

Direct Adaptive Control of Electricity Demand

S. Keshav and C. Rosenberg

School of Computer Science and Department of Electrical and Computer Engineering
University of Waterloo, Waterloo, ON, Canada
CS Technical Report CS-2010-17

ABSTRACT

The legacy electrical grid upper-bounds a customer's energy demand using a circuit breaker. The breaker is conservatively sized so that it rarely 'trips,' interrupting supply and inconveniencing the customer. This results in a system whose peak load can be much higher than its average load. Such a system can be made reliable only by being provisioned, at great cost, for infrequent peaks. The goal of demand management is to reduce the peak load and the resulting infrastructure provisioning cost. The most widely used such approach is to price a unit of energy higher during peak times, giving an incentive to consumers to reduce their demand during these times. Although this does reduce the peak load by 2-10%, this 'human in the loop' approach is onerous and frequently ineffective. Instead, we propose a radically different approach that draws upon the Internet traffic management approaches of proactive and reactive congestion control for fine-grained management of customer demand without human intervention. We show that this direct adaptive control of electricity demand can significantly reduce the peak load and therefore infrastructure costs. Moreover, it allows individual users to place a greater load than in the existing system, as long as it does not hurt other customers or the grid itself, thus increasing revenues. We believe that our approach has the potential to avert grid congestion, reduce capital costs, and eliminate a portion of the carbon footprint of the grid.

1. INTRODUCTION

An electrical grid supplies reliable power to residential, industrial, and commercial customers by dynamically matching the amount of power generated by energy sources to varying demands [11]. In an attempt to control the maximum load from a customer (measured in terms of Amperes of current drawn at a fixed voltage of typically either 120 or 230 Volts), the demand placed by a customer is strictly upper-bounded using a circuit breaker or fuse. If a customer's load exceeds the breaker's rating, the breaker trips, interrupting supply and inconveniencing the customer. Thus, the breaker is conservatively sized so that it rarely trips.

This conservative sizing causes two significant problems. First, customers who obey this limit could still contribute to

overloading the system if their demand happens to be correlated with that of their neighbours. Therefore, generation capacity is currently provisioned, at great cost, for infrequent peaks. For example, in 2009, in Massachusetts, 15% of the generation capacity was used less than 88 hours per year [7]. Second, it is inflexible: customers who could benefit from placing load greater than their breaker rating on a lightly loaded grid are prevented from doing so.

A widely-deployed approach to demand management is to raise the price of a unit of energy during peak times, incentivizing customers to reduce their peak load. This 'human in the loop' approach is onerous and a recent study concludes that it can only reduce peak demand in the US by 2-4.45% in 2010 and 6.9-9.6% in 2030 [6].

Instead, we propose a radically different approach that directly controls load without human intervention. Building on the infrastructure already being widely deployed as part of the 'smart grid,' our solution combines fine-grained demand scheduling with techniques inspired by the Internet traffic management approaches of proactive and reactive congestion control to significantly reduce the peak load and hence the provisioned capacity of the grid [9]. Moreover, it allows users to place a greater load on the grid than the existing system, as long as it does not hurt other customers or the grid itself, increasing revenues. We believe that our approach, therefore, has the potential to avoid grid congestion, increase revenues, reduce capital costs, and eliminate a portion of the carbon footprint of the grid.

Our main contributions are:

- The design of a demand management architecture for a 'smart grid' based on direct adaptive control;
- Proactive and reactive control schemes to reduce the peak load;
- A preliminary analysis of two aspects the performance of our system showing significant gains from our approach.

The rest of the paper is laid out as follows. Section 2 provides the necessary background to this work. Section 3 describes the architecture of our solution and Sections 4 and 5 present the proactive and reactive control schemes, respectively. We analyse two aspects of the performance of our

system in Section 6 and conclude in Section 7 with a discussion of open research issues.

