Factored Models for Multiscale Decision-Making in Smart Grid Customers

Prashant P. Reddy
Machine Learning Department
Carnegie Mellon University
Pittsburgh, USA
ppr@cs.cmu.edu

Manuela M. Veloso
Computer Science Department
Carnegie Mellon University
Pittsburgh, USA
mmv@cs.cmu.edu

Abstract

Active participation of customers in the management of demand, and renewable energy supply, is a critical goal of the Smart Grid vision. However, this is a complex problem with numerous scenarios that are difficult to test in field projects. Rich and scalable simulations are required to develop effective strategies and policies that elicit desirable behavior from customers. We present a versatile agent-based *factored model* that enables rich simulation scenarios across distinct customer types and varying agent granularity. We formally characterize the decisions to be made by Smart Grid customers as a multiscale decision-making problem and show how our factored model representation handles several temporal and contextual decisions by introducing a novel *utility optimizing agent*. We further contribute innovative algorithms for (i) statistical learning-based hierarchical Bayesian timeseries simulation, and (ii) adaptive capacity control using decision-theoretic approximation of multiattribute utility functions over multiple agents. Prominent among the approaches being studied to achieve active customer participation is one based on offering customers financial incentives through variable-price tariffs; we also contribute an effective solution to the problem of *customer herding* under such tariffs. We support our contributions with experimental results from simulations based on real-world data on an open Smart Grid simulation platform.

Introduction

Demand side management (DSM) has been an important focus area for Smart Grid research over the past decade (Strbac 2008). Smart Grid customers are steadily acquiring distributed renewable generation capabilities; the promises and challenges of this evolution have increased the urgency of progress in DSM-related research (Amin and Wollenberg 2005) (Gomes 2009). However, achieving active participation from customers in the management of their electricity demand and supply is a complex problem with numerous scenarios that are difficult to test in field projects, *e.g.*, (Borenstein 2002). Rich and scalable simulations are required to develop effective strategies and policies that elicit desirable behavior from customers.

The emergent behavior observed in simulations is key to understanding the scenarios that need to be considered in developing Smart Grid technology. The granularity at which simulations are conducted and the realism with which the components of the simulations are represented directly influence the types of lessons that can be learned. For example, simulations of the nation-wide grid may be necessary to stress test long-haul transmission line capacities whereas simulation of a single *smart home* may yield lessons on appliance load coordination. Customers can be modeled simply as consumers, producers and storage facilities or more explicitly as suburban homes, office complexes, solar farms, electric vehicles and so on. Along another dimension, residential customers can be represented at the granularity of a single household with each appliance modeled separately, at the aggregate behavior of the entire household, or over collections of many households. The effort required to model this diversity of options is a critical challenge.

We present a versatile agent-based *factored model* that enables rich simulation scenarios across distinct customer types and varying agent granularity by leveraging a generic customer representation that can be suitably parameterized. We formulate the decisions to be made by each Smart Grid customer as a multiscale decision-making problem along temporal and contextual dimensions. We introduce a *utility optimizing agent* as a component that manages the multiscale decisions and briefly describe how it can be embodied differently for various real world deployment scenarios. We further contribute a mechanism that learns how to simulate demand or supply capacity profiles based on prior observed timeseries samples using hierarchical Bayesian timeseries models. We also contribute an algorithm for adaptive capacity control using decision-theoretic approximation of multiattribute utility functions over multiple agents.

Prominent among the approaches being studied to achieve active customer participation in DSM is one based on offering customers financial incentives through variable-price tariffs. Some studies have shown that customer price-response under such tariffs leads to detrimental peak-shifting behavior. Our adaptive capacity control algorithm contributes an effective stochastic solution to this *customer herding* problem. We support our contributions with experimental results from simulations based on real-world data on Power TAC, a large open-source Smart Grid simulation platform.

Related Work

The need to model power grid dynamics using software-based simulation has long been recognized in power systems research, *e.g.*, (Allen et al. 2001). (Sun and Tesfatsion 2007) extend such simulations to include economic considerations. (Karnouskos and de Holanda 2009) illustrate simulation focused on retail customers instead of wholesale markets. The Power TAC simulation platform (Ketter et al. 2011) (Chrysopoulos and Symeonidis 2009) furthers this trend by including competitive benchmarking; we use this platform to evaluate our models although they are conceptually transferable to other equivalent environments.

(Gottwalt et al. 2011) present a customer model that simulates load profiles for households equipped with smart appliances under variable-price tariffs; their methodology extends that of (Armstrong et al. 2009) to also consider real-time pricing (RTP). They find that a considerable amount of load is available for shifting. We build upon this work in our learning-based Bayesian timeseries simulation methodology. (Guo et al. 2008) provide similar fine-grained models for household capacity adaptation using models of occupant comfort. (Kolter and Ferreira 2011) provide a coarse-grained model based on real utility data to predict energy usage of buildings using geographical survey features.

