

Residential Task Scheduling Under Dynamic Pricing Using the Multiple Knapsack Method

Kumaraguruparan N., Sivaramakrishnan H., and Sachin S. Sapatnekar
ECE Department, University of Minnesota–Twin Cities, Minneapolis, MN 55455

Abstract— A key component of the smart grid is the ability to enable dynamic residential pricing to incentivize the customer and the overall community to utilize energy more uniformly. However, the complications involved require that automated strategies be provided to the customer to achieve this goal. This paper presents a solution to the problem of optimally scheduling a set of residential appliances under day-ahead variable peak pricing in order to minimize the customer’s energy bill (and also, simultaneously spread out energy usage). We map the problem to a well known problem in computer science – the multiple knapsack problem – which enables cheap and efficient solutions to the scheduling problem. Results show that this method is effective in meeting its goals.

Index Terms—Cost Function, Dynamic Scheduling, Multiple Knapsack Problem (MKP), Optimization, Smart Grids.

I. INTRODUCTION

With ever-increasing and imbalanced electrical demands, there is a clear realization that further efficiencies can and must be achieved by embedding intelligence into the power grid. Since 1982, the growth in peak demand has exceeded transmission growth by almost 25% every year [1]. This growing demand for electricity and the steep increase in the power consumption of a household arising as a result of various technological improvements, pose a stiff challenge to the future electricity distribution systems. The smart grid has emerged as a solution that provides benefits to all parties involved: to the utilities because it provides better demand management, to the customer because of its potential for economy, and to the environment because of the ability to reduce greenhouse gas emissions.

The European Technology Platform Smart Grids (ETP Smart Grids) defines smart grids as “electricity networks that can intelligently integrate the behavior and actions of all users connected to it – generators, consumers and those that do both – in order to efficiently deliver sustainable, economic and secure electricity supplies” [2]. The smartness in the design of the infrastructure of a smart grid is attributed to allowing the users to play a significantly larger role in optimizing and reducing their power consumption thereby increasing the efficiency of the system. Estimates show that smart electricity grids should reduce CO₂ emissions in the EU by 9% and the annual household energy consumption by 10% [2].

There are numerous technical issues involved in enabling smart grid infrastructure, and this paper will focus on the problem of managing and scheduling residential power demand under a dynamic pricing system.

The idea of dynamic pricing is a key component of the residential smart grid, where the utility creates a time-varying rate structure over time. The user is given more responsibility wherein s/he determines and controls when s/he needs a particular appliance in his household. Moreover, as locally-generated power at the residential level proliferates (e.g., through home solar panels), the residential user is provided with an opportunity to supply excess electricity back to the grid, with monetary compensation. This local generation and use of power through renewable sources of energy helps reduce the overall carbon footprint and greenhouse gas emissions.

Strategies for reducing the cost of residential power of include reducing the power consumption or shifting the temporal profile of power consumption to consumption at cheaper times, where possible [3]. Within the framework of this paper, we do not aim to reduce the power consumption of a household, but instead study how the power consumption of a household can be shifted to off-peak hours so that the user benefit financially, and the load on the system is better balanced through the day. This behavior is incentivized through pricing strategies: incentives are provided to the users who are willing to move in their appliances to off-peak hours through dynamic pricing. Shifting consumption in this way is a powerful strategy: the Pacific Northwest National Laboratory states that existing U.S. power plants could meet the electricity needs of 73% of the nation’s light vehicles (i.e., cars and small trucks) if the vehicles were replaced by plug-ins that recharged at night. Such a shift would reduce oil consumption by 6.2 million barrels per day [1].

There are several variants of dynamic pricing [16]:

- *Time of use* pricing establishes a variable price structure for peak, shoulder, and off-peak hours. These prices are typically established well in advance, e.g., a year ahead.
- *Real-time pricing* allows the prices to change on an hourly basis, based on market demand.
- *Variable peak pricing* is a hybrid of the two, and establishes variable pricing in the day, defined a day ahead.

Various pilot projects have tested the notion of dynamic pricing based on real-time pricing and variable peak pricing, and it has been concluded that customers show exhibit price responses within this environment that provides mutual benefit to all [17].

