$ captures the total load including losses in transmission and storage charge-discharge.

StorageCost is the cost of energy storage deployment across the grid, given by (7.4), where $\beta_{k,u}$ is the amortized cost of the energy storage device k at node u per unit energy adjusted to its lifetime.

$$\text{StorageCost} = \sum_{u,k} C_{u,k} * \beta_{k,u} \quad (7.4)$$

The lifetime of an energy storage device depends on several factors such as the depth of discharge (DoD)—a battery lasts longer for smaller DoD. The value of $\beta_{k,u}$ is an input and is determined by the DoD and the set number of charge-discharge cycles over the time horizon. In this chapter, in addition to the storage costs, $\beta_{k,u}$ includes the cost of the power conversion system, balance of plant, operation and maintenance [35].

Constraints: We assume that the state of charge in all storage devices at the end of the time horizon is same as their state at the beginning, as stated in (7.5).

$$E_{k,u,1} = E_{k,u,T+1}, \forall k, u \quad (7.5)$$

At any time, an energy storage device can only store energy between a lower threshold dictated by its allowed depth of discharge and a maximum capacity; this is captured by (7.6).

$$(1 - DoD_k^{max}) * C_{k,u} \leq E_{k,u,t} \leq C_{k,u}, \forall k, u, t \quad (7.6)$$

For each storage device, the rate at which energy can be drawn from and fed into the device is bounded by its discharge (r_k^{disch}) and charge (r_k^{charge}) rates, as determined by the underlying storage technology. This is captured in equations (7.7) and (7.8).

$$0 \leq D_{k,u,t} \leq C_{k,u} * r_k^{disch}, \forall k, u, t \quad (7.7)$$

$$0 \leq S_{k,u,t} \leq C_{k,u} * r_k^{charge}, \forall k, u, t \quad (7.8)$$

Equation (7.9) is the energy conservation constraint, which states that the total energy drawn out of the energy storage is never greater than the energy charged to the battery multiplied by the storage efficiency (e_k).

$$\sum_{t=1}^T D_{k,u,t} \leq e_k * \sum_{t=1}^T S_{k,u,t}, \forall k, u \quad (7.9)$$

$$\begin{aligned} Demand_{u,t} = & \frac{1}{\eta} \left(\sum_{v:(u,v) \in E} Demand_{v,t} \right) + \sum_k S_{k,u,t} \\ & - \sum_k D_{k,u,t}, \forall u \in (V - \text{LeafNodes}), t \end{aligned} \quad (7.10)$$

Net power consumption at any non-leaf node u at time t , denoted by $Demand_{u,t}$, is determined by the sum of net power consumption at all its child nodes and the net power drawn/delivered by the energy storage devices deployed at the node, equation (7.10). (η takes care of the electricity lost due to transmission inefficiencies between node u and its children.) For example, in Figure 7.2, if transmission efficiency is 100%, the net power draw at node u_1 is given by the sum of net power drawn at its child nodes v_1, v_2 and the storage devices at u_1 . On the other hand, net power draw at the leaf nodes, i.e., homes, is given simply by the sum of home's electricity demand and net electricity drawn/delivered by energy storage at the home, as in equation (7.11).

$$\begin{aligned} Demand_{u,t} = & UsrDmnd_{u,t} + \sum_k S_{k,u,t} \\ & - \sum_k D_{k,u,t}, \forall u \in \text{LeafNodes}, t \end{aligned} \quad (7.11)$$

Equation (7.12) bounds the tallest peak ($Peak_u^{shvd}$) seen by u .

$$0 \leq Demand_{u,t} \leq Peak_u^{shvd}, \forall u, t \quad (7.12)$$

We also model the following energy storage characteristics in our framework, which are presented in 7.A: the rate at which output power of a battery can increase, as




Figure 7.3. Composition of capex and peak penalty costs for Long-Term Contract.

some energy storages (such as compressed air) may take up to few minutes before delivering maximum rated power; battery self-discharge, as batteries lose some energy simply with passage of time; lifetime of the storage, as it affects the costs in the long term; volume needed for deploying energy storage, as some form of storages may need significant space, e.g., flywheels.

7.6 Evaluation

7.6.1 Experimental Setup and Methodology

Avoided Energy Cost (γ)	\$3.53/MWh [12]
CapEx(Low)	\$6/kW/month
CapEx(Medium)	\$15/kW/month
CapEx(High)	\$30/kW/month
Energy cost (Contract pricing)	\$ 0.05/kWh
Peak Penalty (Transmission + Energy)	\$20/kW/month
Volume V_{homes}^{max}	10L (\sim Car Battery)
Volume $V_{transformer}^{max}$	25L (\sim 2.5X Car Battery)

Table 7.2. Experiment Parameter Values.

Configuration and Parameters: Our evaluation uses the grid distribution hierarchy shown in Figure 7.1. As explained in Section 7.3, we assume each distribution transformer supports five homes, each distribution substation serves 500 transformers, and two distribution substations are served by one bulk power substation. Storage devices can be placed at any of the levels in Figure 7.1. For simplicity, we present results at three levels: Homes, Transformers, and Substations, including distribution and bulk power substations.

We evaluate our framework with two op-ex cost models: a *long-term contract* and *day-ahead* model. *Long-term contract* represents the scenario where the utility either owns most of its generation or buys its energy from third parties under contracts. We adapt a real utility contract, available at [20], for evaluation. Since we include distribution costs as part of the distribution cap-ex costs, we subtract the distribution costs from the peak penalty and use the final values in Table 7.2. *Day-ahead* represents the case where a utility buys all of its electricity in day-ahead markets. However, the utility still incurs the peak penalty due to transmission. We use the day-ahead prices for March, 2014 from ISO New England [19]. We consider the cap-ex costs ranging from low to high, where *low* = \$6/kW/month, *medium* = \$15/kW/month, and *high* = \$30/kW/month ([12, 65, 71]).

