Flexibility Meets Variability: A Multiagent Constraint Based Approach for Incorporating Renewables into the Power Grid

Xiaoyue Jiang, Ramgopal Mettu, K. Brent Venable
Dept. of Computer Science
Tulane University
6823 St Charles Ave,
New Orleans, LA 70118
Email: {xjiang,rmettu,kvenabl}@tulane.edu

Geoff Parker
Freeman School of Business
Tulane University
6823 St Charles Ave,
New Orleans, LA 70118
Email: gparker@tulane.edu

Abstract

This paper outlines a new approach to creating value from the Smart Grid by incorporating individual households into the response system that must be deployed to accommodate increasingly large sources of intermittent renewable power. We propose a framework that couples agent-based AI techniques with envelope methods. Envelope methods provide a unified mathematical framework to model intermittent renewable resources, conventional dispatchable resources, demand side response, and storage. The overall goal of our system is to develop a distributed autonomous agent architecture that is able to facilitate market transactions among load serving entities, residential consumers, conventional merchant power producers, and intermittent power producers.

Introduction

The deployment of renewable power resources, such as wind and solar, has grown rapidly because of policy actions, including renewable portfolio standards, that seek to reduce the power industry's carbon emissions. Increased global adoption has led to dramatic cost reductions, creating positive feedback loops that drive further adoption. However, to maintain stability, systems with large renewable portfolios must include dispatchable resources, usually fossil fueled, that are ramped up and down to accommodate renewable variability. These dispatchable resources are often run outside of their optimal operating specifications with the paradoxical result that some systems can experience higher carbon emissions as the penetration of renewable power increases. The adverse effects can be directly observed through increased *negative* nodal prices observed on some systems, especially at night when more wind power is generated than the system demands.

Because power demand and supply must be balanced in real time, grid operators have maintained the balance. However, the assumption that the grid must absorb all variability is becoming less true because building automation technology is on the horizon. Importantly, houses and buildings are an underutilized source of adaptability if properly used. In particular, because of the thermal inertia of buildings and the increasing controllability of power consuming systems, end-users can be directly incorporated into the response system that must be deployed to accommodate renewable power.

In this paper we outline a system that utilizes artificial intelligence to manage power consumption based on end-user preferences and system capacity. Our system is end-to-end in the sense that household preferences for demand are aggregated and communicated up through the supply network in a hierarchical fashion. Variability in supply and demand is incorporated through agents at each level of this hierarchy using *envelope bounds*. Importantly, the envelope method can characterize the quality of service at various levels of aggregation and offers the possibility of a unified framework in which the variability imposed from any supply side resource and any demand side resource can be quantified in terms of the impact on overall system reliability.

Envelope-based model of energy variability

Our approach implements a computational layer over the energy grid that is organized (Figure 1) to hierarchically manage system load. We believe this is the key aspect our approach that makes it possible to compute real-time responses to rapidly changing supply and demand patterns driven by renewable energy generation at both the household and grid. The overall system is organized so that each entity in the hierarchy has a computational agent that plays a role in balancing supply and demand. Within each agent, the basic component of estimating supply and demand bounds is done through the use of the envelope method. This method originates from the study of quality of service guarantees in networking, and has recently been shown to be useful in the context of power system analysis (Jiang, Parker, and Shittu 2010; 2013).

Given a usage schedule (Figure 2(a1)) that gives instantaneous consumption, the first step of the envelope methodology is to convert usage to the energy domain (i.e., cumulative MWh over time) (Figure 2(a2)). Then, an upper envelope $\alpha(t)$ is constructed for the energy flow (Figure 2(a3)), which characterizes the maximum energy demand as a function of the time scale. More specifically, this upper envelope is obtained by applying Legendre transforms to the original energy flow, where the corresponding Legendre conjugates describe the net excessive electricity demand (in MWh) above a constant demand flow at given capacity levels (in MW), and are realized through the leaky bucket mechanism