We focus, in particular, on the model of day-ahead variable peak pricing. Under this model, the benefit to the smart grid

depends on how effectively it is able to shift the appliances to the off-peak hours. In the context of near-future power consumption patterns, smart scheduling approaches are becoming particularly significant. For example, the advent of plug-in hybrid electric vehicles (PHEVs) is anticipated to place a significant and large load on the grid; however, in many cases, customers will need the vehicle to be charged during some period between a customer-specified start and end time (e.g., between 9pm and 9am). Smart scheduling algorithms can reduce the impact on the global grid by distributing this charging period and/or charging the vehicle during noncontiguous time slots.

Several algorithms and techniques have been proposed in the context of scheduling under dynamic residential pricing. Pedrasa *et al.* [4] use a modified form of coevolutionary particle swarm optimization (CPSO) with stochastic repulsion amongst particles, referred to as CPSO-R, for their scheduler. The scheduler determines when various devices in the household need to co-ordinate with each other to increase the benefits (during the peak-load hours) and when they can operate independently (at times when coordinating with each other does not add significant value). However, stochastic optimization techniques can be computationally expensive. If performed locally within the home, they may require significant computational resources; if performed offline, the combined computation associated with a large number of homes may be very significant.

Mohsenian-Rad *et al.* [5] present a game-theoretic approach to solving this problem. As outlined earlier, smart scheduling has the potential to benefit all players. Through their analysis, referred to as energy consumption scheduling (ECS), devices work with current pricing policies to choose a schedule, while the prices are set by the cumulative demand at any time of the day. The consumers within the community aim to reach an equilibrium condition, the Nash Equilibrium, which provides the globally optimal solution to the scheduling problem, while providing considerable cost savings to the community as a whole. However, this work does not consider dynamic effects, such as rates of response for the pricing structure and for consumption. If the pricing structure changes too rapidly, the settling time for achieving the Nash equilibrium may be an issue; if it changes too slowly, it may not predict the demand on a specific day.

In other work, Hatami *et al.* [6] introduce a quasi-dynamic pricing function and build a scheduler that can work with interruptible tasks, i.e., tasks that need not be run in contiguous intervals of time. Sianaki *et al.* [7] introduce a single knapsack problem based approach to achieve energy efficient consumption in smart grids by mapping the scheduling problem to a standard dynamic programming model, the knapsack problem. The authors take into account the preference of the consumers within a community before determining the schedule. The appliances are prioritized by the respective consumers and this priority function of an appliance plays an important role when scheduling during peak-load hours. However, under this framework, appliances are chosen

based on a priority ranking, and completion times for tasks cannot be trivially captured or guaranteed.

In this paper, we develop a formulation that can be used to schedule appliances using the multiple knapsack problem (MKP) approach. Our schedule has several important features:

- We consider the effect of must-run services (such as lights, refrigerators, etc.) that must always be on and cannot be scheduled, as well as tasks that can be flexibly scheduled in time.
- We incorporate the effects of local power generation sources in the objective function
- We consider the case where task deadlines must be imposed and permit tasks to occupy noncontiguous time slots.

The rest of the paper is organized as follows: Section II formally presents the problem statement, the definition of the terminology used in the paper, and shows the mapping from the scheduling problem to an instance of the MKP problem. The formulation of the deterministic model of our problem is explained in section III. The experimental results of are discussed in section IV, and the paper is then concluded in section V.

II. PRELIMINARIES

In this section, we first formally state the parameters of the residential scheduling problem, followed by a definition of the terminology used within the paper. We then consider the multiple knapsack problem (MKP) and discuss the mapping between our scheduling problem and an instance of the MKP.

A. Problem Statement

This section describes the problem statement, stating the requirements on the design of the scheduler.

The consumers within a smart grid community are charged depending upon the net power consumed from the smart grid based on a variable peak pricing model. The energy prices for a specific period of time is set in advance (e.g., a day ahead). Consumers with smart appliances may define allowable periods when their appliances are to be used. Models for the specifications include:

- *Must-run schedules* for appliances that must be operational at all times, such as lights, refrigerators, and winter heating in cold climates.
- *Fixed schedules* that apply to appliances that must run at specific fixed times
- *Flexible schedules* that apply to appliances that must run for a specified amount of time, anytime within a flexible interval.