2. BACKGROUND AND SYSTEM MODEL

The electrical grid in most countries consists of a mesh-like core (the *transmission* network) interconnecting a few hundred electrical energy generators (the *generation* system) with hundreds or thousands of geographically compact *distribution* networks [11]. Distribution networks, in turn, provide electrical energy to residential, industrial, and commercial consumers. Our work focuses primarily on residential customers, although it can be applied to all customer segments.

In a typical distribution network, a few tens of residential customers are connected on a single *lateral line* to a pole-top transformer. A circuit breaker is used to conservatively limit the load that can be generated from a single customer. For instance, in Ontario, Canada, older homes are restricted to 100A and newer homes are restricted to 200A. Moreover, the total current load placed on the lateral is itself limited using another circuit breaker to protect the pole-top transformer.

Let $L_i(t)$, $i = 1, 2, \dots, N$, be the load of the i th residential, commercial, or industrial customer on a lateral at time t . A circuit breaker limits $L_i(t)$ so that $L_i(t) \leq L_i^{max}$, $\forall t$. Moreover, a circuit breaker set to the value of T^{max} protects the transformer by guaranteeing that $\sum_{i=1}^N L_i(t) \leq T^{max}$, $\forall t$. T^{max} is a sub-linear function of N ¹. For example, in Ontario², for $N = 2$, $T^{max} = 208.3A$ corresponding to a T^{max}/N value of 104.2A and for $N = 20$, $T^{max} = 1391A$ corresponding to a T^{max}/N value of 69.6A. This reflects an underlying assumption that the peak loads from each house rarely coincide. In other words, the existing system is already provisioned for a value smaller than the peak load, though, as we will show in Section 6 this system size can be further reduced by relying on direct, adaptive control.

A customer's electrical load can be roughly partitioned into two portions. The *base* load is the load from always-on devices such as set-top boxes, safety lighting, sump pumps etc. This is typically fairly low. Demand sharply increases when heavy-load devices such as air-conditioners, refrigerators, electric ovens, and baseboard heaters are turned on.

From the perspective of a transformer, loads on a lateral line exhibit a high degree of correlation at longer time scales and are essentially uncorrelated at short times scales of one-to-five minutes. The correlation at long time scales is due to shared seasonal and diurnal effects: the load is high on a hot summer day due to air-conditioning, which causes all the houses on the lateral to increase their electrical demand.

¹In practice, the utility assigns each customer with a weight corresponding to their expected load. In Ontario, small houses are given a weight of 1 or 2 and large houses with electric baseboard heating can have a weight as large as 7 [8]. T^{max} is actually a sub-linear function of this weighted sum rather than N .

²For this analysis, we have assumed that all houses have an equal weight of 3.5.

Similarly, the load from all the houses on a lateral typically reduces at night.

At short time scales, however, the load both within a home and from a set of homes is uncorrelated because of random turning on and off of heavy-load devices. We model an appliance as an on-off source which is either off or on and using electrical power at its full capacity. The ratio of the on time to the sum of on and off times is called the *duty cycle*. Assuming we can neglect the base load, the load from both a single home and from a set of homes is the superposition of these on-off loads. We will make the reasonable assumption that, over a short time scale, the performance of an appliance depends only on its duty cycle. Therefore, over a period of five to ten minutes, it should not matter during which portion(s) of this period an appliance is turned on or off, or on how many times it is turned on or off, as long as the duty cycle does not change.

2.1 Goals

The existing grid has two top-level goals:

Reliability Grid operators typically aim to meet a reliability³ target of 99.999% independent of customer behavior.

Revenue Grid operators would like to maximize revenue by encouraging customers to use as much electricity as they want, subject to the reliability constraint.

These goals are met by provisioning the system conservatively at close to the peak load and charging customers for the resulting (high) costs of the necessary infrastructure that the grid operators deploy. This maximizes revenues while maintaining reliability. Most grids are regulated monopolies, so their operators have had no incentive to reduce capital spending that can be recouped from a large and captive customer base.