The exact distribution of cap-ex, i.e., infrastructure cost (α), for the grid’s distribution hierarchy is not available. Thus, in this chapter, we assume these costs are equally split across all the levels. As the space available at homes and distribution transformers is limited, we use a conservative value of 0.01m³ (or roughly the size of a car battery) for the energy storage volume at homes; for transformers we set the volume to 0.025m³. Substations are built on large areas, so we do not constrain the available volume at substations. All results have been amortized to daily costs and savings, which includes the amortized cost of storage over its lifetime. We present results for the five ESDs, i.e., $K = 5$, discussed earlier.




Figure 7.4. Electricity demand on a representative weekday.

We use the terms hybrid or multi-level energy storage to imply a combination of different storage technologies at a given node. Note that by using real-world day-ahead market prices and real utility consumption traces, we experiment with fine-grained time-varying prices and power consumption. Also, as we are solving an optimization problem, the computed storage values can be fractional. Although the computed numbers could differ from the actual storage capacity deployed in practice, we do not expect significant deviations from the computed values.

Workloads: For empirical evaluations, we use power consumption traces obtained from a local electric utility collected over one month (March 2014). Our traces are representative of consumption in a real electric grid. The traces contain power consumption data from 5000 homes at five minute granularity. In aggregate, we have more than 1.4 million unique power measurements. The average daily energy consumption of individual homes in our traces range from 15 kWh to 73 kWh.

Figure 7.4 shows the aggregate grid demand on a representative weekday. The figure shows that the homes peak between 6AM to 9AM (breakfast peak) and 6PM to 9PM, i.e., at dinner time. The pattern is expected based on typical work patterns,

where home’s electrical activity is concentrated after office hours. Throughout this chapter, unless specified otherwise, we compute the peaks at a 30 minute granularity. To gain insights into the results, we present a detailed analysis of our results on a randomly picked weekday. Later, we present the results on traces for March, 2014.

7.6.2 Potential Savings from Storage

Can Energy Storage Reduce Distribution Costs? We first evaluate the savings from deploying only lead-acid batteries at a single level, i.e., either at homes, or transformers, or substations. Figure 7.5(a) shows the percentage distribution cost savings corresponding to a low, medium, and high capital expenditure for the long-term contract pricing plan. Figure 7.5(a) depicts the daily percentage cost savings for *long-term contract*, which shows that even a lead-acid deployment only at homes under a low cap-ex can cut costs by 3.75%. Savings increase as cap-ex increases. Also, note that for all single-level deployments, deployment at homes shows the maximum savings. This happens because peak shaving at the lowest level (homes) provisions the infrastructure at all levels for a lower peak, thereby saving significantly in cap-ex. In contrast, savings from deployments at the transformer-level are the lowest because of the limited volume availability, which is much smaller than at homes and substations.

Result: Deploying energy storage, in this case lead-acid batteries, at a single level of the hierarchy modestly reduces costs. Deploying at the lowest level, i.e., in homes, shows the greatest savings, since it also affects peak demands at higher levels.

Is Multi-Level Energy Storage Deployment Beneficial? Since related work suggests deploying lead-acid batteries only at homes, we next evaluate the impact of deploying lead-acid batteries across multiple levels of the hierarchy on savings. Figure 7.5(b) compares the savings from multi-level lead-acid deployment with its deployment only at homes (single-level). Savings are shown corresponding to low, medium, and high capital expenditures under the long-term contract pricing plan.




Figure 7.5. Savings from deploying lead-acid battery storage at (a) single level and (b) multiple levels under the long-term contract pricing plan.

For all cap-ex values, savings from a multi-level deployment surpasses that of a single-level deployment at homes. In addition, for high cap-ex, the daily cost savings from multi-level lead-acid energy storage deployment shows an increase of more than 60% compared to single-level deployment at homes.

Result: Deploying one ESD type, in this case lead-acid batteries, at multiple levels of the grid’s hierarchy further increases the cost savings up to an additional 60%.

Is Hybrid Energy Storage Deployment Beneficial? We next evaluate the additional savings possible from deploying multiple, i.e., hybrid, storage technologies at any given (single) level over a corresponding lead-acid storage deployment.

Figure 7.6 compares the percentage cost savings from hybrid energy storage deployment at single levels with the corresponding lead-acid deployment. Savings are shown for storage deployment at homes, transformers, and substations. Figure 7.6(a) and (b) shows results for the long-term contract, 7.6(a) is with low capex and 7.6(b) is with high capex.

We find that deploying hybrid energy storage boosts savings compared to lead-acid deployments, e.g., in Figure 7.6(a) hybrid energy storage at substations increases savings by 103%. Note that as opposed to our observation in Figure 7.5(a), in Fig-

ure 7.6(a), savings for hybrid deployment at substations is higher than that of homes. This occurs because under low cap-ex, substations can save more from greater peak shaving with hybrid energy storage, in large part, because there is no volume constraint at substations. However, as cap-ex increases, it becomes a greater component of the cost (as shown in Figure 7.3); as a result, the savings from deployment at homes is more than the substation’s savings for high cap-ex.

Result: Employing multiple ESD technologies in hybrid further increases savings relative to only using lead-acid batteries at any single level to as much as $\sim 10\%$. Hybrid deployments are able to best match the usage pattern at any given level with the characteristics of the energy storage device.




Figure 7.6. Savings from deploying hybrid storage technologies at a single-level for low and high cap-ex costs under long-term contract.

Is Multi-Level Hybrid Energy Storage Deployment worth it? Figure 7.7 shows how a multi-level multi-technology storage deployment can further increase savings over, first, any single level multi-technology deployment (7.7(a)), and second, any multi-level lead-acid (single storage technology) deployment (7.7(b)). Savings are shown for low, medium, and high values of cap-ex. Figure 7.7(a) shows that multi-level hybrid solution outperforms all the single-level hybrid solutions under all cap-ex values. For instance, 52% improvement over best single level solution under

high cap-ex. We further find that a multi-level multi-technology storage deployment saves more than multi-level lead-acid deployment, as shown in Figure 7.7(b): for example, 83% increase in savings under low cap-ex. This increase in savings results from stringent volume constraints at lower levels, where multi-technology solutions gain an advantage by including storage technologies with higher power and/or energy density, as opposed to lead-acid storage.

