

Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration

Jay Taneja

Department of Computer Science

University of California, Berkeley

taneja@cs.berkeley.edu

Abstract

Faced with an uncertain path toward renewables portfolio standard (RPS) goals and the high cost of energy storage, we believe that deep demand side management must be a central strategy to achieve widespread penetration of renewable energy sources. We model loads such as refrigerators and heaters that can be dispatched to match their consumption to these sources. Using *slack*, a generalized measure of dispatchability of energy, we examine a number of techniques for source-following loads as well as the effects of aggregating appliance populations. We intend to develop a machine-learning controller to dynamically manage the electricity consumption of building electricity loads to increase renewables in the generation mix.

1 Introduction and Background

A regulatory push for renewables to play a larger role in the mix of energy sources is underway in many states across the U.S. Already, 36 states have set goals for a Renewables Portfolio Standard (RPS) of anywhere from 10 to 25% of total energy consumed [2] and California has called for renewables to comprise 33% of its energy mix by 2020 [3]. However, as the fraction of energy from renewable, but non-dispatchable, sources like wind and solar increases, it will become more difficult to match supply and demand because today's loads are largely oblivious to supply variations.

Beyond existing demand response programs [1], we propose *deep demand response* – a distributed conjoining of energy information with physical control systems enabled by pervasive sensor/actuator networks – to dynamically match supply and demand down to the appliance-level. The two keys to matching supply and demand in real-time are predicting the output of renewable sources and controlling the consumption of loads in response to these dynamic predictions. The challenge lies in meeting the quality-of-service

requirements of the load while adapting to variations in the source. Though electricity prediction from renewable sources is improving significantly [7], the inherent variability of non-dispatchable sources still presents a barrier to their integration into the energy mix. In order to deal with this unavoidable variability, we envision a future where energy loads adapt to surpluses and shortages on the electricity grid. Specifically, we aim to address the question:

What are the methods, models, and systems necessary to create networks of supply-following energy loads that are responsive to signals from the grid, meet user quality-of-service requirements, and scale from collections of appliances to entire buildings?

As an example of a supply-following load, a refrigerator that is attempting to match incoming energy supply could advance the start time of its cooling cycle to consume excess energy, or a temperature setpoint could be slightly increased to reduce energy consumption. In the first case, pre-cooling changes the phase of subsequent cooling cycles, while in the second case, changing the temperature setpoint (for a cycle or two) reduces the energy consumption. Analogs of these strategies apply to building heating, ventilation, and air conditioning systems, as well as to job-based appliances where operation could be delayed or advanced such as dishwashers, clothes washers, clothes dryers, and electric car chargers.

In our existing work on this topic [15], we develop a metric called *slack* to continuously measure the amount of dispatchable energy to enable system controllers to make quantitative decisions based on existing conditions. Intuitively, slack refers to the amount of time an energy-consuming operation can be advanced or delayed. In this work, we examine a number of energy sinks and quantify the slack available in each system without violating the constraints that govern their operation (e.g. the refrigerator maintains a temperature within an acceptable guardband around a setpoint). By developing an appropriate metric, we then can examine and exploit what is effectively energy storage *already available* within a system to respond to grid fluctuations, allowing collections of supply-following loads to enhance the integration of renewables into the electricity generation portfolio.

2 Integration Results

In this section, we highlight some of the results of our experiments using slack to match energy sources and loads. For our experiments, we deployed networks of motes with

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

climate sensors¹ as well as electricity sensors [10]. Climate data were collected every 10 seconds and electricity data every second using an IPv6-compatible networking layer; a basestation laptop maintains a database of all samples. Having the sensors on a network was important for correlating events between sensors as well as automating the gathering of sensor data.

Experiments are performed by creating an empirical model of the temperature response of the system to compressor operation. We examine both a refrigerator as well as a house heating system. In the case of the heating system, the temperature data is collected by networked embedded sensors and a complementary electricity trace for the heater is synthesized. Note that the magnitude of power consumption is sized for only the fan of the furnace – gas heat is assumed. If an electric furnace was assumed, consumption would be significantly higher. Further details are available in [15].