The main problem with this approach is that it does not take externalities into account. Specifically, the generation of electricity from coal, which was used to generate two-thirds of all electricity in the US in 2008, causes air pollution, radioactive emissions, and an enormous carbon footprint. With the looming danger from anthropogenic climate change, there is tremendous political pressure to decommission coal plants, reducing built capacity. For example, in Ontario, the provincial government has decreed that all coal-fired plants must close by 2014. Thus, to the two goals above, we need to add the following third goal:

Reduce carbon footprint Grid operators should minimize their carbon footprint.

The conservative peak-load provisioning and inflexible demand management in legacy grids makes it difficult to reduce the dependence on coal-powered plants. With the expected proliferation of plug-in hybrid electric vehicles (PHEVs),

³In the context of the grid, 'reliability' means that every customer's demand is fully met.

the electricity demand at homes and businesses is only likely to increase. We now demonstrate that by intelligent demand management based on dynamic, adaptive, and direct control, an operator can increase revenue and decrease the carbon footprint without affecting reliability.

2.2 Solution Approaches

Before presenting our solution, we consider the management of customer demand using pricing. The general idea of peak-load (or congestion) pricing is to charge more for electrical power during time periods when demand is at a correlated peak. Peak-load pricing has been widely used in the context of telecommunications, transportation, and electrical networks. However, when used as a demand management system for the electrical grid, it suffers from three major problems.

First, it puts the human in the control loop. Unfortunately, humans are time-constrained and do not like having to time-shift their load. Moreover, they are *unable* to control demand at a fine time-scale, which, as we show below, can greatly reduce the peak load by negatively correlating peak loads from different houses.

Second, variable pricing—such as peak-load pricing—is nearly always viewed more negatively than fixed pricing: it irks customers who are unable to determine their anticipated monthly bill. Hence, they might be reluctant to participate in such a scheme.

Finally, some demand cannot be time-shifted. In such cases, peak-load pricing does increase utility revenue, but it does not reduce customer demand, which is the purpose of the control scheme in the first place.

Instead, we propose a radically different approach that draws upon the Internet traffic management approaches of proactive and reactive congestion control for fine-grained management of customer demand without human intervention. In the Internet, data sources control their transmission rate in response to network conditions to avert network congestion. With proactive control, a fixed rate transmitter send a signaling (control) message called a reservation request to the infrastructure and receives an admission control decision that either fully or partially grants this request. With reactive control, a source modulates its instantaneous demand in response to implicit or explicit congestion signals.

By analogy, we propose both proactive and reactive schemes to control the electricity demand of a customer [9]. We first describe the system architecture (Section 3). We then describe proactive control in Section 4 and reactive control in Section 5.

3. ARCHITECTURE

Building on existing architectures for the ‘Smart Grid’ [4], our system has four elements (Figure 1): (a) a ‘smart meter,’ (b) a substation controller (SSC), (c) a ‘smart home controller’ (SHC), and (d) a ‘smart circuit breaker.’

A ‘smart meter’ is an embedded computer owned and man-

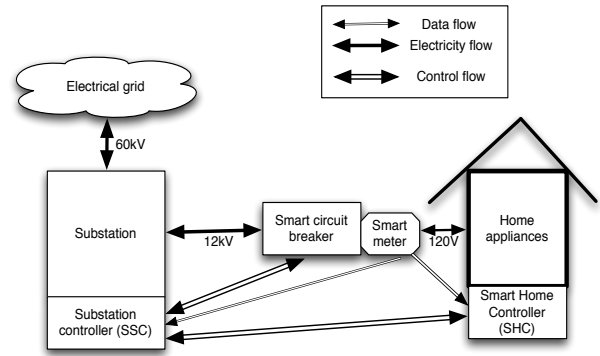


Figure 1: Overview of our architecture. The SSC exchanges control messages with the SHC and the smart circuit breaker.

aged by the utility and installed on customer premises that measures the customer’s load at a fine time scale of one to five minutes and reports it back to the utility. A typical smart meter is the one from Trilliant corporation: 1.4 million such meters will be deployed in Ontario by the end of 2010. This meter communicates to the utility using the ANSI C12.22 protocol [1] over a wireless mesh ZigBee connection from the meter to a pole-mounted base station, which backhauls the data over a dedicated WiMax channel to a substation controller (SSC) at the local substation.

