



Lawrence Berkeley National Laboratory

Fast DR: Controlling Small Loads over the Internet


By: Sila Kiliccote, Steven Lanzisera, Anna Liao, Oren Schetrit and Mary Ann Piette

Lawrence Berkeley National Laboratory

August 2014

Presented at ACEEE Summer Study on Energy Efficiency Buildings

This work was sponsored by ARPA-E through the Highly Dispatchable and Distributed Demand Response for the Integration of Distributed Generation project and in part by the Demand Response Research Center, which is funded by the California Energy Commission (Energy Commission), Public Interest Energy Research (PIER) Program, under Work for Others Contract No. 500-03-026 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Authors would like to thank the City of San Jose and the owners of the residences for participating in the field tests.



Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Fast DR: Controlling Small Loads over the Internet

*Sila Kiliccote, Steven Lanzisera, Anna Liao, Oren Schetrit, Mary Ann Piette
Lawrence Berkeley National Laboratory*

ABSTRACT

Current state-of-the-art demand response systems are expensive due to extreme reliability constraints in telemetry, control, and metering. The typical cost of enabling a site for regulation services is between \$50k and \$80k. These costs are orders of magnitude too high for small loads to participate in high-value grid services. Traditional demand response (DR) schedules events in advance and provides basic on/off control. We developed fast DR measurement and communication technologies, which are scalable cost-effective demand response systems with sufficient reliability for real-time control and monitoring of loads over the internet, and we demonstrated an example system for under \$100. We installed FastDR systems for lighting, motor, and thermostat control along with power metering at a total of 10 sites, including 6 commercial buildings and 4 residential buildings. This demonstration addressed the following important methodological questions: 1) Can we reliably control loads simultaneously and with low latency across multiple sites within a specified response time; 2) How reliable is using the internet or 4G cellular network as the mode of control and metering; and 3) Can existing Smart Meters provide data for real-time telemetry services. Testing shows fast DR is capable of control to response in 4 seconds and loads complete transition in seconds to less than one minute depending on load type. We demonstrated 4s regulation services and synchronized load control, and we successfully met specifications for ancillary services across the US. Telemetry via the 4G cellular network demonstrated 94.5% success rate, and residential Internet demonstrated a 98% success rate.

Background

Deployment of renewable portfolio standards in 29 states in the U.S. increased the focus on the use of demand response (DR) to address four major challenges related to renewable generation penetration: over generation during low-load hours; steep and unpredictable ramps; forecast errors associated with renewable generation; and intra-hour variability of these resources (Kiliccote et al. 2010a). Increased flexibility of demand-side resources and availability of real-time signals from the electricity grid are the key ingredients for successful supply and demand interactions. Automated Demand Response (AutoDR) is machine-to-machine enablement of DR in response to grid signals without a human in the loop. While AutoDR is seen as a critical part of managing the smart grid with high renewable penetration, today's wholesale demand response (DR) systems seem experimental and retail DR systems typically work on slow time scales as open-loop systems to address peak load reduction (Kiliccote et al. 2010b). In contrast, generation resources continually report generator conditions to and accept commands from grid operators to provide real-time closed-loop control and status information of the grid. Any demand-side resource that transacts with the electricity grid, by providing a bid, price and duration, with short

(minutes) or no notification is required to adhere to the same communication requirements as generators. These transactions today happen in the wholesale ancillary services markets.

Ancillary services are support services in the power system and are essential in maintaining power quality and reliability. There are typically two types of ancillary services products that DR participates in: contingency and operating reserves. The specifics of these products are summarized in Table 1. Regulation, the product with the fastest communication requirements in ancillary services markets, allows the system operator to request upward or downward changes in output. Regulation is used to track and balance system wide generator output with system wide load on a sub-minute by sub-minute basis (Kiliccote et al. 2010b). Only in California and Electricity Reliability Council of Texas, regulation is separated into two products; these are regulation up and regulation down. In all other markets, regulation products are symmetric, meaning the generator signs up to deliver as much regulation up as down product.

