CONSERVE: Client Side Intelligent Power Scheduling

Terry L. Zimmerman[1], Stephen F. Smith[2], Atikhun Unahalekhaka[2]

[1]Smart Information Flow Technologies

tzimmerman@sift.info

[2]The Robotics Institute, Carnegie Mellon University

{sfs, atikhun}@cs.cmu.edu

ABSTRACT

We present preliminary work supporting a consumer-side smart grid agent designed to operate autonomously in order to maximize economic reward for its client. The agent must function in an environment of real-time electricity pricing and ubiquitous uncertainty (price profiles fluctuate, client behavior patterns shift) by leveraging advanced scheduling technology that we adapt to residential power management. For those resident loads that can be automatically controlled the scheduling engine will produce advance schedules that respect client preferences and constraints.

In a larger context with multiple such agents within a community, we argue that advance scheduling provides distinct advantages over dispatching systems. Such schedules support collaborative power management between largely self-interested grid agents to realize efficiencies in local power production and consumption.

We report here initial efforts to model and schedule a 'thin slice' of the smart home power management problem, yet one of the more complex aspects; water heating via a heater with advanced control features. We present a model of household water heating patterns together with real-time pricing profiles and demonstrate that an LP approach can rapidly generate cost-optimal solutions to this static problem. Issues and prospects for scaling LP approaches to the larger dynamic problem of interest, as well as trade-offs in adopting less computationally intensive heuristic scheduling methods are discussed.

Categories and Subject Descriptors

I.2.8 [Problem Solving, Control Methods, and Search]: Scheduling

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Algorithms, Performance, Design, Experimentation, Theory

Keywords

Intelligent scheduling, smart grid, linear programming

1. INTRODUCTION

The international effort to modernize and restructure the electric grid has gelled around some core hardware configurations, but there is a range of visions as to operation and control. We focus here on a flexible approach that, while it can support a remotely driven, more centralized power management model in which consumers cede control of loads and distributed energy resources (DER) to a grid management organization, realizes its greatest potential in an environment fostering locally generated solutions to the larger problem of promoting efficient energy production/ consumption. In particular, we report on our progress in implementing a smart grid agent that exploits look-ahead scheduling approaches to implement price responsive demand while maximizing economic reward for consumers

The large scope context and design for the CONSERVE agent (Client Owned Negotiator for Scheduling Energy Resources Very Efficiently) as presented in [1] is rooted in autonomous operation on behalf of consumers, in contrast to utility-side control models for implementing demand response. In brief, the agent is intended to operate on the client side in a context that includes Advanced Metering Infrastructure (AMI) that delivers power pricing signals, 2-way communications, and a variety of residential 'smart appliances' and components that allow some degree of automatic control (Figure 1). The agent fills the role sometimes referred to as an Energy Management System (EMS); Based on the real-time price of power, it is tasked with deciding when to activate/ deactivate those electrical devices that residents agree to cede direct control over in order to reduce costs (maximize profit). The gray panel of Figure 1 envelops the three components of the agent of concern here[1]: Scheduler, Dispatcher, and Distributed State Manager (DSM). All components share the state model of the problem and known state information such as temporal constraints, scheduled activities, client-imposed constraints, etc. The scheduler builds a limited horizon (e.g. 24 hr.), cost minimizing schedule and manages it in the face of environmental change (e.g., new price profiles). The Dispatcher is responsible for triggering energy usage actions according to the extant schedule. Finally the DSM is responsible for communicating relevant aspects of the usage schedule to the grid at large.

This home energy optimization problem is complicated by numerous sources of uncertainty, such as the regional power demand (and thus price), changing wind and insolation (and thus power output from intermittent resources), as well as sporadic changes in residents' routines and behavior. Further complications arise from the presence of consumer's constraints, both "hard" (dishes must be clean by dinner time, the electric vehicle departs on a weekday at 9 AM) and "soft" ('I prefer that the pool water be cleaned at least every other day'). Finally, the great majority of consumers will be unwilling to devote more than brief and cursory attention to electricity consumption in the buildings they inhabit, so an EMS must (re)solve the building power optimization problem in real-time with minimal dependence on human interaction.

We seek to address this problem with scheduling technologies that we have developed for environments featuring high levels of uncertainty, the need to keep pace with execution, a variety of hard and soft constraints, and reward maximization in both

[1] Not shown: the Negotiator/Options Manager that comes into play in collaborations between multiple agents. Cooperation between CONSERVE agents is not addressed in this paper.