The SSC aggregates measurements and reports them to the billing infrastructure. It is also in charge of proactively or reactively scheduling individual loads, determining the overall network congestion state, and directing the smart circuit breaker to enforce a particular load schedule.

We posit the deployment, in each home, of a Smart Home Controller or SHC that can not only read load data, but can also *directly* control, without human intervention, heavy-load appliances such as air conditioners, baseboard heaters, electric dryers, and refrigerators. A SHC can be implemented either as an application running on a smart meter, or as an extension to home gateways or set-top boxes that are currently being developed by Smart Home initiatives in the industry [2, 3]. In addition, SHC may communicate with an I/O unit to receive commands and display system status. Note that the SHC has complete freedom in deciding how to obey instructions from the SSC. This decoupling of functionality should allow independent innovation in each domain.

Finally, we assume the existence of a per-customer ‘smart circuit breaker’ that corresponds to a policer in the Internet IntServ model. This generalization of a standard circuit breaker is under the direct control of the SSC and limits the peak demand from a customer to protect the infrastructure from excess demand. In addition, in some cases, it may report to the SSC the degree to which the load from a customer exceeds a specific threshold.

4. PROACTIVE CONTROL

Proactive control is carried out once every scheduling pe-

riod of about five to ten minutes. It involves four phases: (a) monitoring and prediction, (b) reservation request, (c) TDM scheduling, and (d) enforcement. We will make the assumption that the SSC and all the SHCs and smart circuit breakers are tightly time-synchronized using an appropriate time-synchronization protocol.

4.1 Monitoring and prediction

In this phase, SHC monitors the electrical demand in the home. Measurements may be of the aggregate demand (which may be obtained from the smart meter) or of individual demand from heavy-load appliances obtained from per-appliance sensors. The SHC then predicts the energy load in the next time period. Predictions can be simple—such as predicting that the demand in the next period is the same as in the current one—or complex—such as automatically learning that the load increases or decreases discontinuously at particular times. Specific prediction algorithms promise to be a rich area of future work.

4.2 Reservation request

In this phase, the SHC sends the SSC a reservation request that contains a desired load profile. This profile represents the superposition of on-off appliance loads in the home as a set of tuples of the form (l_i, d_i) that indicates that appliance i requests a load of l_i amps for a duration of d_i minutes over the scheduling period. For example, consider a system where the scheduling period is 10 minutes. Suppose that at a particular home, for the next time period, the base load is 10A, the air conditioner draws 40A and should be on for at least 3 minutes, the refrigerator draws 15A and should be on for at least 4 minutes. Then, the corresponding reservation request would contain the desired load profile $\{(10,10), (40,3), (15, 4)\}$.

4.3 Scheduling and admission control

The SSC responds to each SHC with a load profile that negatively correlates their peaks so that their sum has the lowest-possible peak-to-average ratio. This has the desirable effect of reducing the need for peak-load provisioning in the electrical grid. Recall our assumption that the duration of the requested load d_i can be broken up into a set of smaller sub-durations d_i^k as long as the duty cycle of an appliance is preserved. Thus, load profile returned to an SHC is in the form of a set of tuples (l_i, d_i^k, o_i^k) , which tells the SHC that it can use the load of l_i for durations d_i^k starting at offsets of o_i^k from the start of the scheduling time period. For example, if the SSC were to get the load profiles $\{(10,10), (40,3)\}$ and $\{(15,10), (40,2)\}$ from two SHCs, it could respond to them with the profiles $\{(10,10,0), (40,3,0)\}$ and $\{(15,10,0), (40,2,3)\}$ respectively. This would reduce the peak load from 105A (in case all the loads are on simultaneously) to only 65A.