Table 1. Summary of ancillary services for fast DR participation

Service	Service Description		
	Response Speed	Duration	Market Cycle
Operating Reserves			
Regulating Reserves, or Regulation up/down; AGC; Frequency responsive reserves	Online/Spinning reserve, immediately responsive to Automatic Generation Control (AGC) to allow the Balancing Authority to meet the NERC Real Power Balancing Control Performance.		
	<1 minute; must be able to reach max amount of Reg within 10-30 min	30 min (Real Time); 60 Min (Day Ahead)	Hourly; every 15 minutes looking ahead 2 hours
Load Following or Fast Energy Markets	Similar to regulation but slower. Bridges between the regulation service and the hourly energy markets.		
	~10 Minutes	10 minutes to hours	5 min
Contingency Reserves			
Spinning Reserves	Online Generation, synchronized to the grid, that can increase output immediately in response to a major generator or transmission outage and can reach full output within 10 minutes		
	Instantaneous response; <10 minutes for full output	30 minutes	10 min
Non-Spinning Reserve	Same as spinning reserve, but need not respond immediately; resources can be offline but still must be capable of reaching full bid within the required 10 minutes		
	< 10 minutes	30 minutes	10 min

Resources providing regulation products are certified by the ISOs. ISOs currently have a number of requirements for generators seeking to provide the ancillary service of regulation, including direct communication and control system, which communicate key parameters and allows the generator to respond without operator intervention. These generators must respond to

automatic generation control (AGC) signals to change their operating levels depending upon the service they provide, regulation up or regulation down.

There are specific telemetry requirements, which include remote measurement and communication specifics, stated by each system operator for generators to participate in regulation products. These range from two to six seconds (four seconds in California) granularity in measurement and communication (CAISO 2012).

This paper concentrates on technologies that enable fast DR, which are scalable, cost-effective and fast-acting demand response system with sufficient reliability for real-time control and monitoring of loads over the internet for ancillary services participation of loads. The term “fast DR” is defined as demand-side resources that respond without advanced notification and with fast response time (within minutes to seconds depending on the ancillary services products’ requirements outlined in Table 1). In a project funded by ARPA-E, we demonstrated examples of communication and control systems at a low-cost to enable communication and control of small loads. The control strategies vary based on the end-use and acceptability of the strategy by the consumers. The objectives of the study and demonstration were to answer the following questions:

1. Can we reliably control loads simultaneously and with low latency across multiple sites within a specified response time;
2. How reliable is using the internet or 4G cellular network as the mode of control and metering; and
3. Can existing Smart Meters provide data for real-time telemetry services?

Introduction

The California Independent System Operator (CAISO) provides telemetry specifications for communications between generators and the ISO (CAISO, 2012). These specifications provide a reference case for any demand-side resources that aim to provide generation-like services. The CAISO specifies a specific gateway system and communication must take place over the CAISO’s dedicated Energy Communications Network. Both voice and data communications are required between the generators and CAISO, which is typically done over a single data link with Voice-over-Internet Protocol (VoIP).

There are two relevant DR cases to consider: (1) a large load directly participates, or (2) two or more smaller loads participate through an aggregation service. The aggregation service can be a traditional aggregator or a cloud-based system that aggregates and optimizes resources. Figure 1 shows a conceptual schematic of the different portions of the current telemetry system for real-time demand response using information from the California ISO for direct participation of a large load. Structured wholesale markets around the country have similar ancillary services products and requirements that vary primarily in nomenclature and in timing. This figure shows a load directly connected through the dedicated Energy Communications Network (ECN) as well as loads connected via an intermediary, the aggregator. The timing requirements shown reflect those required for resources providing the ancillary service spinning reserve. The 1 minute rate is increased to 4 seconds for the regulation ancillary service and decreased to 5 minutes or 10 minutes for other services or for wholesale energy generation (e.g., load following).





Figure 1. An overview of the architecture showing communication latency for DR resources with and without aggregation.