Though the primary purpose of exploring slack in energy loads is not for energy conservation but for advancing or deferring operations to better match energy consumption and generation, we also note that flexibility (and thus, slack) can be increased by softening guardband constraints.

In Figure 1, we show three different control scenarios of a house heater, with the rows representing temperature, power consumption, slack energy, and a two-day profile of wind energy from a wind farm in Minnesota. The gray area (lightly shaded) on the wind energy plots represents the maximum amount of “portfolio” or non-renewable energy drawn throughout the two-day period. The wind supply generation is the same for all three scenarios.

The first column represents an oblivious heater, which only considers its measured temperature when deciding if the compressor should be actuated. As such, it maintains a tight guardband and consistent slack energy profile. All of the wind plots have been scaled to provide 50% of the total energy needed to power this oblivious heater.

The second column (“supercool”) represents a heater that allows the temperature to cool beyond its normal lower bound by 2°C when faced with a renewable energy shortage. The threshold for an energy shortage, represented by the maximum power consumption of the heater when operating, is indicated by the dashed line on the supply power graph in Figure 1. Thus, when wind energy is available, the heater operates in its standard tight guardband, consuming more energy.

The third column (“wide guardband”) takes this concept further – this heater not only reduces power consumption when faced with a deficit of renewable energy, it also attempts to increase power when faced with a renewable energy surplus by running the heater for longer. In addition, to reduce the discretization of its response, this fridge scales the change in its guardband boundaries by the magnitude of the deficit or surplus of renewable energy. For example, in the late evening of the first day, wind energy far exceeds the threshold, the heater increases its guardband relatively more as compared to the early evening period, when wind

¹We used Telos motes with Sensirion SHT15 temperature/relative humidity sensors and a Hamamatsu S1087 photodiode.

energy barely exceeds the threshold. This change results in less energy saved than the “supercool” case because the proportional response is always less than or equal to 2°C.

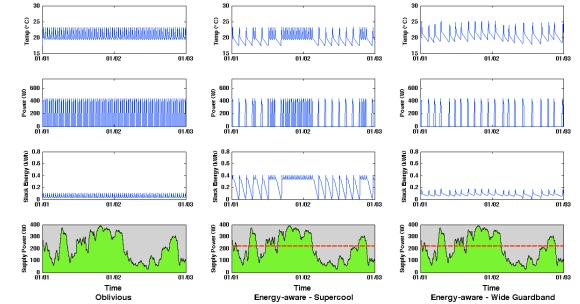


Figure 1. Three scenarios of heater operation. The heater in the first column is oblivious to outside information, the second allows the house to cool during an energy deficit, and the third allows the house to cool during an energy deficit as well as warm during a surplus.

What these scenarios show is that the heater can reduce the frequency of its cycles, and thus its energy consumption, when faced with an energy deficit. A natural follow-up question is whether house occupants are amenable to a wider guardband at certain times – this has been addressed widely in the building comfort and demand response literature. [9]

What is the breakdown of energy sources in these scenarios? Figure 2 compares the three; the first plot shows total energy consumption while the second considers the balance between renewable and portfolio energy. The ability to preferentially loosen guardband constraints saves significant energy – over 33% by reducing only the lower guardband boundary and over 50% while altering both the lower and upper boundaries, even with responses proportional to the energy deficit or surplus. The increase in energy savings from relaxing the lower boundary is intuitive, but the large additional reduction from increasing the upper boundary may not be – it results from lengthened actuation cycles (heating phases, in this case), which ensure a lower overall energy consumption than the oblivious scenario.

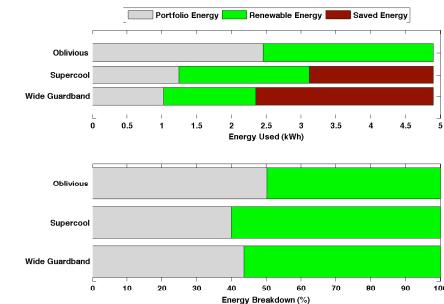


Figure 2. Energy breakdown of three heater scenarios.