The general scheduling problem at the SSC can be posed as the following constrained optimization problem: Given a

set of tuples (l_i, d_i) corresponding to the loads from appliances in a set of homes, find the optimal vector of offsets \vec{o}^* that minimizes the peak load, which is given by

$$\vec{o}^* = \min_{\vec{o}, k} \max_t \sum_i \sum_k \delta_i^k(t) l_i$$

where $\delta_i^k(t) = 1$ in the interval $[o_i^k, o_i^k + d_i^k]$ and 0 elsewhere; and where $\sum_k d_i^k = d_i$. Such problems have been extensively studied in the algorithmic literature and we plan to draw on these solutions in our future work.

It is possible that the peak load, even after optimal scheduling, exceeds the system capacity. In this case, the SSC must reduce the load profiles it grants to each SHC. For example, consider an SSC that gets the desired load profiles (10,10), (40,10) and (15,10), (40,10) from two SHCs. Suppose that the lateral line shared by these SHCs cannot support a load of more than 65A. Then, the SSC could respond to the SHCs with the load profiles (10,10,0), (40,5,0) and (15,10,0), (40,5,5). Reducing the desired load introduces complex issues of fairness, pricing, and gaming which promise to be fruitful areas for future study.

4.4 Enforcement

An SHC is expected to control a home's load so that it conforms the load profile received from the SSC. This reduces to turning the heavy-load appliances on and off at the prescribed times. This can be accomplished using X10- or zWave-controlled power strips.

The interesting problem is when the SHC does not conform to its load profile, because it is malfunctioning or because it has been tampered with. In this case, the SSC can enforce a load profile using the smart circuit breaker. We envisage either a flexible or a strict enforcement action depending on the congestion state of the grid. If the grid is uncongested, then it probably has enough spare capacity to absorb the excess load. In this case, the utility could simply charge extra for out-of-profile load measured by the smart circuit breaker. This would allow a home to obtain excess energy while simultaneously increasing operator revenue. On the other hand, if the grid were close to congestion, it could trip the smart circuit breaker when the load from a home did not meet its profile. This would protect the grid at the expense of inconveniencing the customer. Determining the congestion state of the grid and appropriate enforcement actions will be studied in future work.

5. REACTIVE CONTROL

In the Internet, reactive control allows a source to self-adjust its transmission rate in response to network congestion detected either implicitly (e.g., retransmission timeout or an increase in RTT) or explicitly (e.g., EFCN). Based on the fact that transient electrical overloads are essentially harmless [5], we propose the use of explicit congestion notification in the smart grid as follows: Appliances turn on

(or are manually turned on) as they wish. The SSC immediately senses the increase in load. If this load would cause the system to overload, the SSC sends an explicit congestion notification to the SHC, which then reduces its load. A myriad of demand management schemes can be posed and analyzed in the context of this broad framework. For example, on receiving a congestion notification, the SHC could simply turn off the last appliance that turned on. Alternatively, if it knew which appliances have priority over others, it could reschedule the use of other appliances such that the load was reduced to its prior level. For example, consider a home where an oven is turned on at the same time as the air conditioner, causing the overall load to increase beyond a sustainable level. If so, the SSC sends the SHC a congestion notification signal. The SHC, on receiving this signal, may determine that the oven has priority over the air-conditioner, and could therefore turn off the air-conditioner to reduce the overall load to a sustainable level, or schedule them more carefully.

Other possibilities arise for appliances whose duty cycle can be varied. In this case, the initial increase in load could be cautious, with a small duty cycle, and, if this did not elicit a congestion notification, the load could be progressively increased. This is akin to the standard AIMD approach for window-based flow control. Alternatively, the SHC could begin with a higher duty cycle, then exponentially backoff the duty cycle with each congestion notification. The first approach is well suited to a network that is close to congestion, whereas the second is likely to work well for a relatively uncongested network.