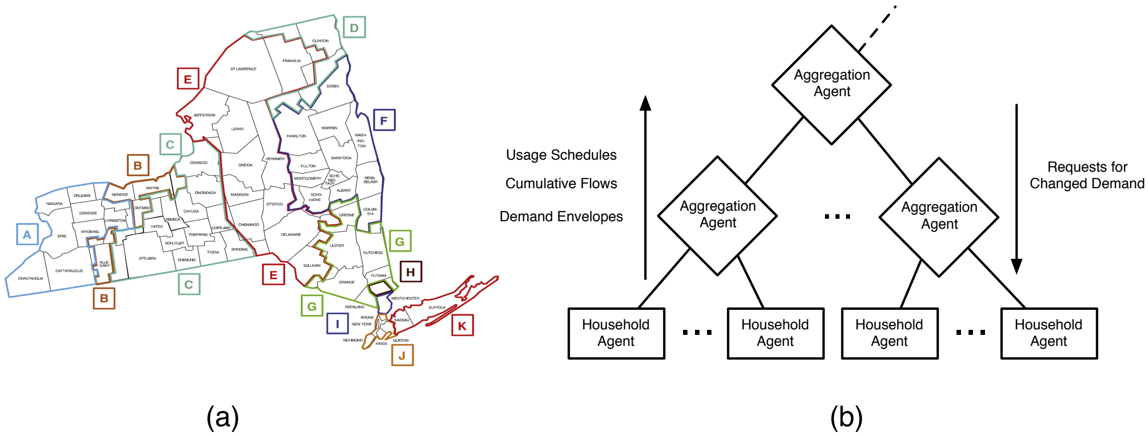

Figure 1: **Overview of Agent Organization.** (a) The regional topography of the New York Independent Operator is shown, which pricing nodes indicated as squares. (b) In our proposed system, agents at the consumer level are leaves of an aggregation tree that is defined according to the regional energy market being modeled. The root of this tree corresponds to a pricing node (indicated by colored letters in (a)).

(Figure 2(a4)) (Jiang, Parker, and Shittu 2013). By further taking variable sources of energy (for example, wind or solar power) as a negative load (Figure 2(a5)), lower envelopes (Figure 2(a6)) of generation can also be derived. The lower envelope of wind generation represents the guaranteed minimum output over any period of a given duration. The time granularity of the boundary calculations can be chosen for specific applications such as planning (years), forward contracting (months/weeks), day-head trading (hours), and operations (minutes).

Figure 2(b) demonstrates the ability of the envelope method to model real-world variability in power generation, showing data from July 2009 for energy from resources including wind farms, solar plants, and geothermal plants within the California ISO region. As described above, we construct the upper and lower envelope of system load to study the demand pattern for capacity. In particular, Figure 2(b) compares these envelopes of system load and net load (e.g., Load - Wind), which shows the impact of wind energy. The shift of the right tail of the conjugate curve quantifies the reduction of the system peak load due to wind power, which is wind's capacity contribution to the system. The shift of the right tail is also accompanied with a shift of the left tail, which quantifies the reduction of system base load, confirming the negative impact on the system. These bounds quantify the intuitive idea that variability in wind generation is a challenge. In Figure 2(b) we can see that the system must be able to handle about 4GWh in load variability in order to accomodate wind generation. Using this knowledge we can deploy a utility-scale pumped storage, for example HELM in California, to "shift" the curves so that surplus and deficit bounds match. This approach can currently be adopted only on a utility-scale level, while our system manages this type of gap in a distributed fashion at the household level.

A hierarchical multi-agent model of an energy market

Our system is organized hierarchically as an *aggregation tree* and has two fundamental types of agents, *consumer agents* (CA) and *aggregation agents* (AA). Consumer agents are leaves of the aggregation tree, and each consumer agent is responsible for managing the energy demands of an individual consumer. The basic role of a consumer agent is to produce an accurate demand plan for a given time window, taking into account possible sources of uncertainty, such as weather or plan changes of the inhabitants, preferences and energy efficiency. The aggregation tree is rooted at pricing nodes of a local energy market and models that particular market (see Figure 1). Each local market has hundreds of thousands of customers, and thus a tree-structured approach is a necessity. That is, it offers the possibility of highly concurrent processing of household preferences, with synchronization staged at each aggregation node.