For our project, we served as an intermediary between the load being shed and the ISO to provide specific data formats, forecasts, capabilities, etc., regardless of the load size. Case (2) is similar to our field studies and implementation. The CAISO requires meter data to be reported every four seconds to the CAISO. This data should be four seconds interval data if it is a single load, but aggregated loads can provide 1m interval data instead of four seconds data. The value must still be reported every four seconds, and the aggregating load meter data server (ALMDS) should update its load estimate with every report it receives (i.e. if there are two loads that report data 20s apart, the ALMDS will report one value for the first 5 samples or 20 seconds then provide an updated value for the last 10 samples or 40 seconds). An example situation is shown in Figure 2 where an ALMDS is used.



CAISO EMS: The energy management system used to manage the supply demand balance
 ECN: Energy Communications Network, a Internet like entity dedicate dto CAISO communication with generators and loads
 DPG: Data Processing Gateway. A device the provides protocol and data support in compliance with CAISO standards.
 ALMDS: Aggregating Load Meter Data Server. A server that aggregates load information from more than one load

Figure 2. Schematic showing communication latency and reporting rate requirements for demand side aggregated loads, specific to CAISO.

Figure 2 shows the schematics of communication latency and reporting rate requirements for a demand side spinning reserve ancillary services system using several aggregated loads. The telemetry links have some general physical requirements. The loads must communicate with the Data Processing Gateway (DPG) or ALMDS (for single or multiple loads respectively) using a broadband communication medium. There are no requirements on the data rate of this link. The protocol to the right of the DPG in Figure 2 is unspecified and can be selected by the implementer, but the communication from the DPG to the CAISO EMS must use DNP3 (Distributed Network Protocol Version 3). DNP3 is a Supervisory Control and Data Acquisition (SCADA) protocol that is commonly used in the electricity utility industry. The connection between the DPG and the Internet or ECN must be an always available connection (i.e. not dial up), but it does not need to be dedicated to only providing ancillary services.

A variety of data types must be reported depending on the size of the load and the needs of the ISO at that time. For loads less than 10 MW, real power in MW, reactive power in MVAR, Voltage in kV, and connectivity status for each load (if source is an ALMDS). For loads over 10 MW, gross values for real power and reactive power are also required. A gross value is one measured at the terminals of the load itself rather than one measured at the connection point to the grid. The DPG must also provide a “heartbeat” that changes from 0 to 100 at one increment per second. The CAISO specifies a target data reliability of 99.0%. Billing settlement is handled using five-minute interval data, and these data are not provided in real-time to the ISO. Discrepancies between telemetry data and settlement data are reviewed by the ISO, but there is no defined standard for the maximum discrepancy allowed between the settlement and telemetry data. The ISO specifies meters that are allowable for settlement, and it reserves the right to inspect facilities to review how the metering is conducted.

For the remainder of this paper, we discuss communication technology implementation options, including cellular, wired Internet, ISO networks, advanced Metering Infrastructure (AMI) networks, and unidirectional broadcast radio. Next, we describe our implementation in detail and share the load impact results from the Fast-DR tests and evaluate latency and

reliability of Internet connection. Finally, we present conclusions and explain next steps for this research.

Technology Implementation Options

One approach for developing a real-time demand response system would be to think about how one could be built without considering existing requirements or limitations of the ISO or load facility systems. Target reliability above 90% is low for enterprise data communication, and achieving better than 8s round-trip times using standard, wired IT infrastructure is relatively easy.

Using data from (Aikat et al. 2003), we see that roughly 90% of round trip TCP connections using a high quality, wired Internet connection exhibit less than one second in latency. About 99% of links exhibit less than four-second latency. In the four years since these data were collected, performance has improved. Using existing wired LAN technology in buildings, round trip latencies in excess of 100ms represent less than 0.1% of the time. Therefore, a system using standard, wired Internet connectivity from the site to the ISO and standard wired LAN connectivity from the Internet connection directly to the sensing and actuation points would most likely meet the requirements. However, for large commercial buildings with connectivity through energy management and control systems (EMCS), latencies related to control infrastructure need to be considered. If the Internet and associated IT infrastructure installation is not already present at a site, this method can be cost prohibitive for widespread deployment.