Result: A hybrid, multi-level deployment results in the greatest savings, by as much as 12%, since it is able to exploit different energy storage device characteristics at each level of the grid hierarchy, which exhibits different usage patterns.




Figure 7.7. Multi-Level Hybrid ESDs v/s: (a) Single-level Hybrid ESDs and (b) Multi-level lead-acid batteries under long-term contract.

How do the savings change under other pricing plans? While the above results depict savings for the long-term contract pricing plan, we have repeated all of the above experiments for the day-ahead pricing plan. In each case, we find similar cost savings for the day-ahead pricing when compared to the long-term contract plan. For example, Figure 7.8(a) and (b) depicts savings from a multi-level LA deployment and a multi-level hybrid deployment, respectively. As can be seen, the corresponding savings under long-term contract pricing, depicted in Figure 7.5(b) and 7.7(a), are similar to that depicted in Figure 7.8(a) and (b) under day ahead pricing. Since all

experiments show similar results, we omit the remaining graphs for brevity (see [80] for detailed results). **Result:** Overall our experiments show that the savings due to energy storage are not specific to a pricing plan and hold for both long-term contract and day ahead pricing.





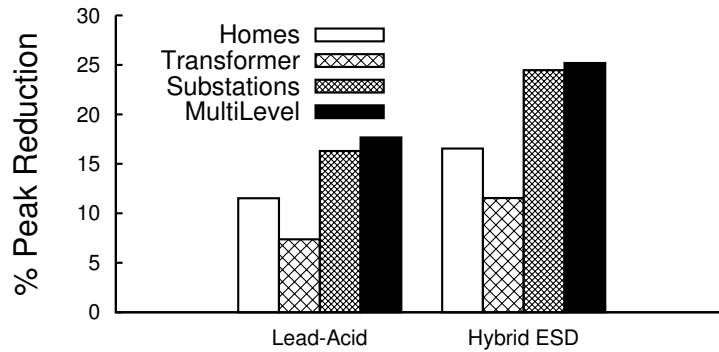
Figure 7.8. Savings under day-ahead pricing for multi-level lead acid batteries and multi-level hybrid storage.

Peak Reduction: Figure 7.9 shows the aggregate percentage peak reduction across the grid with lead-acid only, and multi-technology storage deployments. For each of the cases, we present results for both single-level deployments at homes, transformers, substations, and multi-level deployment across the hierarchy. Results are presented for long-term contract (7.9(a)) and day-ahead pricing (7.9(b)). Only medium cap-ex numbers are shown, as the numbers for other cap-ex are similar. Figure 7.9(a) shows even a lead-acid deployment at homes achieve a peak reduction of 11.6%, which is further increased to 16.7% by a hybrid energy storage deployment.

As we have seen, hybrid solutions have the advantage of choosing storage technologies with greater power/energy density and discharge rates. In addition, note that peak reductions achieved by substation and multi-level deployments (both hybrid and lead-acid) are very close; however, the savings achieved by them are much different (e.g., 32% difference in Figure 7.7(a)). Although they get similar aggregate peak reductions, the multi-level approach saves more in cap-ex by deploying energy stor-



(a) Long-term Contract



(b) Day-Ahead Pricing

Figure 7.9. Aggregate Peak Reduction from Lead-Acid only and Hybrid ESDs for CapEx(Medium).

age devices at the lower levels. Figure 7.9(b) shows equivalent results for day-ahead pricing.

Result: Hybrid ESD deployments at multiple levels results in the greatest reduction in peak demand (by as much as 25%) compared to deploying hybrid ESD at individual levels or using only lead-acid batteries as the ESD.

Optimal Energy Storage Placement and Configuration: To give insight about the different types and configurations of storage technologies selected by our framework, Table 7.3 presents the energy storage configuration under the contract pricing

for medium cap-ex. Configuration for the other cap-ex values and pricing are similar. To give an idea about absolute numbers, we include dollar savings and cost values. First, we see that if we are to use a single-level solution—just lead-acid batteries—deployment at homes does provide the best savings, because of the cap-ex gains at all levels. Second, volume constraints play an important role in limiting the benefits of lead-acid batteries in the lower level of the hierarchy; therefore, a single-level hybrid storage solution is able to increase savings by deploying a higher energy and power dense storage device, such as lithium-ion batteries, e.g., 42.5% increase at the homes. In addition, at substations where there is no volume constraint hybrid solutions employ a combination of lithium-ion, ultra-capacitors, and compressed air energy storage and further increases the savings by 15%.

Compressed air is the cheapest form of storage, however, it has a long start-up delay; ultra-capacitor and lithium-ion can be used to bridge this delay; ultra-capacitors have a very high power density, which helps in shaving tall narrow peaks, and their low energy density is complemented by lithium-ion energy storage. Finally, with the freedom of hybrid storage for multi-level deployment, we get maximum savings by deploying lithium-ion at lower levels and compressed air storage at the top level of the hierarchy.

Result: Using different ESDs at different levels of the grid’s hierarchy result in significant differences in costs and savings.

Energy Storage Costs: As the numbers in Table 7.3 show, the cost of energy storage is a small fraction of the total daily costs without storage devices. For example, a hybrid storage solution at the substations is only 1.83% of the total daily cost. A hybrid storage deployment at lower levels costs more than deploying lead-acid batteries because lithium-ion batteries are more expensive. However, due to their higher energy and power density they also save more. Even the most expensive

	LA Single (\$ save, \$ cost)	Hybrid Single (\$ save, \$ cost)	Hybrid MultiLevel (\$ save, \$ cost)
Substations	LA (880.25,416.74)	LI + UC + CAES (1544.73,394.27)	CAES
Transformer	LA (558.87,125.0)	LI (841.70,236.25)	LI
Homes	LA (940.06,249.77)	LI (1340.36,472.20)	LI (2039.53,766.21)

Table 7.3. Storage Configuration and Placement (Long-term Contract, Medium capex): (Savings(\$/day), Storage costs (\$/day)). Total cost without storage is \$21.5k/day.

energy storage deployment, i.e., multi-level hybrid, incurs less than 3.6% of the total costs; most of its costs are from the lithium-ion deployment at the lower levels.