We also examine aggregation of the responses of a population of source-following loads. In Figure 3, we show the proportional change between renewable and portfolio energy as an aggregated population of loads increases. In this scenario, each load is a fridge that is initialized at a random phase with warming, cooling, and electrical behavior ran-

domly chosen from our empirical model distribution. Additionally, fridge guardband boundaries are scaled to the magnitude of the energy deficits and surpluses, just as in the “wide guardband” heater described previously. Further, each time a guardband boundary is approached, its value is increased or decreased by a random value within $\pm 0.2^\circ\text{C}$. This small randomness reduces the “herd effects” when a number of agents with identical control logic respond to the same stimulus, as well as partially captures the variation among a large population of appliances, the measurement error of sensors, and local climatic differences. Additionally, the wind energy supplied in this experiment is scaled to satisfy 25% of the energy needed by an “oblivious” fridge. This restriction highlights the benefits of statistical multiplexing in a population.

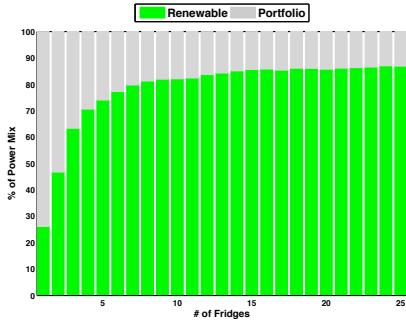


Figure 3. Effects of population size on the proportion of total energy that is renewable. Variations of fridge period and phase allows for cooperation without direct communication, making better use of renewable resources.

Looking at the results of the population study, the proportion of energy used by the fridges improves rapidly as the number of fridges increases, leveling off around 85%. With such limited wind energy supplied per fridge – only 25% – there is significant improvement as each fridge is added. This is because fridges with slightly different guardbands and varied warming and cooling behavior operate with slightly different periods and phases, allowing individual fridges to use excess renewable energy unused by the population. As the population grows, the collection of power spikes resulting from the compression phases of each fridge aggregate into a flatter curve that begins to resemble the wind supply curve, reducing the overall difference between the consumption and the supply. We feel that this result coupled with the amplified variation among an actual population of fridges and similar gains from adapting other energy loads to be source-following forebears the potential for a substantial substitution of renewable energy in place of portfolio energy.

3 Related Work

There is significant literature dealing with wind integration into the electricity grid. Much of this work looks at Europe where penetration is beyond 5% in at least 5 countries, aided by less diffuse population centers and higher overall density as compared to the USA. [4] Experience has shown that beyond 4% penetration, “sudden drops in wind output require some users to accept their electricity being cut off, often by pre-arrangement, or rolling blackouts may result.” [8]

Though grid-scale energy storage is a potential panacea for variability problems from all non-dispatchable sources, existing forms of energy storage do not have the scale or economics for widespread deployment. [14, 13]

The techniques described in this proposal advocate employing slack as widely distributed storage at the points of electricity consumption, an idea that is prevalent in past work as well. For example, Ma, *et al.* describe a system that uses thermal energy stored in water chillers on a college campus [12], Whirlpool analyzes different appliances for their potential to be demand-responsive [11], and Google uses the smart charging of batteries in plug-in hybrid electric vehicles to match an incoming wind signal as well as perform frequency regulation service for the grid [6]. Further, a generalized treatment of the issues with incorporating demand response into appliances is provided in [5].

4 Future Work

Going forward, we will construct empirical models of physical spaces as well as deploy thermostats capable of remote actuation. This combination of monitoring and control will allow us to exploit slack empirically rather than in a model. To construct the controller for a population of appliances and thermostats, we are looking into the promising machine learning technique of reinforcement learning, where a control agent aims to maximize a reward function by observing and inferring the transition matrices between the states in a system. An adaptive system is appropriate in the face of variable weather, energy supply, and energy cost.