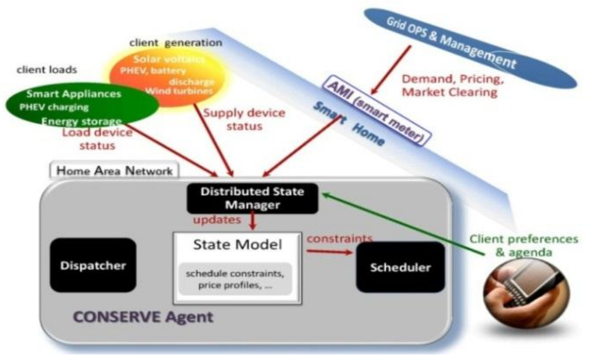

a) Collecting client and environment constraints

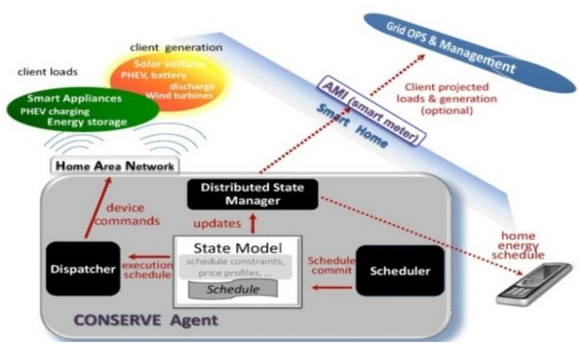

b) Communicating & executing power schedule

Figure 1. Agent architecture and data flow for a smart home client.

centralized and distributed problem-solving contexts [2,3]. To this end, we formulate client-side power scheduling as a dynamic, cost-minimizing optimization problem and anticipate that a robust scheduling approach will seek an optimal solution if conditions permit, and transition to a faster incremental methods [4] when they do not. In this paper we consider a thin but challenging slice of the consumer power management problem typical of a *non-collaborative* smart home environment; residential hot water heating. The feasibility of generating optimal advance heating schedules is explored with LP techniques, to support an effort to determine conditions for which approximate and incremental scheduling techniques may be warranted.

In contrast to advance scheduling, most EMS devices implemented to date are fundamentally based on dispatch methodologies. An electric load under EMS control is activated when the price of power is 'sufficiently low' and no user-specified requirements on the task are violated. Similarly, for some smart devices, they can be temporarily deactivated when the price exceeds some threshold. Such reactive strategies are prone to producing sub-optimal results as they make little or no use of the predictive strengths inherent to advance scheduling. Furthermore, dispatch approaches to energy management offer little support for coordination between clients as the smart grid matures and supports the formation of cooperatives or micro-grids.

We defer this assessment to a future paper, but note here that CONSERVE agents endowed with two-way communication and allowed to cooperate, can leverage advance scheduling in ways that are beneficial to both their clients and overall health of the regional power grid. For example, a CONSERVE agent-arbitrated agreement to purchase power from a local cooperative member can reduce participant costs, reduce transmission costs, and contribute to greater grid stability. From a grid operations perspective, cooperating CONSERVE agents can magnify the effectiveness of real-time pricing in flattening demand peaks.

2. POWER MANAGEMENT PROBLEM

The somewhat simplified general smart home power management problem we describe here can be formulated in a variety of ways; as a constraint satisfaction or optimization problem (CSP, COP), a satisfiability (SAT) problem, or incremental constraint posting search problem, to name a few. We outlined in prior work the appeal of viewing it as a scheduling problem in which discrete energy consumption and generation 'tasks' are to be allocated so as to maximize reward [1].

Figure 2 depicts key constraints and influences for a 24-hour window of a smart home's power-consuming activities (gray box) for which the client cedes control to an EMS. These include any client requirements on start or completion times for those activities, a *projected* load profile of activities outside automated control (non-dispatchable loads), and a day-ahead power pricing profile supplied via the smart meter. If the client owns a distributed energy resource (DER), such as a wind turbine, then the projected output of the turbine based on the wind speed forecast over the day is an important factor. The power activities constitute the decision variables –intuitively they are tasks that a scheduling engine could seek execution times for such that the cost of electric power consumed is minimized. To generate an advance schedule for a smart home requires detailed characterization and classification of all typical daily electric loads. This includes household loads not suitable for automatic dispatch such as on-demand room lighting, televisions, and microwave ovens which are directly controlled by the consumer.