(Paatero and Lund 2006) artificially generate domestic consumption timeseries and use a simple DSM simulation to estimate the effects on peak load from shifting; they do not consider time-based pricing. (Barbose, Goldman, and Neenan 2005) and (Hammerstrom 2008) describe large field projects that study the impact of time-based pricing amongst various customers. They both find that RTP has some impact on peak-shaving but no significant impact on peak-shifting. The latter study finds that time-of-use (TOU) pricing does produce significant peak-shifting behavior.

(Vytelingum et al. 2011) and (Voice et al. 2011) describe the problem of load-shifting peaks under RTP amongst micro-storage agents such as plugin electric vehicles. In this problem, many agents independently converge their loads on short time intervals with lower expected prices thus leading to undesirable load peaks. (Gottwalt et al. 2011) observe the same phenomenon in their simulations and call it the *avalanche* effect. (Ramchurn et al. 2011) illustrate this problem with Figure 1 and refer to it as *herding*, terminology which we adopt here, and also propose an adaptive algorithm that imposes inertia on proportions of customers to mitigate the impact of this effect. (Voice et al. 2011) propose a mechanism based on penalties for deviation from past behavior to achieve a similar result. We propose an alternate solution to this problem based on decision-theoretic agent behavior.

The benefit of formulating agent decisions as multiscale decision-making problems is addressed in (Barber 2007). (Wernz and Deshmukh 2010) formalize the concept in the domain of organizational behavior using an analytical game-theoretic approach. The complexity of our simulation-based model makes it unsuitable for a simple analytic solution and we instead rely on an algorithmic approach. Our utility optimizing agent is related to distributed control with constrained reasoning, *e.g.*, (Modi et al. 2005), and also team coaching in adversarial settings, *e.g.*, (Riley 2005). Our stochastic approximation algorithm is based on classic multiattribute utility theory (Wellman 1985) and decision theory (Horvitz, Breese, and Henrion 1988).

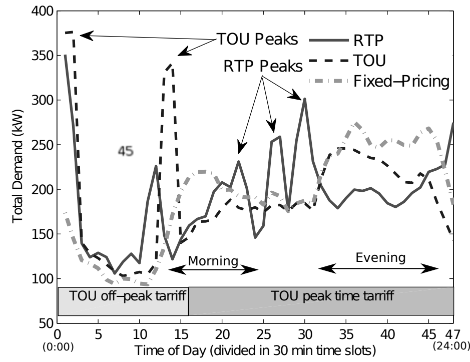

Figure 1: TOU and RTP tariffs both result in undesirable shifted peaks. (Ramchurn et al. 2011)

Our work on hierarchical Bayesian timeseries simulation is based on Seasonal ARIMA models (Cryer and Chan 2008) combined with multilevel Bayesian models (Gelman and Hill 2007). Some examples of autoregressive Bayesian prediction with latent variables are described in (West and Harrison 1997). Hierarchical Bayesian models (HBMs) and Dynamic Bayesian Networks (DBNs) have been studied extensively, *e.g.*, (Murphy 2002). We apply Gibbs-sampling based inference techniques (Geman and Geman 1984) as is typical with complex DBNs (Koller and Friedman 2009).

Our use of the term *factored* is analogous, although not strictly similar, to factored state representations for reinforcement learning, factor graphs in probabilistic graphical models, and factored (*i.e.*, multiplicative) ARIMA models. We intend the term to convey that the model's behavior is characterized by a composition of its determining factors.

Multiscale Decision-Making

We model Smart Grid customers in markets with retail competition where they have a choice of tariffs offered by several energy aggregators (Braun and Strauss 2008) or brokers (Reddy and Veloso 2011). We observe that Smart Grid customers are faced with a *multiscale decision-making* problem along at least the following two dimensions:

1. *Temporal*: Customers must simultaneously optimize their current capacity levels given their tariff prices and also their tariff choices given their expected capacity levels. While capacity optimization occurs at high frequency, tariff selection occurs at a lower frequency.

2. *Contextual*: For example, a single household must consider the optimal behavior of each appliance individually but also of all appliances together and similarly of the household unit versus its *neighboring* households. While this contextual dimension is loosely related to the spatial dimension, it can also be applied more broadly using alternate definitions of *neighborhood*.