Both consumer and aggregation agents, regardless of level, interact in order to synthesize demand information in a bottom up fashion: information from each consumer flows up the agent tree, being aggregated at each level. The synthesis of information at each aggregation agent yields a demand envelope and cumulative energy flow that is passed up to the next level. Once information passes upward to the pricing node, the aggregate demand is compared with the desired demand and negotiation requests are then computed in a downward fashion. Thus each aggregation agent may receive a request to change its usage from its parent, necessitating the calculation of a new demand envelope. These requests are in the form of pricing (or usage) changes and are used to compute a new cumulative energy flow (and resulting demand envelope) for that agent. The basic role of an aggregation agent is to fairly split requests for changes among its children, based on their respective energy needs. Bottom-up and top-down interaction between agents proceeds as long as some change is initiated at the consumer level (due to new production or consumption of energy), or

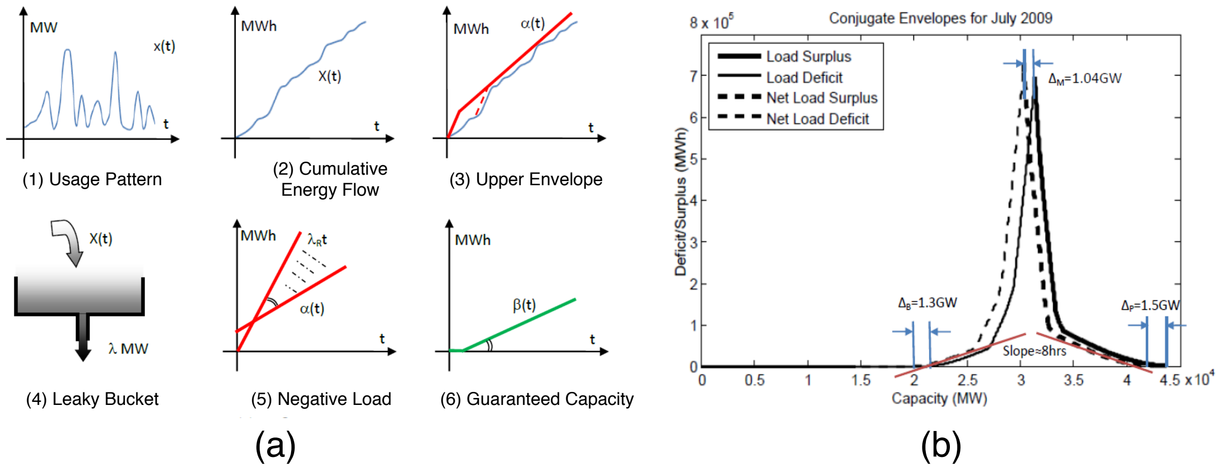

Figure 2: **Envelope Method.** (a) Overview of the computation of upper and lower envelopes; (b) Example of envelope calculations on Cal-ISO data quantifying the effect of variability from wind power generation with respect to capacity.

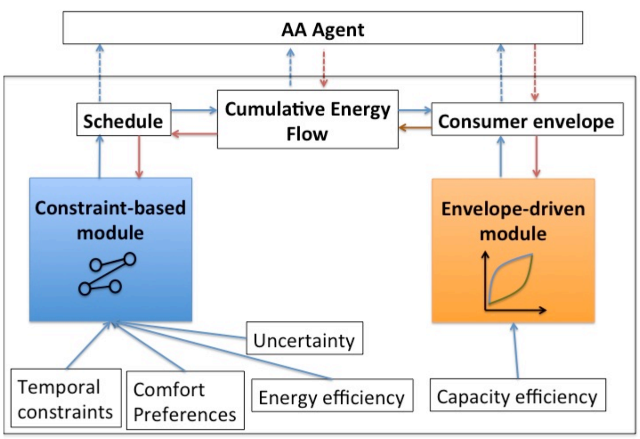

Figure 3: Architecture of the consumer agent.

a change in supply (and thus pricing) at a higher level necessitates new downward requests. We describe each agent in detail below.