The goal of the approach is to optimally schedule the operating times of all appliances so that they meet the constraints imposed by the list above, while minimizing the electricity bill for the household. The household may contain a mix of smart appliances that can be scheduled and appliances without inbuilt intelligence, which are manually operated. For the purposes of this work, we consider only the scheduling of smart appliances. The schedules for manually

operated appliances, if available, can be provided as limited time fixed must-run applications.

Thus the goal of the scheduler is to schedule the appliances within the household to minimize the electricity bill, while contending with constraints such as the preferences of the user, the price model, and the requirement that all the appliances must be scheduled within the given time frame.


B. Key terms and Notation

Cost function: The cost function $Cost(t)$ represents the cost of consumption of a certain amount of energy at a time instant t . It is assumed that variable peak pricing is followed where the cost function varies with the time in the day. Generally the cost of power consumption is high during the peak-hours (evenings and mornings) and is considerably low during periods of low power usage (afternoons and late nights); exceptions to these scenarios correspond to examples such as heat waves, where energy consumptions in the afternoon may be large. Day-ahead pricing, combined with weather forecasts, provides utilities with the tools to vary these prices in a quasi-dynamic way.

Practically speaking, the cost function is implemented as a discrete piecewise-constant cost function, $Cost_{ij}$, as illustrated in Figure 1. The per-unit price is altered to discrete values at discrete time steps, e.g., at the top of every hour. In our exposition, we assume that the time slots are of equal length, but it is rather simple to extend this approach to the case where the time slot lengths are not the same.

FIGURE 1: AN EXAMPLE COST FUNCTION, CORRESPONDING TO THE COST DATA THAT WILL BE SHOWN IN TABLE II.

Locally-generated power: This term may include any renewable sources of energy such as wind energy or solar energy, that is generated on-site. These sources of energy are not deterministic but may be reflected by their expected values. Again, weather forecasts may be used to predict these



values on a day-ahead basis.

C. The Multiple Knapsack Problem

In this paper, we map the scheduling problem to a multiple knapsack problem. We begin by providing a brief overview of knapsack problems.

The (single) knapsack problem is a combinatorial problem where a number of objects, each associated with a *value* and a *weight*, must be packed into a knapsack of a specific capacity, such that the value of the objects within the knapsack is

maximized. Single knapsack problems are a standard problem template in computer science, and numerous practical problems have been mapped on to this problem. The fractional version of the knapsack problem can be solved using dynamic programming techniques in polynomial time. The integer problem is NP-complete, but a number of efficient polynomial-time approximation algorithms are available to provide near-optimal solutions [8].

The multiple knapsack problem (MKP) [8–12] is a generalization of the single knapsack problem. In essence, the MKP is a resource allocation problem wherein there is a set of m resources (i.e., m knapsacks) and a set of n objects. As in the single knapsack problem, each object i in this set of n objects has two important attributes associated with it: the value of the object and its weight. Every resource (knapsack) j has a capacity constraint C_j , which represents the maximum weight that the resource can support. The objective of the MKP is to find a subset of the objects that can be packed within the bins such that the net value of all objects within all knapsacks is maximized. Although the integer MKP problem is NP-complete, in practice optimal or near-optimal solutions are available in polynomial time or for very reasonable runtimes [9,14].

D. Mapping the Scheduling Problem to the MKP

The mapping between our scheduling problem and the MKP is now outlined:

- The m knapsacks correspond to m time intervals where the energy price is fixed.
- The n appliances correspond to n objects that must be packed.
- The weight of each object is the energy E_i consumed by appliances A_i in each time slot. Note that E_i is independent of t .
- The value of the object in a specific time slot is the cost of power consumption of the appliance in that time slot, defined by the $Cost_{ij}$ function described in Section IIB.
- The bin capacity, or the capacity of the knapsack j , $capacity_j$, is the maximum energy that can be drawn from the grid in any time slot j . For the customer, this can help ensure that the electricity bill can be controlled, and for the utility, enforcing this limit ensures that the grid is not overstressed; in the absence of this, the customer may well always choose the least expensive time slot! For our purposes, we consider this limit on a per-household basis. A variation of the problem formulation applies this on a per-community basis, but this requires simultaneous scheduling of all devices in all households in a community, which is onerous and probably not worthwhile. Such constraints can be enforced at other points of the smart grid.

III. MODELING THE OBJECTIVE FUNCTION AND CONSTRAINTS

The description of the problem formulation is divided into two subphases: initially we describe a formulation that schedules only the smart appliances, as specified by the consumer. Next, we extend to formulation to account for must-run services and

locally-generated power.