Therefore, at any timeslot, t, an appliance or customer must perform the following optimization:

$$\underset{y_t}{\operatorname{argmax}}\ U_S(p_t, y_t, U_N(y_t)) \tag{1}$$

where y_t is the capacity level, U_S is a self-utility function, p_t is the applicable tariff price structure, and U_N is the neighborhood-utility function. Then at certain less frequent timeslots, t', that occur every τ timeslots, the following optimization is needed:

$$\underset{z \in \mathcal{Z}_{t'}}{\operatorname{argmax}}\ V_S(P^z_{t'}, Y_{t'}, V_N(Y_{t'})) \tag{2}$$

where $\mathcal{Z}_{t'}$ is the set of applicable tariffs available at t', Y_τ is the expected *capacity profile* over the horizon τ, P^z_τ is the vector of expected prices under a tariff z over τ at t', and the utility functions, V_S and V_N, are evaluated over τ.

Factored Customer Representation

Let a Smart Grid customer, C, be defined as:

$$C = \langle \{\mathcal{B}_i\}_{i=1}^N, \{S_i\}_{i=1}^N, U \rangle,\ \mathcal{B}_i = \{O_{ij}\}_{j=1}^{M_i}$$

$\mathcal{B}_i$ is a *capacity bundle* which contains one or more *capacity originators*, O_{ij}. S_i is a *tariff subscriber* and U is a *utility optimizer*. Figure 2 illustrates this composition; the 1-to-1 correspondence between $\mathcal{B}_i$ and S_i is shown with solid lines while the dotted arrows indicate recommendations from U.

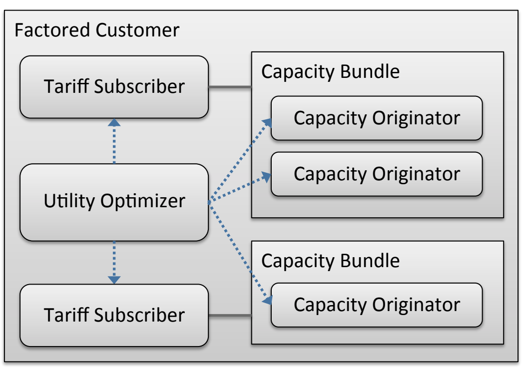

Figure 2: An example factored customer modeled with three capacity originators in two capacity bundles.

Capacity Originator The behavior of each capacity originator, O, is determined by a *base capacity generator* and several *influence factors*. The base capacity generator is either an arbitrary probability distribution or a *timeseries generator*. The capacity originator generates an original capacity level, y^o_t, for each discrete timeslot, t, by drawing from the base generator distribution or by obtaining the next prediction from the timeseries generator.[1] The next section describes the prediction method of the timeseries generator.

The original capacity level, y^o_t, is then adjusted according to the following influence factors:

- **Calendar**: The subfactors, time-of-day, day-of-week, and month-of-year are given relative weights.
- **Pricing**: Th value of this factor is computed based on absolute tariff prices or a price elasticity function applied to deviation of current prices from a benchmark price.
- **Weather**: Factor values based on segmented real values or elasticity functions applied to benchmark deviations are computed for the following subfactors: temperature, wind speed, wind direction, and cloud cover.

The adjusted capacity level, y^a_t, obtained as the product of y^o_t and each factor value is then used to forecast *capacity profiles*, which are used for adaptive capacity control as described later. A capacity originator can be viewed variably as an appliance, or recursively as one or more customers. It can also be an autonomous control agent or a decision-support interface to humans who manually control capacity. Important additional influence factors that enable adaptive capacity control are also described in a later section.

Capacity Bundle A capacity bundle is an aggregation of capacity originators with the constraint that all originators in the bundle must be of the same *capacity type*, which can be categorized coarsely as {*consumption, production, storage*} or more finely with types such as *household consumption*, *wind production*, and *electric vehicle storage*. Typically, one bundle is assigned to a single tariff, however when the bundle represents a collection of grid-connected entities as in a farming cooperative, the bundle can allocate segments of its population to different tariffs.

Tariff Subscriber A tariff subscriber is an autonomous or human agent that manages the assignment of a capacity bundle to one or more of the available tariff choices. The agent is modeled using a multinomial logit choice model where the utility of each tariff choice is assumed to be given. Two additional factors determine the tariff selection process:

- **Inertia**: This is modeled as a probability distribution, a draw from which decides whether the subscriber maintains its corresponding capacity bundle in its current tariff subscriptions or whether it considers reassignment.
- **Rationality**: This factor, $\lambda \in [0, 1]$, determines the degree to which the utility values, $U^S_\tau(z)$, associated with the tariff choices influence tariff selection:

$$Pr(z) = \frac{e^{\lambda U^S_\tau(z)}}{\sum_{\mathcal{Z}} e^{\lambda U^S_\tau(z)}} \tag{3}$$

The probability values thus derived can be used for random selection of a single tariff or for proportional allocation to multiple tariffs.