A wireless connectivity option provides the needed data link while reducing installation costs. The primary options that exist are cellular data or a wireless local area networks with wired (or wireless) data backhaul to the Internet (either over an existing connection or a dedicated connection for DR). Cellular data appears attractive because it eliminates the need for additional infrastructure in most cases. Wireless LAN connections, such as WiFi systems or wireless sensor networking, are also potential solutions that bring their own challenges and opportunities.

The simplest cellular data is Short Message Service (SMS), but latencies for SMS are often in excess of several seconds per message (Keynote 2003). Cellular communication using 2G,3G, and 4G all perform much better and are candidate technologies. Legacy 2G data typically can provide round trip latencies of less than one second, and 90% of latencies are less than two seconds. Cellular 3G data has typical latencies of 200 ms and 90% latencies of 300 ms (Xu et al. 2011). A key remaining constraint is on the overall availability of the cellular connectivity at the particular location in question. Given that we expect to have a large number of loads at our disposal spread across a wide geographical area, it's likely that the diversity we would achieve in this arrangement will mitigate the connectivity issues we have due to periodic outages of a carrier's service.

This cellular data connection approach is shown in Figure 3. Four connections (A through D) are identified in this drawing, and each has it's own latency issues. We expect, based on available data that links A, B and C are reasonably less than 300ms in latency for connections to an external server. The connection from Demand Response Optimization and Management System (DROMS) to the ISO (connections C and D) has unknown characteristics because we do

not have data on the capacity, latency or other performance aspects of any ISO system. It is highly likely that this connection will meet the four-second requirement for this set of links.




Figure 3. A telemetry solution based on cellular data connections.

Unfortunately, the amount of data that must pass over the cellular data connection is significant. If a device participates in the market on an ongoing basis and receives and transmits a nominal 250 byte (B) packet every four seconds, over 300 megabyte (MB) of data traffic is required each month. This places the cost of data services at about \$30/month for typical business data plans for machine-to-machine communication. This high price is a potential deal breaker for all but moderate to large loads. In California, spinning reserve is worth roughly \$3/MWh and non-spinning reserve is worth roughly \$1.5/MWh (MacDonald et al. 2012). If a load provided spinning reserve 100% of the time, the load would need to have an excess of 15 kW available for real-time DR. A more realistic situation is for the load to only participate a fraction of the time. This would lower the data costs, but the required load size would remain about the same due to the reduced utilization. It is likely that cellular data prices will decrease in the coming years. A roughly 10x decrease will be required before this method is appropriate for typical residential loads and many smaller commercial loads. For these smaller loads, a different connection mechanism is required.

The lowest cost method of providing low latency data connectivity is to use an existing wireless LAN (i.e. WiFi) deployment at a site. Broadband Internet has high penetration in US households today (~70%, Connected Nation 2014), and most of these homes already have a WiFi network (Lanzisera Nordman and Brown 2012). Even adding this to a home for providing ancillary services is a low, one-time cost. Use of an existing business WiFi network (even if an access point must be added to cover some loads) may prove to be a lower life-cycle cost because there are no special, recurring monthly costs. Business data connections are widely considered to be highly reliable with availability above 99%, but home data connection reliability is less well known. Our experience with these and other experimented site indicate that reliability may be less than 90%. Using a wired connection to the home or business Internet connection is more attractive than wireless if it can be done inexpensively, but it is likely this cost will be too high in many cases where wiring is difficult.

It is attractive to consider alternative wireless LAN technology other than WiFi, and the primary examples are Zigbee or other wireless sensor network technologies (e.g. 6LoWPAN-based networks). These networks provide improved coverage because they provide a multi-hop network that is either a tree, partial mesh (Zigbee), or full mesh. The primary issues with these networks are unknown reliability and latency performance in real world networks. Networks that achieve high reliability and battery operation often provide median latencies in excess of a few seconds. Worst case latencies can exceed 30 seconds. Zigbee networks based on all full-function devices will have lower latencies but early results suggest they will also provide lower overall reliability.