Figure 2 also depicts one feasible solution that minimizes usage cost. Each power task (gray box) is shown as a block with height proportional to its power demand, length proportional to task duration, and allocated within the limits of its specified execution window. Since the client's wind turbine power is conceptually (but not exactly) free, an optimal solver for this situation will seek to schedule all tasks such that the combined power demand from dispatchable and non-dispatchable loads is lower than the projected turbine output at all points over the planning time horizon. Until that capacity is reached, grid power pricing has no

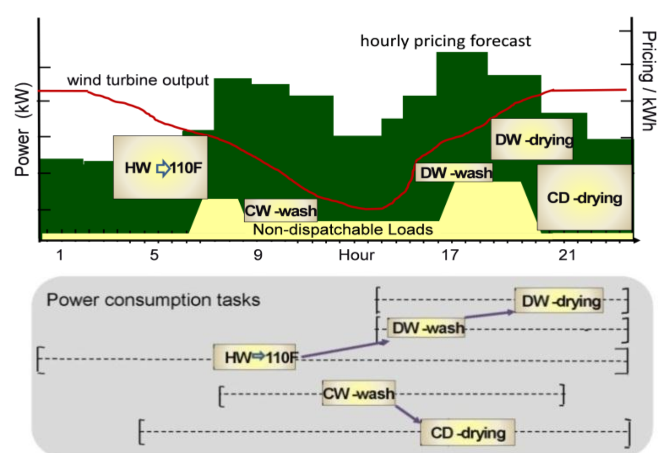

Figure 2. Home energy management: Constraints, environment, and pricing profiles (HW –hot water heater, DW -dishwasher, CW/CD –clothes washer, dryer)

impact on choosing cost-minimizing task start times.

For this study we focus on the water heating power task due to both the size of the impact on consumption (Heating hot water for domestic use typically accounts for 14 – 25% of home energy consumption) and the fact that hot water usage presents a more interesting modeling challenge than other smart appliances coming to market. As described below, the water heater tank can essentially function as an energy storage reservoir that an intelligent control system can use to leverage against forecast or historic electricity pricing profiles.

Key physical aspects of a *conventional* storage water heater include the ready reservoir of hot water and operation by release of hot water from the top of the tank when a hot water tap is opened. Cold make-up water enters the bottom of the tank, ensuring that the tank is always full. Thermostats in the tank control the heater to maintain the tank water temperature within a narrow band. Since the tank contents exceed ambient temp, there is some energy loss (called *standby heat loss*) even when a hot water tap isn't running. The smart water heater we model differs from the conventional primarily in its ability to support super heating of the tank water. We loosely define this as heating above the conventional upper limit of 130 ^{0}F. This requires a thermostatically controlled *water mixing valve* at the hot water outlet of the tank that functions to prevent scalding by sensing the water temperature and mixing enough cold water to achieve the nominal (safe) water heater outlet temperature. Intelligent control of such a heater can realize cost efficiencies by shifting heating cycles to off-peak times via a combination of 'super-heating' when the electricity price is low and 'sub-heating' when it is high.

We cast the power scheduling problem for our experiments as the challenge of finding a minimum cost, day-ahead schedule for cycling an electric water heater subject to the condition that the hot water demand of residents is always satisfactorily met. Specifically, we assume:

1. a day-ahead 24 hour profile/forecasts for power price and household hot water usage
2. a 50 gal. water heater tank
3. water enters tank at T_{in} = 50 °F (10 °C), required hot water tank temp (conventional heater) is T_{max} = 120 °F.
4. client comfort does not degrade for supply temperatures down to T_{out} ~105 ^{0}F, (faucets are manually adjusted).
5. max tank water temp limit (smart heater only) T_{max} = 180 ° A water heating cycle will terminate when temperature reaches max.
6. electric heater power(on) = 11.72 kW (40,000 BTU/hr) and heater is 98% efficient. Heater power is high enough to keep pace with maximum proposed demand.
7. energy required to raise a volume of water by a degree, (*regardless of* T_{in}): .00245 kWh/ gal -°F x .98 (eff) = .0025 kWh/gal °F (1.185 kWh/liter -°C) (first value is the avg heat capacity of water over range T_{in} to T_{out})
8. standby heat loss (E_{loss}) is equivalent to a temperature loss in 50 gallons of water = 2 °F/hr.