Utility Optimizer In their survey of customer behavior under real-time pricing tariffs offered by over 20 utilities in the United States, (Barbose, Goldman, and Neenan 2005) observe limited responsiveness to price changes. They note that this may be due to inadequate customer-side automation that can manage price volatility risk and capitalize on opportunities arising from real-time price changes. We propose our utility optimizer component as an optionally deployed intelligent autonomous agent to serve this goal. We

[1] We use the term *capacity* to describe both demand and supply.

extend its scope to automate the frequent tariff subscription decisions that have become more important for customers in recent years with increasing competition in retail markets. We describe its algorithms in detail in a later section.

We conclude this section with some examples of how our factored model can be instantiated to represent varying customer types and agent granularities.

1. *Fine-grained Household Model*: Individual appliances are represented as separate capacity originators with each originator drawing its base capacity from an appropriate probability distribution.
2. *Coarse-grained Household Model*: All consumption appliances are represented in aggregate as one capacity originator that draws from a timeseries generator, and rooftop solar panels form another originator in a separate bundle.
3. *Farming Cooperative Model*: All consumption for all farms is modeled as one originator in its own bundle, and the windmills on each farm are modeled as separate originators but collected in one wind production bundle.

In each of the above examples the tariff subscriber and utility optimizer may be autonomous agents integrated into the capacity control automation or be separate services that help inform humans and execute their commands.

Bayesian Timeseries Simulation

We now focus on the *timeseries generator* component of a capacity originator. Our goal here is to learn information from one timeseries to approximately replicate it using only the generic factors we consider in our model. As a concrete example, consider the top panel of Figure 3 which shows 12 days of hourly consumption for 10 households from the MeRegio pilot project (Hirsch et al. 2010) as simulated by the fine-grained customer model of (Gottwalt et al. 2011). Such data can often be collected from the real world for various customer types with manageable effort. However, incorporating the distinct customer types into a simulation model typically requires considerable analytical and programming effort. We instead use statistical machine-learning techniques to automatically *learn* the model parameters for a reusable *hierarchical Bayesian timeseries model* so that noise-added replications of the timeseries can be produced for online simulation using only our factored model.

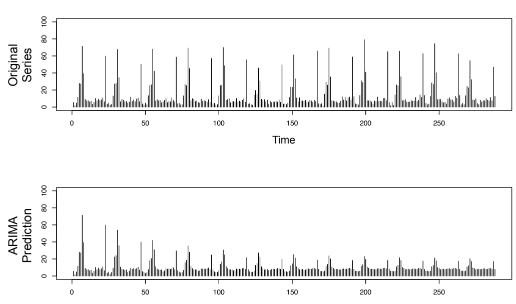

Figure 3: The top panel shows an observed timeseries; the bottom panel shows the decaying ARIMA prediction.

A standard approach to this problem is to predict from a Seasonal ARIMA model.[2] However, such models exhibit fast exponential decay towards a grand mean when the predicted values are themselves used sequentially for further prediction. We provide a solution to this long-range prediction problem by fitting a hierarchical Bayesian model with autoregresssive covariates and other influence factors. We use Gibbs sampling to fit models like the example in Equations 4-9, which includes ARMA(1,1)$\times$(1,1)$_{24}$ factors (Eq. 6 and 7) and daily and hourly factors (Y_d and Y_h).

$$Y_{1,t} \sim N(Y_0 + Y_{d,t} + AR_t + MA_t, \sigma^2) \quad (4)$$

$$Y_{2,t} \sim N(Y_0 + Y_{h,t} + AR_t + MA_t, \sigma^2) \quad (5)$$

$$AR_t \leftarrow \phi_1 Y_{t-1} + \Phi_1 Y_{t-24} \quad (6)$$

$$MA_t \leftarrow \theta_1 \varepsilon_{t-1} + \Theta_1 \varepsilon_{t-24} + \theta_1 \Theta_1 \varepsilon_{t-25} \quad (7)$$

$$Y_d \sim N(0, \eta_1^2);\ d = 1..7 \quad (8)$$

$$Y_h \sim N(0, \eta_2^2);\ h = 1..24 \quad (9)$$

We eliminate the labeling confusion between hourly and daily fixed-effects by fitting the same data to two output variables, Y_1 and Y_2 (Eq. 4 and 5). This overestimates the hourly and daily coefficients but we compensate for that by taking a combination later in the model. We include separate variance components, η_1^2, η_2^2 and σ^2, for the daily and hourly intercepts and for the output variables (Eq. 4, 5, 8 and 9). We can include informed priors, *e.g.*, from an ARIMA fit, on the AR and MA coefficients, ϕ_1, Φ_1, θ_1 and Θ_1.