A. Initial Scheduling

Let us consider n appliances within a smart grid that has to be scheduled in m time slots. Recall that the cost function, $Cost(t)$, is a piecewise constant function that reflects the cost of power consumption at a time slot t , and that $capacity(t)$ is the maximum power capacity in time slot t . Let $T_{req,i}$ be the time that is required by an appliance A_i to complete its task: this is specified in units of the number of time intervals. This value of $T_{req,i}$ will be the same across all time slots for an appliance i since the power consumed by the appliance is a constant.

We now present the optimization formulation as a discrete linear program that has the form of the MKP. We introduce the Boolean integer variable, $X_i(t)$, defined as

$$X_{ij} = \begin{cases} 1 & \text{if the appliance is on in time slot } j \\ 0 & \text{if the appliance is off in time slot } j \end{cases} \quad (1)$$

The objective function can now be formulated as

$$\min \sum_{i=1}^n E_i \left(\sum_{j=1}^m X_{ij} \cdot Cost_j \right) \quad (2)$$

Here, X_{ij} is an indicator variable that states whether the $Cost_j$ term in slot j should contribute to the cost function or not. If so, this term, which represents the unit cost in that slot, is multiplied by the energy usage during that period.

There are several constraints that must be obeyed. Specifically:

Completion constraints are designed to ensure that each task is completed during the time period in question (e.g., one day). For each appliance i , these can be formulated as:

$$\sum_{j=1}^m X_{ij} = T_{req,i} \quad (3)$$

which states that an appliance i must be provided with enough time slots to complete its function. Clearly, this constraint support noncontiguous time slots for various appliances as long as the total $T_{req,i}$ constraint is met.

Capacity constraints may be used to control the total energy usage of a home and help incentivize energy usage that is more distributed in time and uniform. For each time slot j , these are stated as

$$\sum_{i=1}^n E_i X_{ij} \leq capacity_j \quad (4)$$

B. Incorporating Must-run Services and Locally-Generated Power

The formulation is now extended to include the power consumed by the must-run services and locally-generated power. Let us consider E_{mj} to be the power that is consumed by the must-run services at a time slot j . As a result, the energy consumption in slot j must be incremented by E_{mj} for all slots j . The effects of this change are:

- The total cost of energy goes up by $E_{mj} Cost_j$ which should affect the objective function; however, since this

adds a constant quantity to the objective function, the function to be minimized is unchanged from (2).

- The capacity constraint is altered since an amount E_{mj} of the capacity is utilized in time slot j by must-run services. This implies that the capacity constraint changes to:

$$\sum_{i=1}^n E_i X_{ij} \leq capacity_j - E_{mj} \quad (5)$$

Clearly, for a meaningful solution, it is necessary to ensure that

$$capacity_j - E_{mj} \geq 0$$

i.e., that the must-run services use energy that is under the capacity at any time slot j .

The role of locally-generated power is dual to the role of must-run services. If P_{gj} represents the power that is generated by the local power source at time slot j , then this amount subtracts from the total cost of energy, and using the same argument as above, the objective function is left unchanged.

An interesting case is the situation where the value of the locally-generated power, P_{mj} , exceeds the demand (equation (2)) in slot j . In this case, the power is sold back to the utility, and the goal of minimizing the objective function still achieves the goal of maximizing the dollar gains to the homeowner of selling electricity back to the grid.

In addition to the above, locally-generated power alters the capacity constraint, which now becomes

$$\sum_{i=1}^n E_i X_{ij} \leq capacity_j - E_{mj} + P_{gj} \quad (6)$$

The overall optimization problem can be stated using the objective in (2), and constraints corresponding to (1), (3), and (6).




FIGURE 2: RESULTS FOR THE SIMPLE EXAMPLE

IV. AN EXAMPLE

We test our formulation initially for a small case study with a small number of appliances and time grids and then enlarge the case study to schedule a larger set of appliances within an entire day. The specifics of the small example are as follows:

Appliances: Two appliances, A_1, A_2

Time discretization: 4 time slots, T_1, T_2, T_3, T_4

Energy usage: $E_1 = 10$ units, $E_2 = 20$ units

Task time requirements: $T_{req,1} = 3, T_{req,2} = 2$

Capacity of each time slot: $capacity_1 = capacity_2 = 30$ units

Cost during each time slot:

$$\begin{aligned} Cost_1 &= 10 \\ Cost_2 &= 20 \\ Cost_3 &= 30 \\ Cost_4 &= 20 \end{aligned}$$

For simplicity, we set

Must-run usage = Locally-generated power = 0 in all slots

As seen earlier, these parameters simply affect the right-hand side of the capacity constraints, and can be considered as alterations to the capacity.