The price and usage profiles are not strictly required, but the advantage of advance scheduling becomes more pronounced as forecasts or historical data are made available (and improve in accuracy). We adopt a day-ahead 24 hour power price profile that is representative of typical forecasts published by the PJM Interconnect[2]. We also adopt a hot water usage profile that is loosely based on several sources characterizing household hot water consumption; highest demand occurs in the morning as residents rise and in the evening as they return home, prepare the evening meal, etc (Figure 3). In practice, resident-specific usage patterns should be learned by a companion computational process.

This model ignores several physical realities; seasonal variations in input water temperature, the variance of water's heat capacity with temperature, and heat loss from the tank increases with temperature difference between tank water and ambient air temperature. These simplifications do not significantly alter our results over the temperature ranges of interest.

Optimal advance scheduling The problem can now be restated as: "Find a minimum cost day-ahead schedule for cycling an electric water heater given 24-hr forecasts of power pricing and hot water usage, and subject to the condition that the tank outlet temperature never fails below the lower limit nor exceeds the specified upper limit. Here the lower limit is 105 ^{0}F, and the upper limit is 120 ^{0}F for the conventional heater and 180 ^{0}F for the smart heater (measured at inlet to mixing valve)."

Intelligent dispatch For comparison purposes we also consider power management similar to many current EMS system strategies. Here the heater control logic can be stated as: 1) Heater ON when the electricity price is below the day's projected median and tank temp does not exceed the high limit (180 ^{0}F) 2) Heater ON when tank temp falls below the lower limit (105 ^{0}F) regardless of electricity price. 3) Heater OFF when the price of electricity is above the specified threshold, subject to #2 above.

3. ENERGY MODELING ALTERNATIVES

We experimented with different ways of encoding of the water heating process as a problem for the Cplex LP solver. An intuitive approach is to model the actual tank water temperature as a function of the heater energy input, hot water usage, the resulting inflow of cold water makeup, and the standby heat loss. This proved adequate for quickly solving a problem featuring a conventional water heater, but introducing the water mixing valve for the smart heater model leads to a problematic set of LP constraints. The difficulty arises due to the iterations needed to determine the amount of super-heated water required to meet a demand when it must first be mixed with an unspecified amount of cold water to reduce the mix temperature to the acceptable band (105 - 120 °F). This quantity, in turn, determines how much cold water at the inlet must now be heated to the target tank temperature. We found that even when accepting several additional simplifications and precision losses, optimal solutions to the smart heater problem could not be found for this encoding within several hours of run time.

We side-stepped this difficulty by adopting, instead, an energy balance model of the tank. (The water tank is conceptually a rechargeable energy storage device.) At any given time tank energy content, E_t, is defined as the net tank water energy above the energy content of an equivalent volume of water at T_{in}:

$$E_t = E_{init} - E_{sub} - E_{loss} + E_{add} \qquad \textit{where}$$

E_{init} is the energy stored in tank at start of analysis cycle

E_{sub} is the energy withdrawn from tank due to hot water usage

[2] http://www.pjm.com/markets-and-operations/energy/day-ahead.aspx

E_{loss} is the standby energy lost from tank due to heat transfer to environment through tank wall

E_{add} is the energy added to tank by heating cycles

Under this model, hot water consumption/demand can be translated into energy demand (E_{sub}). Hot water demand is expressed as gal/min but this is translated into energy demand by assuming a constant energy content value for each gallon of hot water: the energy contained in 1 gallon of water at the current bulk tank temperature after it's heated from the tank inlet temperature. This assumption is reasonable as long as assumption 6 above concerning the heater's power rating holds.

We found that the LP solver efficiently generated cost-optimal solutions for this energy balance constraint model in roughly a minute for the smart heater case.

4. RESULTS AND DISCUSSION

The top two plots of Figure 3 depict the chosen day-ahead power price profile and smart home hot water demand forecast for a 24 hour period (x-axis gives minutes since midnight). The bulk tank temperature for a dispatcher-controlled water heater (Section 2 logic) appears in the third plot while the tank temperature for the cost-minimizing scheduling approach is shown in the bottom plot.

As expected, for the optimal the heater is run primarily during the intervals of lowest cost power to raise bulk tank temperature (energy) to the limit in anticipation of future demand. The exceptions are associated with high water demand intervals where the heater is run in order to keep tank temperature above the lower limit, at times based on the *projected* temperature going forward.

The intelligent dispatch approach does reasonably well by comparison except in its over-eagerness to reheat the tank to the maximum after the morning high water demand period. Here it suffers from reheating at a higher power price than if it had delayed an hour or so. We note that the dispatch mode performance is a bit overstated due to the simplifying assumption about constant standby heat loss. Once we augment the water heater model to reflect the fact that heat loss increases with higher bulk tank temperature, there will be a penalty associated with maintaining high water temperatures unnecessarily.