Result: The cost of energy storage is a small fraction of the total daily distribution costs without any energy storage capacity.

7.6.3 Longer-term Savings

For computational tractability, so far, we have presented results on a single day. However, to show that the savings hold over longer periods, we conducted experiments over an entire month. Figure 7.10 shows the average daily cost savings for the month of March, 2014 from our traces. Due to space constraints, savings are shown only for low and high cap-ex for day-ahead pricing; three types of deployments are shown: lead-acid at homes, multi-level lead-acid, and multi-level hybrid. Here, hybrid multi-level can save up to 11.7% for high cap-ex, and up to 9.8% for low cap-ex, which outperforms the multi-level lead-acid (low cap-ex) by 190%. The general trends in the figure are similar to what we have already seen. As peaks become taller, there is a greater opportunity for savings.




Figure 7.10. Average daily savings for March (Day-Ahead).

7.7 Conclusion

In this chapter, we study the novel problem of ESD deployment across distribution grid hierarchy for enabling automated Demand Response (DR). We present a generalized optimization framework for ESD deployment and control across the distribution grid hierarchy. Our framework can provide ESD provisioning solutions for a range of consumption profiles, electricity pricing plans, ESD technologies, and distribution networks. We showed that ESD provisioning can save up to 12% daily costs in distribution for the utility companies. In addition, we also present several key insights regarding ESD deployment in the distribution network.

Our work has some limitations, which we plan to address as part of future work. For example, our current model assumes the marginal value of reducing peak usage is constant, whereas in practice the marginal value varies with the magnitude of the peak. We also do not consider the impact of renewable generation, which may alter the cost of reducing peak demand. Our models assume linearity to keep the problem tractable, although there are many characteristics of ESDs, and particularly batteries, that are non-linear, e.g., capacity as a function of discharge rate due to Peukert's

law. Finally, while our capital and operational cost estimates are based on publicly available sources, and we evaluate our system over a wide range of possible costs, e.g., high, medium, and low cap-ex, these estimates may vary widely across utilities, which would effect the possible savings in the real world. However, our methodology is general and can be applied to utilities with different costs and distribution hierarchies.

7.A Appendix

Below are the additional constraints of the framework presented in section 7.5. Table-7.4 defines the notations. Constraint (7.13) limits the rate at which an energy storage's output power can increase, this is similar to acceleration of vehicles. Here $constant_k = \frac{r_k^{disch} * I}{R_k^{ramp}}$, and R_k^{ramp} is the power ramp-up time. As batteries lose some energy simply with passage of time, we model this battery self-discharge in constraint (7.14). constraint (7.15) bounds the lifetime of the storage. As the lifetime is primarily determined by the number of charge-discharge cycles and the depth of discharge, (7.15) bounds the number of times an energy storage can be discharged to its allowed depth of discharge in the given time horizon. Finally, we restrict the maximum volume for storage deployment that might be available at node u in (7.16) and (7.17).

$$D_{k,u,t} - D_{k,u,t-1} \leq C_{k,u} * Constant_k, \forall k, u, t \geq 2 \quad (7.13)$$

$$E_{k,u,t} = (1 - \mu_k) * E_{k,u,t-1} + e_k * S_{k,u,t-1} * I - D_{k,u,t-1} * I, \quad (7.14)$$

$$\forall k, u, t \geq 2$$

$$\frac{I * \sum_{t=1}^T D_{k,u,t}}{DoD_k^{max} * C_{k,u}} \leq NumChDischCycles_k, \forall k, u \quad (7.15)$$

$$\sum_{k=1}^K \frac{C_{k,u}}{\phi_k^{energy}} \leq V_u^{max}, \forall u \quad (7.16)$$

$$\sum_{k=1}^K \frac{C_{k,u} * r_k^{disch}}{\phi_k^{power}} \leq V_u^{max}, \forall u \quad (7.17)$$

Symbol	Notation
$C_{k,u}$	Capacity of the k -th energy storage device (ESD) at node u of the grid hierarchy
C_e^{unit}	Unit cost of energy
a_t	Unit cost of electricity in interval t
$D_{k,u,t}$	Average power drawn from the k -th energy storage at node u of the grid hierarchy in time interval t
DoD_k^{max}	Maximum depth of discharge for k -th type of storage
$E_{k,u,t}$	Energy stored in k -th storage device at node u of the grid hierarchy in interval t
e_k	Efficiency of storage type k
I	Length of each time interval
$Demand_{u,t}$	Net power demand on grid at node u in interval t
$UsrDmnd_{u,t}$	User consumption at home node u in interval t
$Peak_u^{shvd}$	Maximum shaved peak seen by the node u
$Peak_u^{max-orig}$	Maximum original peak seen by the node u
b	\$/kW penalty on the tallest consumption peak
r_k^{charge}	Storage charging E-rate for the k -th energy storage
r_k^{disch}	Storage discharge E-rate for the k -th energy storage
R_k^{ramp}	Output power ramp up time of storage k
$S_{k,u,t}$	Average power fed into the k -th storage at node u in interval t
T	Total number of time intervals
ϕ_k^{energy}	Energy density of k -th storage technology; nominal energy per unit volume
ϕ_k^{power}	Power density of k -th storage technology; nominal power per unit volume
V_u^{max}	Maximum volume available for energy storage deployment at node u
α_u	Cost savings for each watt of under-provisioning at u
$\beta_{k,u}$	Amortized cost of storage k per unit energy adjusted to its lifetime
γ	Avoided electric energy cost (\$/MWh)
η	Electric transmission efficiency in the distribution
μ_k	Self discharge rate of storage technology type k

Table 7.4. Optimization framework notations

CHAPTER 8

SUMMARY AND FUTURE WORK

8.1 Thesis Summary

This thesis has explored energy optimization techniques in buildings and distribution networks to make them energy efficient, cut end user’s electricity bills, and cut electric utility’s expenses in electricity distribution. We have proposed a set of techniques to achieve these optimizations without active user involvement and user inconvenience, while making aggregate electricity consumption profile grid-friendly.