Timing has a critical role to play in reactive control. If a home imposes an unsustainable load, the congestion notification signal must be delivered and acted upon quickly. According to IEEE standard 1547, ‘islanding’ to disconnect an improperly operating subsystem should take no more than 2 seconds [5]. We believe that it is eminently feasible to build reactive control schemes that react within this time.

Reactive control has the appealing property that it does not require onerous monitoring and prediction. Moreover, if the system is mostly underloaded, then customers would rarely receive congestion notifications, and therefore would not even be aware of the existence of the control system. Therefore, we anticipate that the study of reactive control in the grid to be a rich area for future work.

6. ANALYSIS

We now analyse two aspects of our proposed architecture: what is the expected reduction in the peak load due to proactive control, and the extent to which a grid operator can pass on its savings to its customers. We have made many simplifying assumptions to ease the analysis: we intend to remove these in future work.

6.1 Reduction in the peak load

Consider a system with N homogeneous homes, each with

N		t/T			
		0.01	0.10	0.90	0.95
10	Uncontrolled	3	6	10	10
	Controlled	1	1	9	10
100	Uncontrolled	7	25	100	100
	Controlled	1	10	90	95
500	Uncontrolled	17	81	476	493
	Controlled	5	50	450	475
1000	Uncontrolled	26	143	938	977
	Controlled	10	100	900	950

Table 1: System sizing in units of per-home peak load as a function of the number of homes (N) and the mean activity duration at a home (t/T) for the controlled and the uncontrolled case. For simplicity, $b=0$ and $p=1$. For example, when $N=100$, if the mean activity duration is 0.9 , then in the uncontrolled case the system size is 100 units and in the controlled case, the system size is 90 units, resulting in a gain of 10% .

the same load profile, and with the same scheduling period T . For simplicity, assume that the load profiles are of the form (b, T) , (p, t) , where the base load b has the duration T and the peak load p has the duration t .

For the purpose of our discussion, we can ignore the base load Nb . If the loads are only upper-limited by a circuit breaker, we can model each home as a Bernoulli process with the probability t/T of being in the peak state. Because the homes are uncorrelated in the short time scale, the aggregate load is given by a binomial distribution with parameters Np and t/T . To guarantee a reliability of ‘five nines’ or 99.999% , we need to size the system so that the cumulative area under the aggregate load distribution to the right of the chosen system size is smaller than 0.00001 of the total area in the positive X axis (and has a size of at least $Nb + p$).

With proactive control and proper TDM scheduling, the SSC can guarantee 100% reliability because the system is deterministic. We upper-bound the peak load of the system as follows: note that in an interval of length T , the scheduler can schedule up to $\lfloor T/t \rfloor$ homes with a net load of only p . First, consider the case when the peak durations are small compared to the scheduling interval, so that $\lfloor T/t \rfloor \gg 1$. In this case, the peak load is upper-bounded by $Nb + \max(p, \frac{Np}{\lfloor T/t \rfloor})$. On the other hand, when the peak durations are a significant fraction of the scheduling interval, we can compute the gain from scheduling as follows: define $s = \lceil \frac{T}{T-t} \rceil$. Note that if we align the schedules of s homes so that their off periods are consecutive, then the net load from these homes declines from ps to $p(s-1)$. There are N/s such sets of homes, so that the peak load is bounded by $Nb + (\max(p, Np(1 - \frac{1}{\lceil \frac{T}{T-t} \rceil}))$.

This analysis allows us to compute Table 1, which shows the system sizing for 99.999% or greater reliability for some representative values of the parameters. Note that the controlled system always has a lower peak load than the uncontrolled system. Figure 2 illustrates the fractional reduction in the peak for a larger set of parameters. The peak can be reduced by up to 90% : the gain is most for smaller values




Figure 2: Reduction in the peak load as a function of N and t/T . The peak is reduced by as much as 90%. The gain is the most for small values of t/T .

of t/T . This strongly suggests that our approach to reducing the peak load allows a grid operator to realize significant reductions in infrastructure sizing and cost. We therefore propose to analyze more complex heterogeneous systems in future work.