Before numerically comparing the 24-hr heating costs of the two power management approaches we compensate for an artifact visible at the end-of-day in the temperature plots. Given the somewhat artificial 24 hour limit in the demand forecast it schedules against, the optimal scheduling model was somewhat constrained to finally return to the beginning-of-day temperature. The dispatch model, on the other hand reheats tank contents to the maximum in the late evening due to its second cycling heuristic. However, even if we adjust for this end-of-day reheating artifact the 24 hour energy cost of the dispatch approach is \$2.25 versus \$1.80 for the optimal scheduling solution.

5. CONCLUSIONS AND FUTURE WORK

We have summarized the design of a client-side agent capable of operating autonomously to minimize smart home power costs. In support of the advance scheduling approach at its core, we developed the physical models and constraints needed to generate cost-optimal schedules for controlling a smart water heater and reported results investigating the feasibility of generating optimal day-ahead solutions. A modest advantage over a representative intelligent dispatch mode is demonstrated even in this non-collaborative static analysis. After extending the physical models

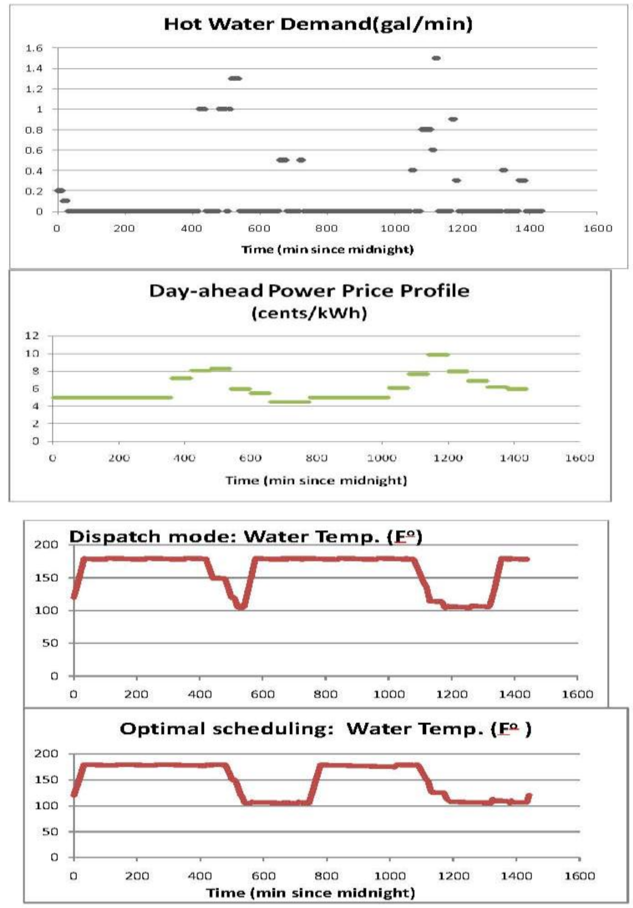

Figure 4. Static solutions for smart water heater under pricing and hot water demand forecasts:

to a complete set of smart home devices that are both more representative and more realistic, we plan to investigate strategies for leveraging the CONSERVE agent power management schedules in local multi-agent cooperatives.

6. REFERENCES

[1] Zimmerman, T. and Smith, S. 2010. "Intelligent Power Scheduling on the Client Side, Design of a Smart Grid Agent", 6th Annual Carnegie Mellon Conference on the Electricity Industry, March 9-10, 2010. (http://www.ece.cmu.edu/~electricityconference/2010/presentations.html)

[2] Gallagher, A., Zimmerman, T. and Smith, S. 2006 "Incremental scheduling to maximize quality in a dynamic environment". In Proceedings of the 16th International Conference on Automated Planning and Scheduling , 2006.

[3] Smith, S., Gallagher, A., Zimmerman, T., Barbulescu, L., and Rubinstein, Z., 2007. "Distributed management of flexible times schedules." In Proceedings 6th International Joint Conference on Autonomous Agents and Multi-Agent Systems, Honolulu, 2007.

[4] Policella, N., S.F. Smith, A. Cesta and A. Oddi, "Generating Robust Schedules through Temporal Flexibility", Proceedings 14th International Conference on Automated Planning and Scheduling, Whistler BC, June, 2004.