We then use the learned coefficients of the calendar factors to generate a complementary series, which we add to the prediction Y_t^f from the hierarchical model to *augment* the signal (Eq. 11). The complementary series is computed as a convex combination of the daily and weekly factors, scaled by a factor ν that is a logarithmic function of the prediction horizon. It is added to Y_t^f in proportions determined by a joint optimization of the parameters λ and γ using least squares loss or KL-divergence. All the model parameters thus estimated can then be combined with Gaussian noise to generate many replicating timeseries, Z_t^{bf} (Eq. 13).

$$Y_t^f \leftarrow Y_0 + Y_{d,t} + Y_{h,t} + AR_t + MA_t \quad (10)$$

$$Y_t^{bf} \leftarrow Y_t^f + \lambda\nu((1-\gamma)Y_{d,t} + \gamma Y_{h,t}) \quad (11)$$

$$\lambda^*, \gamma^* \leftarrow \operatorname*{argmin}_{\lambda,\gamma} \sum_t (Y_t^{bf} - Y_t)^2 \quad (12)$$

$$Z_t^{bf} \sim N(Z_t^f + \lambda^*\nu((1-\gamma^*)Y_{d,t} + \gamma^* Y_{h,t}), \sigma^2) \quad (13)$$

We note that while the above example only uses two influence factors, daily and hourly effects, we can include any additional influence factors for which we have labeled data with the training timeseries, Y_t. One significant advantage of using the Bayesian framework for simulation in this manner is that we can now obtain generated timeseries that deliberately behave differently than the training series by adding *a priori* or *a posteriori* bias to the influence factors.

[2] We observe that the series is stationary with stable volatility, so we do not include any integrative or GARCH components.

Stochastic Utility Optimization

In an earlier section, we defined the multiscale decision-making problem faced by Smart Grid customers as two utility maximization problems with different temporal and contextual components. We now present how we define those utility functions and describe the approximate algorithm to maximize them, which also addresses the herding problem.

We define a *capacity profile*, ρ_H as a vector of capacity values up to some horizon, H. We define the distance between two profiles as the sum of squared point deviations:

$$D(\rho_H, \tilde{\rho}_H) = \sum_{t=1:H} (\rho_t - \tilde{\rho}_t)^2$$

We define an *admissible permutation* of a given profile ρ_H as any profile $\tilde{\rho}_H$ that satisfies two constraints: (i) $\tilde{\rho}_H$ must have the same aggregate capacity as ρ_H, and (ii) $min(\tilde{\rho}_H) \geq min(\rho_H)$. We provide two methods for generating admissible permutations for capacity shifting:

- **Temporal Shifts**: A lag-δ permutation, $\rho_{H,\delta}$, is obtained by rotating the capacity elements of ρ_H by δ timesteps. $\mathcal{R}_T^{\rho_H}$ is then the size H set of all temporal shifts for ρ_H.
- **Balancing Shifts**: Let ρ_j be the permutation obtained from ρ_i by setting $\rho_j = \rho_i$, computing $x = 0.5\, range(\rho_i)$, subtracting x from $max(\rho_H)$ and adding x to $min(\rho_H)$. ρ_k can then be similarly obtained from ρ_j and so on. $\mathcal{R}_B^{\rho_H}$ is then the set of balancing shifts obtained by recursively computing ρ_j and adding it to $\mathcal{R}_B^{\rho_H}$ until $range(\rho_j) < \epsilon$ or until the size of $\mathcal{R}_B^{\rho_H}$ reaches a threshold. Intuitively, this procedure generates permutations that are *flatter* than the previous profile at each iteration.

Temporal shifts of a capacity profile are more appropriate when the duty cycle of a capacity originator has a fixed pattern. For always-on appliances of variable capacity and for collections of appliances or customers, it is often possible to shift individual peaks as we do with the balancing shifts.