We note in passing that the analysis for reactive control is substantially more complex, requiring us to compute the expected state of each home. This can be done by modeling each home as a Markov chain. Because each home is independent over the short time scale, the joint distribution can be obtained simply as the product of the marginals. We intend to study this further in future work.

6.2 Customer Payback

Demand management reduces the peak load and therefore grid operator costs. The cost reduction comes at no expense to customers if their load profiles are unmodified or if they never receive explicit congestion signals. If not, it comes at the expense of customer disutility. Therefore, demand management will either be legislated or a grid operator would have to compensate its customers by making a side payment to them. Here, we present a rough and ready estimate of how large this side payment could be.

It is estimated that the capital, depreciation, and operational cost of a coal plant amounts to about \$35.67/MWh [10]. Therefore, a 1GW coal plant that produces 8760GWh/yr of electricity costs approximately \$312M/yr. Given that a typical North America circuit breaker is set to 100A, we estimate that the typical North American household has a peak load of approximately 50A or 6KW. This means that a 1GW plant could support about 167,000 homes, resulting in a per-home operating cost of \$1874/year. Therefore, if the peak were reduced by 5% due to demand management, this could result in a side payment of some fraction of roughly \$95/year. With more stringent demand management, it may be possi-

ble to reduce the peak by, say, 15%, resulting in a side payment of up to \$280/home/year. We believe this makes the use of direct demand management potentially interesting to the typical home owner.

7. OPEN ISSUES

We have already remarked on several topics that need far greater consideration than we have afforded here. We summarize a few other areas that also require additional work:

- Our analysis would greatly benefit by the measurement and analysis of customer load profiles that are already being measured by smart meters.
- How should proactive and reactive control be modified for homes with local storage, for example, in form of a plugin hybrid electric vehicle (PHEV)?
- Should scheduling be done only at the SSC or as a cooperative decision jointly by the SHC and the SSC?
- At what time scale should the policer operate at?
- What would be the architecture and operation of a smart home gateway?
- What are the protocols for communication between SSC, SHC, and the smart circuit breaker?
- How can we model the effect of storage resources at a substation?
- What would be the effect of distributed energy generation on demand management?

8. REFERENCES

- [1] An Overview of ANSI C12.22. www.electricenergyonline.com/?page=show_article&mag=18&article=1388.
- [2] Digital Living Network Alliance. www.dlna.org/home/.
- [3] UPnP Forum. www.upnp.org/.
- [4] International electrotechnical commission, technical committee 57. www.iec.ch/dyn/www/f?p=102:7:0::::FSP_LANG_ID,FSP_ORG_ID:25,1273,2010.
- [5] T. Basso and R. DeBlasio. IEEE 1547 series of standards: interconnection issues. *IEEE Transactions on Power Electronics*, 19(5):1159–1162, 2004.
- [6] Electric Power Research Institute. Assessment of Achievable Potential from Energy Efficiency and Demand Response Programs in the U.S. mydocs.epri.com/docs/public/00000000001018363.pdf, 2009.
- [7] P. Giudice. Our energy future and smart grid communication, Testimony before the FCC Field Hearing on Energy and Environment. www.broadband.gov/fieldevents/fh_energy_environment/giudice.pdf, 2009.
- [8] Hydro One, Ontario, Canada. Transformer Sizing at Hydro One. Personal communication from Mr. B. Singh, 2009.
- [9] S. Keshav and C. Rosenberg. How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid. In *Proc. ACM SIGCOMM Green Networking Workshop*, 2010.
- [10] US Energy Information Administration. Average Power Plant Operating Expenses for Major U.S. Investor-Owned Electric Utilities. www.eia.doe.gov/cneaf/electricity/epa/epat8p2.html, 2010.
- [11] A. Von Meier. *Electric power systems: a conceptual introduction*. Wiley-IEEE Press, 2006.