Appliance Scheduling for Demand Flattening: Demand-side energy management is challenging, since it often requires active consumer involvement. Forcing people to think about how they use power is not effective in encouraging broader adoption of demand-side management. Therefore, first, we focused on quantifying the benefits of scheduling transparent background loads (such as A/Cs, refrigerators). We found that Least Slack First (LSF) algorithm is able to flatten household demand over each day. Importantly, we showed that LSF is useful over shorter (4-hour) peak usage periods, where demand is “peaky” and deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.

Employing Energy Storage to Cut Electricity Bills: Next, we explored how to lower electric bills by storing low-cost energy for use during high-cost periods. Here we developed an intelligent battery charging-discharging system called SmartCharge. We found that typical SmartCharge savings today are 10-15% per home with the potential for significant grid peak reduction (up to 20% with our data).

We also analyzed SmartCharge’s costs, and showed that recent battery advancements combined with an expected rise in electricity prices may make SmartCharge’s return on investment positive for the average home within the next few years.

Integrating Renewable Energy in Smart Buildings: We then extended SmartCharge to GreenCharge, which integrates on-site renewables in building’s electricity consumption. We showed that with GreenCharge cost savings are boosted to near 20% per home while retaining the potential for significant grid peak reduction (20% with our data). Finally, we analyzed GreenCharge’s costs, and showed that recent battery advancements combined with an expected rise in electricity prices and decrease in solar panel prices may make GreenCharge’s return on investment positive for the average home within the next few years.

Scaling Storage for Peak Reduction: After SmartCharge and GreenCharge, we examined the effects of using large scale distributed energy storage at buildings and homes throughout the grid to flatten grid demand. We found that as more consumers adopt energy storage, a number of problems arise that impact grid stability and generation costs. As a result, we proposed to augment traditional variable rate electricity pricing plans with a substantial peak demand surcharge, which incentivizes consumers to flatten their demand rather than shift it all to a low-price period. We then designed PeakCharge, which includes an online algorithm to minimize electricity costs in the presence of variable rates and the peak demand surcharge. We showed, with PeakCharge, the aggregate energy storage capacity to flatten grid demand by incentivizing consumers to flatten their own demand is within 18% of the minimum, optimal capacity to flatten grid demand in a centralized system.

Integrating Storage in Distribution Networks: Energy storage can help cut end-user electricity bills and electricity generation, and transmission & distribution costs for the utilities. However, most of the prior work has investigated deploying a single storage technology at a single level of grid hierarchy. To overcome these

limitations, we finally studied the novel problem of storage deployment across distribution grid hierarchy for enabling automated Demand Response (DR). We presented a generalized optimization framework for storage deployment and control across the distribution grid hierarchy. The presented framework can provide storage provisioning solutions for a range of consumption profiles, electricity pricing plans, storage technologies, and distribution networks. We showed that energy storage provisioning can save up to 12% daily costs in distribution for the utility companies.

8.2 Future Work

Here we present some of the future research directions that have emerged from the work in this dissertation.

Energy Optimizations for Group of Buildings: Modern buildings are getting increasingly smarter by having integrated building management systems (BMSs) that enables finer monitoring and control of the mechanical and electrical equipment in the buildings. In the future, we would like to leverage the emerging technology trends such as ubiquitous wireless communication, mobile computing, and cloud computing to design and implement large-scale extended-BMSs. These extended-BMSs would, additionally, coordinate consumption among a group of smart buildings to cut electricity bills for the buildings and make the aggregate demand grid-friendly.

Large-Scale Renewable Integration in Grid: In order to reduce carbon emissions and make electricity generation more sustainable, states are requiring electricity companies to produce an increasingly greater fraction of electricity from renewable sources. However, the intermittent and unpredictable nature of renewable generation requires a shift from the traditional electricity generation paradigm of ‘supply following demand’ to ‘demand-supply matching.’ To this end, we would like to investigate several approaches including: 1) intelligent demand-response, 2) combination of intelligent demand-response with geographically dispersed renewable resources

that compliment each-other to match the demand and supply, e.g., wind farms, solar farms, biogas systems, 3) energy storage integration in the grid.

Motivating End-User Sustainable Behavior: The greatest challenge and opportunity for sustainability lies in encouraging the end-users to adopt sustainable behavior. Traditionally, it has been difficult to encourage this transition because, 1) users do not have context for their energy consumption, 2) most users don't know how to cut their energy use and costs, 3) many users are not sufficiently motivated to change. However, with the advent of smart-meters, smart-phones, and cloud-computing, it is now possible to provide users with real-time feedback about their consumption, and its environmental, and economic impact, thereby encouraging and enabling them to make sustainable choices. Real-time consumption data from smart-meters can be shown to the users on smart-phones to make them aware of their usage. Further, insightful environmental and economic statistics about user's consumption, and tips and tricks to make it more sustainable can be shared using smart-phones.

Data Analytics for Sustainability: With ubiquitous energy monitor/sensor deployments, we now have access to fine-grain, real-time electricity consumption data. Data analytics on this data can answer several relevant questions for buildings, distribution networks, and the grids, such as: do occupancy patterns align with building schedules? What fraction of a building's energy is spent on heating/cooling? Does a home have inefficient appliances? Which transformers are likely to be overloaded? Can we use energy storage to extend lifetime of a transformer? Which customers are likely to benefit from a new pricing plan? What pricing plans incentivize customers to flatten peaks?