We extend the representation of a capacity originator to define an *adaptive capacity originator*, which can receive a *profile recommendation* from the utility optimizer and adapt its capacity accordingly. A recommendation is a set of permutations to the currently forecast profile, $\hat{\rho}_H$, of the capacity originator. A permutation is only included in a recommendation if its associated expected payment (debit if consumption, credit if production) is better than that of the forecast profile. The self-utility of a permutation is defined as:

$$U_S(\tilde{\rho}_H) = \Delta f_p(\tilde{\rho}_H) + w_D\, D(\tilde{\rho}_H, \hat{\rho}_H) + w_N\, U_N(\tilde{\rho}_H)$$

Δf_p is a function that computes the change in expected payment relative to the forecast profile and U_N is a neighborhood utility function described below. The weighted distance from $\hat{\rho}_H$ to $\tilde{\rho}_H$ represents the *shifting disutility* to the capacity originator. Thus, the utility of each permutation is a weighted combination of the change in payment, the shifting disutility and the neighborhood utility. We then compute probability values for each permutation by scaling the utility values to $[-3, +3]$ and taking their exponents; this yields smoothly decaying probability values for the permutations included in the recommendation, which are then used in a multinomial logit choice model to randomly choose the profile to be executed. This procedure forms an approximate solution to the optimization problem in Eq. 1.

In implementation, we refine this method to separate the roles of the capacity originator and the utility optimizer. The capacity bundle containing the originator is assumed to be the relevant neighborhood. The utility optimizer first computes local utilities, *i.e.*, U_S with $w_N = 0$, for each permutation and then performs Monte Carlo sampling over the profile recommendations being submitted to each capacity originator in the bundle to obtain an expected aggregate profile for the whole bundle. It then iterates over the permutations in each capacity originator, holding the expected profiles for all other capacity originators constant, to compute the aggregate payment benefit for the bundle given that permutation, $\tilde{\rho}_H$, and assigns it as the *bundle value*, $U_N(\tilde{\rho}_H)$. When sampling, the utility optimizer computes the expected profile of each capacity originator using three *responsiveness* factors, each of which takes on a value between 0 and 1:

- **Reactivity**: The probability that the capacity originator will at least consider the recommended permutations.
- **Receptivity**: The probability that the capacity originator will adopt one of the permutations with the highest utility value in the recommendation.
- **Rationality**: This factor is applied to the multinomial logit choice model by the originator to get new probability values for the given utility values of each permutation.

These responsiveness factors are intended to capture the possibility that the capacity originator may in fact be modeling a human decision-maker. Periodically, the expected profiles for each capacity originator are also used to compute the tariff utility values, $V_S(z)$ as in Eq. 2, to be used in the probabilistic choice model of Eq. 3 by the tariff subscribers.

The adaptive capacity control approach described in this section avoids the herding problem typically seen with shifting under variable-price tariffs. This is at least partly due to the probabilistic multinomial logit choice model we employ, which ensures that shifted capacities are assigned equitably to equivalent future timeslots as opposed to the typical approach where capacity is shifted simply to the next timeslot with low expected prices. Furthermore, our weighted utility function that explicitly accounts for bundle value further ensures that many capacity originators do not converge on the same timeslot. The adapted capacity profiles generated by our approach are not only cost-efficient for customers but also less volatile than the raw profiles, thus making it easier for energy brokers and physical service providers to anticipate them, which in turn leads to greater social welfare.

Experimental Results

We present results from experiments with our timeseries simulation method and also our adaptive capacity control mechanism. The bottom panel of Figure 4 shows the prediction of our augmented hierarchical Bayesian model over 12 days. We see that it is a much better approximation of the original series shown in the top panel of Figure 3 compared to the ARIMA prediction shown in the top panel here.

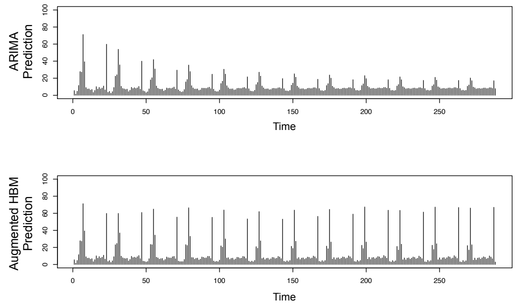

Figure 4: The ARIMA prediction in the top panel exhibits fast decay while our prediction in the bottom panel does not.

Figure 5 shows a tolerance vs. accuracy plot for variations of our methodology. We see that, for example, at 20% tolerance for prediction error, the base ARIMA model achieves 40% accuracy. Our method achieves 75% testing accuracy, *i.e.*, initializing our prediction sequence with a timeseries sample distinct from the one we used for learning. We also show training accuracies for variations of our methodology using ARIMA priors on the hierarchical model vs. vague priors and also with and without the augmentation step.

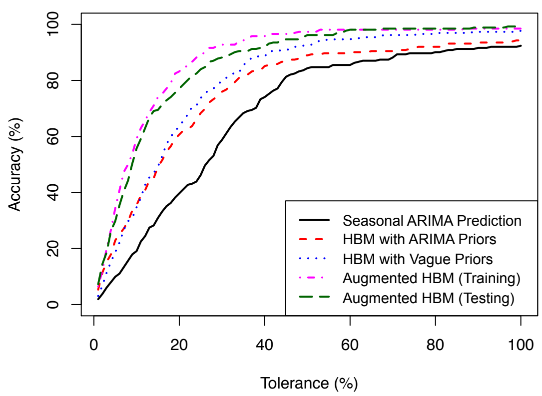

Figure 5: Various incremental steps of our simulation methodology steadily improve prediction accuracy.