Consumer Agents

The model of the consumer agent we envision is depicted in Figure 3 and has two fundamental components: a constraint-based model including temporal constraints on loads, preferences and uncertainty, and an envelope-driven module in charge of optimizing with respect to capacity efficiency within the agent. The constraint-based model of the domestic energy domain will represent: (a) the *temporal constraints* involving energy consuming or producing activities within the consumer; (b) the *preferences* in terms of comfort of the consumer inhabitants and in terms of energy efficient behaviors; and, (c) the *uncertainty* arising from a dynamically changing environment (such as weather) as well as changes in the usage patterns of the inhabitants.

The core technology underlying the *constraint-based module* are temporal constraint satisfaction problems (Dechter, Meiri, and Pearl 1991) where activities are modeled by their end and start time and by intervals containing all possible durations. This very general representation paradigm is amenable for incorporating temporal information and requirements of both loads with fixed durations, such as washing machines and dish washers (also called shiftable static loads in the literature (Ramchurn et al. 2011)), as well as thermal loads controlled by smart thermostats.

We exploit several appealing features of this paradigm. First of all, different types of preferences can be easily incorporated to model comfort and energy efficiency, despite their clearly different semantics (Khatib et al. 2007; Bartak, Morris, and Venable 2014). For example, the $(\min, +)$ approach to preferences, where they are treated as penalties the sum of which should be minimized, appears as a natural choice to model energy efficiency. Moreover, solutions to temporal constraint problems are not fixed schedules for the activities, but rather, a compact representation (known as the minimal network) from which all the consistent schedules can be computed very efficiently. We exploit this to minimize the need of schedule recomputation while adapting to change. Finally, different types of uncertainty, such as the presence of events which are outside of the control of the automated agent or conditional information on the presence of certain activities can be handled (Morris, Muscettola, and Vidal 2001; Vidal and Fargier 1999). This is essential to be able to incorporate the actions of inhabitants into the model.

The output of the constraint-based model is a compact representation of dynamically dispatchable, comfort optimal and energy efficient schedules obtained by applying techniques which extend those described in (Rossi, Venable, and Yorke-Smith 2006). While the execution of the schedule is done incrementally we need to produce an estimate of the energy demand of the household in the future. In the temporal constraint literature the concept of energy envelopes has been used to model the bound on energy consumption associated to the set of schedules represented compactly by a minimal dispatchable network (Muscettola 2002; 2004). We extend these results to identify a particular future schedule to pass on to the *envelope-driven module* which then computes the cumulative energy flow and associated envelope. Other options we are investigating are to consider worst case schedules in terms of energy consumption or most probable schedules assuming additional probabilistic information on uncertain events.

Given the envelope and a desired goal in terms of capacity efficiency, the envelope module computes a new cumulative energy flow satisfying such a requirement. This is done exploiting a linear time complexity algorithm described in (Jiang, Parker, and Shittu 2013). Similarly, aggregation agents will provide energy flows computed in order to meet requests coming from higher levels in the hierarchy. In either case, the system needs to understand if there is a feasible schedule with that cumulative flow and to produce an optimal such schedule in terms of preferences and robustness to uncertainty. We plan to address this problem by starting from the original schedule produced by the constraint-based module and obtaining the required schedule by iterated modifications of the current one.

Aggregation Agents

The computation of a demand envelope is computationally efficient given several forms of information: the actual schedule of planned usage, a cumulative energy flow, and a collection of demand envelopes. Moreover, cumulative energy flows can be combined additively, and thus the upward step of aggregation is relatively straightforward. However, once a change in demand occurs and there must be a reorganization of supply, downward requests to aggregation agents must be incorporated, and decomposed further until individual consumers can make changes to their usage patterns.

Thus, the two key computational steps at each aggregation agent are to: 1) compute a new cumulative energy flow and demand envelope upon receiving a request to change usage, and 2) distribute the change in usage to children aggregation nodes. The aggregation tree will be structured according to the energy market and number of households, but changes to usage will always be sent downward level by level. Once we calculate the new cumulative energy flow at a particular agent, we must then consider the demand envelopes of its children agents, compute change requests (in the form of new cumulative energy flows) and distribute them downward. Given a target cumulative energy flow D, and a set of demand envelopes $\alpha_1, \alpha_2, \ldots, \alpha_k$, we seek to compute a set of modified demand envelopes that respect the target cumulative energy flow D. We outline an exact and a heuristic method that can be used to solve this problem.