The optimization problem can then be formulated as:

$$\begin{aligned} \min E_1(X_{11} \cdot Cost_1 + X_{12} \cdot Cost_2 + X_{13} \cdot Cost_3 + X_{14} \cdot Cost_4) + \\ E_2(X_{21} \cdot Cost_1 + X_{22} \cdot Cost_2 + X_{23} \cdot Cost_3 + X_{24} \cdot Cost_4) \end{aligned}$$

subject to:

$$\begin{aligned} E_1X_{11} + E_2X_{12} + E_3X_{13} + E_4X_{14} &\leq capacity_1 - E_{m1} + E_{g1} \\ E_1X_{21} + E_2X_{22} + E_3X_{23} + E_4X_{24} &\leq capacity_2 - E_{m2} + E_{g2} \\ X_{11} + X_{12} + X_{13} + X_{14} &= T_{req,1} \\ X_{21} + X_{22} + X_{23} + X_{24} &= T_{req,2} \\ X_{21}, X_{22}, X_{23}, X_{24} &\in \{0,1\} \end{aligned}$$

The optimal solution for this problem is shown in Figure 2: appliance A_1 operates in slots 1, 2, and 4, while appliance A_2 operates in slots 1 and 2. Clearly, it meets the requirements $T_{req,1}$ and $T_{req,2}$ as well as the capacity constraints. It can be seen that the schedule evades any consumption in the most expensive slot, slot 3, and fully uses the least expensive slot, slot 1, and that the tasks are scheduled in interruptible fashions.

V. SIMULATION AND RESULTS

The optimization problem formulated above was solved using the What'sBest solver [14]. What'sBest is an add-on to Microsoft Excel and allows building large linear, non-linear, integer, or stochastic optimization models within a spreadsheet.

	A_1	A_2	A_3	A_4	A_5	A_6	A_7	A_8	A_9	A_{10}	A_{11}	A_{12}	A_{13}	A_{14}	A_{15}	A_{16}	A_{17}	A_{18}
$T_{req,i}$	3	2	1	2	3	2	1	2	3	1	2	3	3	1	2	3	1	2
E_i	100	200	150	250	50	100	200	150	250	50	200	300	100	300	200	150	200	150

TABLE I: APPLIANCE PARAMETERS

Time slot j	$Cost_j$	$capacity_j$	E_{mj}	E_{gj}
1	100	500	50	20
2	200	500	0	10
3	300	500	50	20
4	200	500	0	0
5	100	500	100	0
6	400	500	50	0
7	200	500	50	0
8	300	500	0	0
9	100	500	50	0
10	200	500	0	30
11	200	500	0	30
12	300	500	0	30
13	200	500	0	20
14	200	500	50	20
15	200	500	50	10
16	300	500	0	10

TABLE II: DESCRIPTION OF THE TIME SLOTS

We apply our solution to a testcase with 18 appliances and 16 time slots. The parameters for the appliances are described in Table I, and the characteristics of the time slots are described in Table II. The capacity constraint at each of the 16 time slots is uniform and is set to 500 units.

The results of the optimization are shown in Figure 3. The figure at left is a stacked bar chart, where the x-axis lists the time slots from 1 to 16, and the y-axis lists the energy usage. The small bars below the x-axis correspond to locally-generated power. The energy usage of all appliances is added up to create the height of the stack: it is easily seen that the capacity constraint of 500 units is met at each time slot, and that electricity usage is uniform, except for a few time slots. This achieves the goal of forcing the customer to make energy usage more uniform across the day.




FIGURE 3: RESULTS FROM AN 18-APPLIANCE, 16-TIME-SLOT PROBLEM. THE STACKED BAR CHART AT LEFT SHOWS THE ENERGY USAGE OF THE APPLIANCES, AND THE FIGURE AT RIGHT SHOWS THE COST PER UNIT (BLUE) AND THE TOTAL USAGE PER TIME SLOT.