BIBLIOGRAPHY

- [1] Overview of CON EDISON System and LIC Network. <http://bit.ly/1SB6YXA>.
- [2] TED 5000. <http://www.theenergydetective.com/store/ted-5000>.
- [3] Utilities. http://www.nyc.gov/html/sirr/downloads/pdf/final_report/Ch_6_Utillities_FINAL_singles.pdf.
- [4] A Guide to Understanding Battery Specifications. http://web.mit.edu/evt/summary_battery_specifications.pdf, 2008.
- [5] Biennial Report of the North Carolina Utilities Commission. <http://www.ncuc.net/reports/CostAllocationReport09.pdf>, 2009.
- [6] U.S. Energy Information Administration. Annual Energy Review. www.eia.gov/totalenergy/data/annual/pdf/aer.pdf, 2009.
- [7] U.S. Department of Energy. Building Energy Data Book. <http://www.buildingsdatabook.eere.energy.gov/ChapterIntro2.aspx>, 2010.
- [8] Ontario Energy Board. www.ontarioenergyboard.ca/OEB/Consumers, 2011.
- [9] Assessment of Transmission and Distribution Losses in New York. Tech. Rep. PID071178 (NYSERDA 15464), EPRI, Palo Alto, CA, 2012.
- [10] Electric Generator Dispatch Depends on System Demand and the Relative Cost of Operation, August 17th 2012.
- [11] IEC 61850. http://en.wikipedia.org/wiki/IEC_61850, 2012.
- [12] Avoided Energy Supply Costs in New England: 2013 Report. <http://bit.ly/1LF0SzV>, July 2013.
- [13] Electric Bill Breakdown. <http://bit.ly/1KyBhLT>, 2013.
- [14] Beacon Power. <http://beaconpower.com/>, 2014.
- [15] Distribution Transformer. Wikipedia, 2014.
- [16] Electric Utility. Wikipedia, 2014.
- [17] Explanation of Unbundled Electricity Charges Residential Customers. <http://bit.ly/1KyBhLT>, July 2014.

- [18] Flywheel. <http://bit.ly/1FQbqwa>, 2014.
- [19] ISO New England - Hourly Data. http://www.iso-ne.com/markets/hrly_data/index.html, 2014.
- [20] Large General Service Contract Schedule. <http://www.hged.com/customers/payment-billing-rates/rates/electric-rates/183E%20-%20Large%20General%20Service.pdf>, 2014.
- [21] Lead-acid Battery. http://en.wikipedia.org/wiki/Lead%E2%80%93acid_battery, 2014.
- [22] Lithium-Ion Battery. Wikipedia, 2014.
- [23] Maxwell Ultracapacitor Uninterruptible Power Supply (UPS) Solutions. <http://bit.ly/1Brz01K>, 2014.
- [24] Summary of Rates for Public Service. <http://bit.ly/1eA6QaZ>, 2014.
- [25] PowerWall: Tesla Home Battery. <http://www.teslamotors.com/powerwall>, June 2015.
- [26] What is a Smart Building? <http://www.institutebe.com/smart-grid-smart-building/What-is-a-Smart-Building.aspx>, June 2015.
- [27] 24M. 24M. <http://24-m.com/>, 2015. [Online; accessed 25-June-2015].
- [28] Akay, Diyar, and Atak, Mehmet. Grey prediction with rolling mechanism for electricity demand forecasting of turkey. *Energy* (2007).
- [29] Bakker, V., Bosman, M., Molderink, A., Hurink, J., and Smit, G. Demand Side Load Management Using a Three Step Optimization Methodology. In *SmartGridComm* (October 2010).
- [30] Bar-Noy, A., Feng, Y., Johnson, M., and Liu, O. When to Reap and When to Sow: Lowering Peak Usage With Realistic Batteries. In *WEA* (June 2008).
- [31] Bar-Noy, A., Johnson, M., and Liu, O. Peak Shaving Through Resource Buffering. In *WAOA* (September 2008).
- [32] Bar-Noy, Amotz, Johnson, Matthew, and Liu, Ou. Peak Shaving Through Resource Buffering. In *WAOA* (September 2008).
- [33] Barker, S., Mishra, A., Irwin, D., Cecchet, E., Shenoy, P., and Albrecht, J. Smart*: An Open Data Set and Tools for Enabling Research in Sustainable Homes. In *SustKDD* (August 2012).
- [34] Barker, Sean, Mishra, Aditya, Irwin, David, Shenoy, Prashant, and Albrecht, Jeannie. Smartcap: Flattening peak electricity demand in smart homes. In *PerCom* (March 2012).

- [35] Bradbury, Kyle. Energy Storage Technology Review. http://people.duke.edu/~kjb17/tutorials/Energy_Storage_Technologies.pdf, August 2010.
- [36] Bringing Residential Real-Time Pricing to Scale in Illinois: Policy Recommendations, 2010.
- [37] Carpenter, T., Singla, S., Azimzadeh, P., and Keshav, S. The Impact of Electricity Pricing Schemes on Storage Adoption in Ontario. In *e-Energy* (May 2012).
- [38] Cauchon, D. Usatoday, household Electricity Bills Skyrocket. <http://www.usatoday.com/money/industries/energy/story/2011-12-13/electric-bills/51840042/1>, December 13th 2011.
- [39] Chang, C., and Lin, C. LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology 2*, 27 (April 2011).
- [40] Chen, Haisheng, Cong, Thang Ngoc, Yang, Wei, Tan, Chunqing, Li, Yongliang, and Ding, Yulong. Progress in electrical energy storage system: A critical review. *Progress in Natural Science* (2009).
- [41] CNT energy. Dynamic Pricing and Smart Grid. <http://www.cntenergy.org/pricing/>, 2011.
- [42] Controls, Johnson. Building management — johnson controls inc., 2015. [Online; accessed 23-June-2015].
- [43] Daryanian, B., Bohn, R., and Tabors, R. Optimal Demand-side Response to Electricity Spot Prices for Storage-type Customers. *TPS 4*, 3 (August 1989).
- [44] Database of State Incentives for Renewables and Efficiency. <http://www.dsireusa.org>, 2010.
- [45] Davidoff, S., Lee, M., Yiu, C., Zimmerman, J., and Dey, A. Principles of Smart Home Control. In *UbiComp* (September 2006).
- [46] de ven, P. Van, Hegde, N., Massoulie, L., and Salonidis, T. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *ENERGY* (May 2011).
- [47] DeNavas-Walt, C., Proctor, B., and Smith, J. Income, Poverty, and Health Insurance Coverage in the United States: 2010. Tech. rep., U.S. Census Bureau, September 2011.
- [48] EcoGrid EU FP7 project. <http://www.eu-ecogrid.net/>, 2012.
- [49] The Economist, Lead-acid Batteries Recharged: A 150-year-old Technology Looks to the Future. http://www.economist.com/node/12924031?story_id=12924031, January 14th 2009.