We implemented our factored model representation in the Power TAC environment. In one representative experiment, we modeled 10000 customers as profiled in the MeRegio study. We instantiated 100 capacity originators each representing 100 customers as modeled by a replicating timeseries generator trained on the output of the fine-grained simulation model of (Gottwalt et al. 2011). We then offered TOU tariffs with higher prices in the afternoon and evening as is typical in many real-world TOU tariffs. We observed that when the adaptive capacity control was enabled, our customers typically obtained 5-12% lower usage charges. Perhaps more importantly, the resulting shifted capacities have significantly lower volatility as shown in Figure 6. The dashed line shows the original capacities, the long-dashed line shows shifting with only temporal shifts enabled, and the solid line shows the results with only balancing shifts. Figure 7 depicts these results as a boxplot and also shows the effect of combining temporal and balancing shifts.

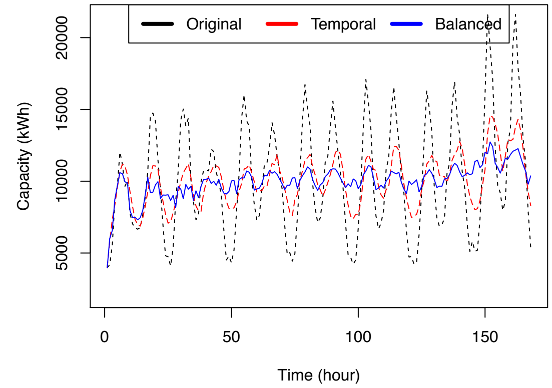

Figure 6: The adaptive capacity control implemented by our utility optimizing agent achieves significant *flattening* and does not exhibit herding behavior under typical TOU tariffs.

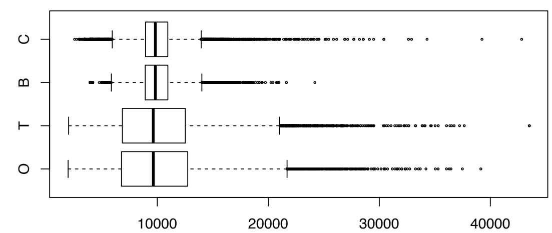

Figure 7: Capacity ranges of an original series (O) and our TOU price-adapted series with temporal shifts (T), balancing shifts (B) and both shifts combined (C).

Conclusion

In this paper, we have contributed (i) a formulation of the decisions to be made by Smart Grid customers as a multiscale decision-making problem, (ii) a versatile factored customer representation, (iii) a learning-based hierarchical Bayesian timeseries simulation method, and (iv) a stochastic adaptive control algorithm which achieves usage charge benefits and capacity smoothing and also does not exhibit any herding behavior under typical variable-price tariffs. We have also supported our contributions with experimental results from the implementation of our model on the Power TAC platform. We plan to further evaluate the behavior of our model with additional real-world data and in official Power TAC tournaments. In future work, we intend to study extensions to the problem where the utility optimizer and the capacity originators are engaged in semi-cooperative relationships.

Acknowledgements

This research was partially sponsored by the Portuguese Science and Technology Foundation. The views and results contained in this document are those of the authors only.

References

Allen, E.; LaWhite, N.; Yoon, Y.; Chapman, J.; and Ilic, M. 2001. Interactive Object-Oriented Simulation of Interconnected Power Systems using SIMULINK. *IEEE Transactions on Education, vol 44*.

Amin, M., and Wollenberg, B. 2005. Toward a smart grid: Power delivery for 21st century. *IEEE Power and Energy*.

Armstrong, M.; Swinton, M.; Ribberink, H.; Beausoleil-Morrison, I.; and Millette, J. 2009. Synthetically derived profiles for representing occupant-driven electric loads in Canadian housing. *Journal of Building Performance Simulation 2 (1), 1530*.

Barber, K. 2007. Multi-Scale Behavioral Modeling and Analysis Promoting a Fundamental Understanding of Agent-Based System Design and Operation. Technical report, DARPA Final Report by University of Texas, Austin (AFRL-IF-RS-TR-2007-58).

Barbose, G.; Goldman, C.; and Neenan, B. 2005. Electricity in real time–a survey of utility experience with real time pricing. *Energy 30 (1), 1418*.