Suppose that we are considering D over n time points. Then, for each time point $x \in [1, n]$, we must modify $\alpha_i(x)$ $(1 \leq i \leq n)$ so that $\sum_{i=1}^{n} \alpha_i(x) \leq D(x)$ for all time points x. Since we are starting with a set of envelopes α_i, we can structure this problem as a linear programming problem in which we introduce variables that modify the envelope points to satisfy the above constraints for D. In this basic form, we have a linear program on nk variables and n constraints. To enforce fairness, we can simply add k^2 additional constraints that attempt to balance the change in usage among all pairs of children agents.

Depending on how the aggregation tree is structured, a mathematical programming approach may not be computationally efficient, and we must examine approximate methods. An alternate strategy is to use a greedy approach to computing modified demand envelopes. This algorithm seeks to compute the same weighting factors discussed above, but can do so utilizing a round-robin scheduling approach. In particular, we can consider each child envelope in turn, considering usage peaks (i.e., largest sequential increases in the cumulative energy flows) and adding a request to trim these for each child. Fairness is naturally enforced by the round-robin nature of the computation, and the computation ends when the requirements of the target cumulative energy flow D have been met.

Now, it is important to note that regardless of the approach, at some stage of the aggregation tree it is possible that we have an infeasible problem. That is, neither the exact or heuristic method can compute a set of requests that will meet the requirements of D. There are two ways to deal with this. First, we can initiate a higher demand request to be send upward, and wait for a modified D. Second, we can loosen the restriction to enforce D exactly, in both the exact or heuristic approaches.

The exact approach described above is most feasible when $nk + k^2$ yields a linear program that can be solved on the order of milliseconds. In general, we must enforce a time budget so as to perform all computations within the time necessary to meet a desired time-resolution. Each agent can select between an exact or heuristic approach depending on these calculations. The specific choice of time budget will depend on the desired overall response time of the system (i.e., how quickly we may wish to make changes on the supply side of the grid), as well as the level of the aggregation agent (i.e., if the tree structure is more sparse near the pricing node).

References

Bartak, R.; Morris, R. A.; and Venable, K. B. 2014. *An Introduction to Constraint-based Temporal reasoning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artif. Intell.* 49(1-3):61–95.

Jiang, X.; Parker, G.; and Shittu, E. 2010. A two-sided electricity capacity market with quality-of-services guarantees. Technical report, Tulane University.

Jiang, X.; Parker, G.; and Shittu, E. 2013. Envelope modeling of renewable resource variability and capacity. Technical report, Tulane University, Available at SSRN: http://ssrn.com/abstract=2271177.

Khatib, L.; Morris, P.; Morris, R.; Rossi, F.; Sperduti, A.; and Venable, K. 2007. Solving and learning a tractable class of soft temporal constraints: Theoretical and experimental results. *AI Commun.* 20(3):181–209.

Morris, P. H.; Muscettola, N.; and Vidal, T. 2001. Dynamic control of plans with temporal uncertainty. In Nebel, B., ed., *17th International Joint Conference on Artificial Intelligence (IJCAI 2001)*, 494–502. Morgan Kaufmann.

Muscettola, N. 2002. Computing the envelope for stepwise-constant resource allocations. In *Principles and Practice of Constraint Programming - CP 2002, 8th International Conference, CP 2002*, 139–154. Springer.

Muscettola, N. 2004. Incremental maximum flows for fast envelope computation. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004), June 3-7 2004, Whistler, British Columbia, Canada*, 260–269. AAAI.

Ramchurn, S. D.; Vytelingum, P.; Rogers, A.; and Jennings, N. R. 2011. 10th international conference on autonomous agents and multiagent systems (aamas 2011). 5–12. IFAAMAS.

Rossi, F.; Venable, K. B.; and Yorke-Smith, N. 2006. Uncertainty in soft temporal constraint problems: A general framework and controllability algorithms forthe fuzzy case. *J. Artif. Intell. Res. (JAIR)* 27:617–674.

Vidal, T., and Fargier, H. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. *Journal of Experimental and Theoretical Artificial Intelligence* 11(1):23–45.