- [50] eGauge Energy Monitoring Solutions. <http://www.egauge.net/>, 2012.
- [51] EIA, U S. How much electricity is lost in transmission and distribution in the united states? - FAQ, 2014. [July-2014].
- [52] Electric, Schneider. Building management systems - schneider electric corporate, 2015. [Online; accessed 23-June-2015].
- [53] Energy, NEC. Grid Energy Storage; Battery Systems — NEC Energy Solutions, 2015. [Online; accessed 25-June-2015].
- [54] Faruqui, A., Hledik, R., and Tsoukalis, J. The Power of Dynamic Pricing. *Electricity Journal* 22, 3 (February 2009), 42–56.
- [55] State of the Markets Report 2008, August 2009.
- [56] Furukawa, J., Takada, T., Kanou, T., Monma, D., Lam, L., Haigh, N., Lim, O., and Louey, R. Development of UltraBattery. Tech. rep., CSIRO, December 2006.
- [57] Gantenbein, D., Binding, C., Jansen, B., Mishra, A., and Sundstrom, O. EcoGrid EU: An Efficient ICT Approach for a Sustainable Power System. In *SustainIT* (October 2012).
- [58] Ganu, T., Seetharam, D., Arya, V., Kunnath, R., Hazra, J., Husain, S., DeSilva, L., and Kalyanaraman, S. nPlug: A Smart Plug for Alleviating Peak Loads. In *e-Energy* (May 2012).
- [59] Govindan, S., Sivasubramaniam, A., and Urgaonkar, B. Benefits and Limitations of Tapping into Stored Energy for Datacenters. In *ISCA* (June 2011).
- [60] Hall, E. On a new action of the magnet on electric currents. *American Journal of Mathematics* 2, 3 (November 1879), 287–292.
- [61] Hall, M. Correlation-based Feature Selection for Discrete and Numeric Class Machine Learning. In *ICML* (June 2000).
- [62] Hammerstrom, D., Oliver, T., Melton, R., and Ambrosio, R. Standardization of a Hierarchical Transactive Control System. In *Grid Interop* (September 2009).
- [63] Hammerstrom, DJ, Ambrosio, R, Brous, J, Carlon, TA, Chassin, DP, DeSteese, JG, Guttromson, RT, Horst, GR, Järvegren, OM, Kajfasz, R, et al. Pacific Northwest Gridwise Testbed Demonstration Projects. *Part I. Olympic Peninsula Project* (2007).
- [64] Harle, R., and Hopper, A. The Potential of Location-Aware Power Management. In *UbiComp* (September 2008).
- [65] Horii, Brian, and Cutter, Eric. Energy Efficiency Avoided Costs 2011 Update, December 2011.

- [66] Insteon. Insteon hub insteon.com, 2015. [Online; accessed 23-June-2015].
- [67] Jansen, B., Binding, C., and Mishra, A. Input on the Real-Time Price distribution protocol for Ecogrid EU WP 3 Task 3.6. http://www.zurich.ibm.com/pdf/ecogrid/price_distribution_protocol_1.1.pdf, July 2012.
- [68] Kaplan, Stan. Power Plants: Characteristics and Costs. Tech. rep., Congressional Research Service Report to Congress, November 2008.
- [69] Keshav, S., and Rosenberg, C. Direct Adaptive Control of Electricity Demand. Tech. Rep. CS-2010-17, University of Waterloo, September 2010.
- [70] Keshav, S., and Rosenberg, C. How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid. In *GreenNet* (August 2010).
- [71] Knapp, Karl, Martin, Jennifer, Price, Snuller, and Gordon, Frederick. Costing Methodology for Electric Distribution System Planning, November 2000.
- [72] Koutsopoulos, I., Hatzi, V., and Tassioulas, L. Optimal Energy Storage Control Policies for the Smart Power Grid. In *SmartGridComm* (September 2011).
- [73] Laboratory, Lawrence Livermore National. U.S. Energy Flowchart 2013. <https://flowcharts.llnl.gov/>, August 2014.
- [74] LG. Smart appliances: Discover lg’s smart thing appliances — lg usa, 2015. [Online; accessed 23-June-2015].
- [75] Linden, D. *Linden’s Handbook of Batteries*, Fourth ed. McGraw Hill, 2011, pp. 17.1–17.39.
- [76] McLaughlin, S., McDaniel, P., and Aiello, W. Protecting Consumer Privacy from Electric Load Monitoring. In *CCS* (October 2011).
- [77] Mishra, A., Binding, C., Gantenbein, D., Jansen, B., and Sundström, O. Initial Inputs to Ecogrid Task 3.5 on DER and Aggregator Communication and Standards Contributions. http://www.zurich.ibm.com/pdf/ecogrid/SOTA_price-protocols.pdf, July 2012.
- [78] Mishra, A., Irwin, D., Shenoy, P., Kurose, J., and Zhu, T. SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *e-Energy* (May 2012).
- [79] Mishra, A., Irwin, D., Shenoy, P., and Zhu, T. Scaling Distributed Energy Storage for Grid Peak Reduction. In *e-Energy* (May 2013).
- [80] Mishra, A., Sitaraman, R., Irwin, D., Zhu, T., Shenoy, P., Dalvi, B., and Lee, S. Integrating energy storage in electricity distribution networks. Tech. Rep. UM-CS-2015-012, UMass Computer Science Tech Report, May 2015.