Borenstein, S. 2002. The Trouble With Electricity Markets: Understanding California's Restructuring Disaster. *Journal of Economic Perspectives*.

Braun, M., and Strauss, P. 2008. Aggregation approaches of controllable distributed energy units in electrical power systems. *Journal of Distributed Energy Resources*.

Chrysopoulos, A., and Symeonidis, A. 2009. Improving Agent Bidding in Power Stock Markets Through A Data Mining Enhanced Agent Platform. In *Agents and Data Mining Interaction workshop, AAMAS'09*.

Cryer, J., and Chan, K. 2008. *Time Series Analysis with Applications in R.* Springer.

Gelman, A., and Hill, J. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models.* Cambridge University Press.

Geman, S., and Geman, D. 1984. Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 6 (6): 721741*.

Gomes, C. 2009. Computational Sustainability: Computational Methods for a Sustainable Environment. *The Bridge, National Academy of Engineering, 39 (4)*.

Gottwalt, S.; Ketter, W.; Block, C.; Collins, J.; and Weinhardt, C. 2011. Demand side management - A simulation of household behavior under variable prices. *Energy Policy 39, 8163-8174*.

Guo, Y.; Li, R.; Poulton, G.; and Zeman, A. 2008. A Simulator for Self-Adaptive Energy Demand Management. In *IEEE Conf. on Self-Adaptive and Self-Organizing Systems*.

Hammerstrom, D. 2008. Pacific Northwest GridWise Testbed Demonstration Projects; Part I. Olympic Peninsula Project. Technical report, PNNL-17167, Pacific Northwest National Laboratory.

Hirsch, C.; Hillemacher, L.; Block, C.; and Schuller, A. 2010. Simulations in the smart grid field study MeRegio. *Information Technology 52 (2), 100106*.

Horvitz, E.; Breese, J.; and Henrion, M. 1988. Decision theory in expert systems and artificial intelligence. *Int. Journal of Approximate Reasoning, 2, 247-302*.

Karnouskos, S., and de Holanda, T. 2009. Simulation of a smart grid city with software agents. In *Computer Modeling and Simulation*.

Ketter, W.; Collins, C.; Reddy, P.; and Flath, C. 2011. The Power Trading Agent Competition. Technical Report ERS-2011-011-LIS, Erasmus Research Institute of Management.

Koller, D., and Friedman, N. 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.

Kolter, J., and Ferreira, J. 2011. A large-scale study on predicting and contextualizing building energy usage. In *AAAI Conf. on Artificial Intelligence (AAAI-11)*.

Modi, P.; Shen, W.; Tambe, M.; and Yokoo, M. 2005. ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees. *Artificial Intelligence Journal*.

Murphy, K. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. Dissertation, University of California, Berkeley.

Paatero, J., and Lund, P. 2006. A model for generating household electricity load profiles. *International Journal of Energy Research 30 (5), 273290*.

Ramchurn, S.; Vytelingum, P.; Rogers, A.; and Jennings, N. 2011. Agent-Based Control for Decentralised Demand Side Management in the Smart Grid. In *Int. Conf. on Autonomous Agents and Multiagent Systems*.

Reddy, P., and Veloso, M. 2011. Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. In *AAAI Conf. on Artificial Intelligence (AAAI-11)*.

Riley, P. 2005. *Coaching: Learning and Using Environment and Agent Models for Advice*. Ph.D. Dissertation, Computer Science Dept., Carnegie Mellon University.

Strbac, G. 2008. Demand side management: benefits and challenges. *Energy Policy 36 (12), 44194426*.

Sun, J., and Tesfatsion, L. 2007. An Agent-Based Computational Laboratory for Wholesale Power Market Design. In *IEEE Power and Energy Society General Meeting*.

Voice, T.; Vytelingum, P.; Ramchurn, S.; Rogers, A.; and Jennings, N. 2011. Decentralised Control of Micro-Storage in the Smart Grid. In *AAAI Conf. on Artificial Intelligence (AAAI-11)*.

Vytelingum, P.; Voice, T.; Ramchurn, S.; Rogers, A.; and Jennings, N. 2011. Theoretical and Practical Foundations of Large-Scale Agent-Based Micro-Storage in the Smart Grid. *Journal of Artificial Intelligence Research*.

Wellman, M. 1985. Reasoning about preference models. Technical report, MIT, MIT/LCS/TR-340.

Wernz, C., and Deshmukh, A. 2010. Multiscale Decision-Making: Bridging Organizational Scales in Systems with Distributed Decision Makers. *Journal of Operational Research 202 828-840*.

West, M., and Harrison, P. 1997. *Bayesian Forecasting and Dynamic Models*. Springer-Verlag, 2nd ed.