We explore the low-usage time slots in the right part of Figure 3. The x-axis represents time and shows the 16 time slots. The two plots show the cost (blue) and the total energy usage (red) in each time slot. We focus on the four time slots (#3, #6, #12, and #16) with nonmaximal usage: it can be seen from the blue bars and from Table II that these correspond to the highest $Cost_t$ values. Since the goal of the optimization is to minimize customer cost, these slots are left unused or underused.

VI. CONCLUSION

We have presented an approach to enable residential energy scheduling under a dynamic pricing paradigm. Our experiments show that the method is effective in encouraging uniformity in the energy profile in the home.

Further constraints can easily be built into this system: although these have not been shown here. One example is when a customer wishes to constrain the time range in which a specific task is scheduled by specifying a preference. This is quite simple to incorporate into the current system by reducing the number of X_{ij} variables: if an appliance is not to be scheduled in a specified slot, the corresponding Boolean variable is set to zero (and is therefore not a variable). This, in fact, makes the optimization simpler since it reduces the number of variables. Another example: we have enabled scheduling with interruptible tasks and sufficiently large time slots, but a customer may wish for some appliances to be scheduled in contiguous time slots (e.g., laundry tasks). In this connection, for finer-grained scheduling, fractional time slots with contiguous scheduling may be desirable. We expect to include this in future work

REFERENCES

- [1] U.S. Department of Energy, The Smart Grid: An Introduction, 2009.
- [2] "What is a smart grid?" SmartGrids, European Technology platform for the Electricity Networks of the Future. <http://www.smartgrids.eu/>
- [3] Energy Conservation Committee Report and Recommendations, *Reducing Electricity Consumption in Houses*. Ontario Home Builders' Association, May 2006.
- [4] M. Pedrasa, T. Spooner, and I. MacGill, "Coordinated scheduling of residential distributed energy resources to optimize 'smart' home energy services," *IEEE Transactions on Smart Grids*, Vol 1, p134-143, 2010.
- [5] A. H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich and R. Schober "Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid," *IEEE PES Conf. Innovative Smart Grid Technology*, 2010.
- [6] S. Hatami and M. Pedram, "Minimizing the Electricity Bill of Cooperative Users under a Quasi-Dynamic Pricing Model," in *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010, pp. 421–426.
- [7] O. A. Sianaki, O. Hussain, A.R Tabesh "A knapsack problem approach for achieving efficient energy consumption in smart grid for end-users' life style," *IEEE Conference on Innovative Technologies for an Efficient and Reliable Electricity Supply (CITRES)*, 2010, pp. 159–164.
- [8] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*, Springer, 2004.
- [9] C. Chekuri and S. Khanna "A PTAS for the multiple knapsack problem," *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2000, p213.
- [10] S. Fidanova, "Evolutionary algorithm for multidimensional knapsack problem", *PPSNVII-Workshop*, 2002.
- [11] F. Diedrich and K. Jansen, "Improved approximation algorithms for scheduling with fixed jobs". *Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms*, p675, 2009.
- [12] F. Diedrich, K. Jansen, F. Pascual, and D. Trystram. "Approximation algorithms for scheduling with reservations," in S. Aluru, M. Parashar, R. Badrinath, and K. Prasanna, editors, *HiPC*, volume 4873 of *LNCS*, pages 297-307, Springer, 2007.
- [13] LINDO API 6.0 User Manual, LINDO Systems Inc.
- [14] "What'sBest" LINDO Systems Inc., http://www.lindo.com/index.php?option=com_content&view=article&id=3&Itemid=11
- [15] M. Grant, S. Boyd and Y. Ye *cvx: Matlab Software for Disciplined Convex Programming*, <http://www.stanford.edu/~boyd/cvx/>
- [16] T. Flaim, R. Levy, and C. Goldman, "Dynamic Pricing in a Smart Grid World," NARUC webinar, 2010, <http://www.naruc.org/FERC/LBNL-Webinar3-Dynamic%20Pricing%20in%20a%20Smart%20Grid%20World.pdf>
- [17] B. Neenan and J. Eom, "Price Elasticity of Demand for Electricity: A Primer and Synthesis." *Electric Power Research Institute*, Palo Alto, CA, 2007. Available at: www.epri.com