- [81] Molina-Markham, A., Shenoy, P., Fu, K., Cecchet, E., and Irwin, D. Private Memoirs of a Smart Meter. In *BuildSys* (November 2010).
- [82] Motors, Tesla. Tesla Powerwall, 2015. [Online; accessed 25-June-2015].
- [83] National Public Radio. How Smart is the Smart Grid?: Interview with Dan Delurey, President Demand Response Smart Grid Coalition, July 7th, 2010.
- [84] Nest. Meet the nest thermostat — nest, 2015. [Online; accessed 23-June-2015].
- [85] New York ISO. Understanding the Markets. http://www.nyiso.com/public/about_nyiso/understanding_the_markets/energy_market/index.jsp, 2012.
- [86] New england ISO. <http://www.iso-ne.com/>, 2012.
- [87] Ontario Energy Board: Electricity Prices. <http://www.ontarioenergyboard.ca/OEB/Consumers/Electricity/Electricity+Prices>, 2012.
- [88] Pelley, Steven, Meisner, David, Zandevakili, Pooya, Wenisch, Thomas F, and Underwood, Jack. Power Routing: Dynamic Power Provisioning in the Data Center. In *ACM Sigplan Notices* (2010).
- [89] Pentland, William. Perverse Economics of the Electric Grid: As Generation Gets Cheaper, Transmission Costs Soar. <http://onforb.es/1HJMsAP>, 2013.
- [90] Power, Sparkplug. Sparkplug Power — Bridging The Gap, 2015. [Online; accessed 25-June-2015].
- [91] Roozbehani, M., Dahleh, M. A., and Mitter, S. K. On the Stability of Wholesale Electricity Markets under Real-Time Pricing. In *CDC* (December 2010).
- [92] Roozbehani, M., Rinehart, M., and Dahleh, M. A. Real-Time Pricing with Ex-post Adjustments in Competitive Electricity Markets: Volatility and Efficiency Analysis. In *Allerton* (September 2011).
- [93] Schoenung, Susan. Energy Storage Systems Cost Update: A Study for the DOE Energy Storage Systems Program. Tech. rep., Sandia National Laboratories, April 2011.
- [94] Schoenung, Susan M, and Hassenzahl, William V. Long-vs. Short-Term Energy Storage Technologies Analysis. A Life-Cycle Cost Study. A Study for the DOE Energy Storage Systems Program.
- [95] Schülke, A., Bauknecht, J., and Häußler, J. Power Demand Shifting with Smart Consumers: A Platform for Power Grid friendly Consumption Control Strategies. In *SmartGridComm* (October 2010).
- [96] SERC: State Environmental Resource Center. <http://www.serconline.org/netmetering/stateactivity.html>, 2012.

- [97] Sharma, N., Gummesson, J., Irwin, D., and Shenoy, P. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *SECON* (June 2010).
- [98] Sharma, N., Sharma, P., Irwin, D., and Shenoy, P. Predicting Solar Generation from Weather Forecasts Using Machine Learning. In *SmartGridComm* (October 2011).
- [99] Sickinger, Ted. PGE to Test Peak-pricing for Electricity. *The Oregonian* (September 22, 2009).
- [100] Siemens. Building Management System - Building Technologies - Siemens, 2015. [Online; accessed 23-June-2015].
- [101] Smart Grid Information Clearinghouse: Legislation and Regulation. <http://www.sgiclearinghouse.org/Legislation?q=node/1705>, 2011.
- [102] Sonneborn, L., and Vleck, F. Van. The Bang-Bang Principle for Linear Control Systems. *SIAM Control Series A*, 2 (1965).
- [103] Srikantha, Pirathayini, Rosenberg, Catherine, and Keshav, Srinivasan. An analysis of peak demand reductions due to elasticity of domestic appliances. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet* (2012), ACM, p. 28.
- [104] State of California Executive Order S-21-09. <http://gov.ca.gov/executive-order/13269>, 2009.
- [105] Technical Manual for Sun Xtender Batteries. http://www.sunxtender.com/pdfs/Sun_Xtender_Battery_Technical_Manual.pdf, July 2011.
- [106] Taneja, J., Culler, D., and Dutta, P. Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration. In *SmartGridComm* (October 2010).
- [107] Tso, Geoffrey KF, and Yau, Kelvin KW. Predicting Electricity Energy Consumption: A Comparison of Regression Analysis, Decision Tree and Neural Networks. *Energy* (2007).
- [108] Electric Power Annual 2010, November 2011.
- [109] Uргаonkar, R., Uргаonkar, B., Neely, M., and Sivasubramaniam, A. Optimal Power Cost Management Using Stored Energy in Data Centers. In *SIGMET-RICS* (March 2011).
- [110] Vytelingum, P., Voice, T., Ramchurn, S., Rogers, A., and Jennings, N. Agent-based Micro-storage Management for the Smart Grid. In *AAMAS* (May 2010).

- [111] Wang, Di, Ren, Chuangang, Sivasubramaniam, Anand, Urgaonkar, Bhuvan, and Fathy, Hosam. Energy Storage in Datacenters: What, Where, and How Much? *SIGMETRICS* (2012).
- [112] Wikipedia. Frequency scaling — wikipedia, the free encyclopedia, 2013. [Online; accessed 23-June-2015].
- [113] Wikipedia. Building management system — wikipedia, the free encyclopedia, 2015. [Online; accessed 23-June-2015].
- [114] Wikipedia. Dynamic voltage scaling — wikipedia, the free encyclopedia, 2015. [Online; accessed 23-June-2015].
- [115] Wikipedia. Grid energy storage — wikipedia, the free encyclopedia, 2015. [Online; accessed 25-June-2015].
- [116] Wikipedia. Internet of things — wikipedia, the free encyclopedia, 2015. [Online; accessed 23-June-2015].
- [117] Yang, W., Li, N., Qi, Y., Qardaji, W., McLaughlin, S., and McDaniel, P. Minimizing Private Data Disclosures in the Smart Grid. In *CCS* (October 2012).
- [118] Zhu, T., Mishra, A., Irwin, D., Sharma, N., Shenoy, P., and Towsley, D. The Case for Efficient Renewable Energy Management in Smart Homes. In *BuildSys* (November 2011).
- [119] Zivan. Battery Discharge at Constant Current: Technical Features and User Manual. <http://www.zivanusa.com/pdf/SBM%20Gb.pdf>, 2012.