Additionally, we will examine the implications of directly incorporating energy storage (e.g. a battery) into a system. This allows an additional degree of freedom where a system can decide when to consume electricity from the grid, when to charge local storage, and when to operate from local storage, enabling exploration of widespread arbitrage of electricity price variability.

Further, with an eye on making models of entire building consumption, we will examine the schedule of large-scale building loads with slack such as chillers and air handlers. We believe there is significant potential in reducing peak load and better matching variable electricity supply by more intelligently managing when these sizable loads operate.

Our existing work only touches the surface of what demand-side management using sensor/actuator networks can do to incorporate energy loads. We believe that a combination of deep demand response through modeling and empirical experimentation will help this technology to achieve its promise and enable many new avenues of research in the building energy space.

References

- [1] Automated Demand Response Program. <http://www.auto-dr.com/>.
- [2] Database of State Incentives for Renewables and Efficiency. <http://www.dsireusa.org>.
- [3] State of California Executive Order S-21-09. <http://gov.ca.gov/executive-order/13269>.
- [4] E. W. E. Association. Wind Energy - The Facts. <http://www.wind-energy-the-facts.org/en/>, 2009.

- [5] Association of Home Appliance Manufacturers. Smart Grid White Paper – The Home Appliance Industry’s Principles and Requirements for Achieving a Widely Accepted Smart Grid. <http://www.aham.org/ht/a/GetDocumentAction/i/44191>, December 2009.
- [6] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl. Demand dispatch: Using real-time control of demand to help balance generation and load. *Power and Energy Magazine, IEEE*, 8(3):20–29, May/June 2010.
- [7] W. Grant, D. Edelson, J. Dumas, J. Zack, M. Ahlstrom, J. Kehler, P. Storck, J. Lerner, K. Parks, and C. Finley. Change in the air. *Power and Energy Magazine, IEEE*, 7(6):47–58, November/December 2009.
- [8] K. Hawkins. Why Industrial Wind Power? Limitations to Wind Penetration. <http://whitherindustrialwindpower.wordpress.com/1-limitations-to-wind-penetration-preliminary-version/>, 2009.
- [9] T. Hoyt, K. H. Lee, H. Zhang, E. Arens, and T. Webster. Energy savings from extended air temperature set-points and reductions in room air mixing. In *Proceedings of the International Conference on Environmental Ergonomics 2009*, 2009.
- [10] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler. Design and implementation of a high-fidelity ac metering network. In *The 8th ACM/IEEE International Conference on Information Processing in Sensor Networks Track on Sensor Platforms, Tools, and Design Methods (IPSN/SPOTS 09)*, April 2009.
- [11] T. J. Lui, W. Stirling, and H. O. Marcy. Get smart. *Power and Energy Magazine, IEEE*, 8(3):66–78, May/June 2010.
- [12] Y. Ma, F. Borrelli, B. Hencey, A. Packard, and S. Bortoff. Model predictive control of thermal energy storage in building cooling systems. In *Proceedings of the 48th IEEE Conference on Decision and Control*, 2009.
- [13] D. MacKay. *Sustainable Energy – Without the Hot Air*. <http://www.withouthotair.com>, 2009.
- [14] B. Roberts. Capturing grid power. *Power and Energy Magazine, IEEE*, 7(4):32–41, July/August 2009.
- [15] J. Taneja, D. Culler, and P. Dutta. Towards cooperative grids: Sensor/actuator networks for renewables integration. In *The 1st IEEE International Conference on Smart Grid Communications (SmartGridComm10)*, October 2010.

Biography

Jay Taneja is a Ph.D. candidate advised by Professor David Culler at the University of California, Berkeley. His research interests include applications of networked sensing and embedded systems. He received his M.S. in Computer Science from Berkeley in 2007 and his B.S. in Electrical and Computer Engineering from The Ohio State University in 2005. He expects to finish his dissertation in 2011.