The variety of network configurations allows for tailoring a network solution to a specific site. This approach is unattractive, however, because it requires as site evaluation and the use of many hardware configurations. We strive to have a one-size-fits-most approach to our system, where a standard set of tools apply to every site with an emphasis on minimizing customization. It appears that existing wireless LAN connectivity is the best approach under this constraint, but we will continue to evaluate perspective solutions throughout this work.

Field Tests

Lawrence Berkeley National Laboratory (LBNL) has been developing low cost telemetry options to be utilized for an ARPA-E project titled “Highly Dispatchable and Distributed Demand Response for the Integration of Distributed Generation” led by AutoGrid. The objective of the project is to design and demonstrate automated control software that helps manage real-time demand for energy across the electrical grid. The system is designed and developed to allow personalized price and control signals to be sent to millions of customers and incentivizing them to change their demand in response to changing grid conditions. This system is envisioned to enable grid operators to better manage supply and demand fluctuations. In this section, we describe the field demonstrations conducted by LBNL that demonstrate low-cost telemetry options to enable fast DR.

Inside each facility participating in the project, we deployed an LBNL-designed local control processor called the “B-Box”. The B-Box is responsible for collecting data from inside the facility, receiving grid information from an OpenADR server, and making local control decisions. Our pilots rely on the Simple Measurement and Actuation Profile (sMAP) to collect and visualize data. sMAP was developed at UC Berkeley and extended for this project. We use a standard sMAP server to provide data archiving and as a common interface for data. A commercial power meter measures the load, and the B-Box retrieves the power data and passes it back to the sMAP server. All of these actions can happen every second and take only a fraction of a second to complete. The B-Box is a simple platform that provides a high-level environment to manipulate data, make control decisions, connect to a sMAP server, and control connected loads. The high level architecture of these systems is shown in Figure 4.




Figure 4. Overview of fast DR deployment at participating sites.

For each small; commercial site installation, power meters were installed at the breaker panel for building mains power and breakers connected to power for the HVAC units. The installed meters are Dent PowerScout 3 Plus meters with serial Modbus communication via Neteon RS-485 gateways and Split-Core current transformers. The thermostats to control these roof top units were replaced with Cloudbeam-enabled Wi-Fi Radio Thermostats. The core of the B-Box was implemented using an embedded computer (eBox-3350MX-AP or BeagleBone Black). The B-Box serves as the interface to stream data (via HTTP posts) from meters to our server located at LBL. For Internet connectivity to devices, a Cradlepoint mobile broadband router serves as a gateway for the wired serial connections from the power meters and wireless connections for the Wi-Fi thermostats. Internet connectivity is enabled on the Cradlepoint gateway with either a Verizon 3G/4G data-card or wired connection into the network located at the site. The cost of site enablement is dominated by the metering electronics. Out of the \$1000 hardware cost to enable these sites, \$600 was for the metering electronics, \$150 was for the Cradlepoint, \$80 for the RS-485 gateway, \$50 for the embedded computer, \$50 for the cellular modem, and miscellaneous costs made up the remainder. Installation costs were typically \$1000.

The residential sites used the PG&E Smart Meter for power measurement, and a Rainforest Automation RFA-Z106 RAVeN USB Zigbee transceiver provided the data link to the meter. A B-Box was deployed in each site with a RAVeN Zigbee device, and the B-Box passed data back to the server at LBL using the existing home Internet connection. Custom software ensures the RAVeN receives a new meter reading every 4s, and faster data rates were not possible with this setup. This setup was quite inexpensive at less than \$100 total.

The large commercial building is made up of offices and classrooms. In this building, there is a building automation system (BAS) that controls the zone temperature setpoints while

the VAVs, fans and lighting system can also be controlled through the use of sMAP. Internet is used to bring grid transaction signals to the premise and within the building; the implementation of Fast-DR is completed through the wired infrastructure.

Table 2 summarizes the field implementations at selected sites and presents the facility size, the Fast-DR strategy the site employed, communication methods and measurement device.

Table 1. Summary of field implementations

Site	Location	Size (ft ²)	Fast DR Strategy	Communication	Measurement
Large Commercial	Berkeley	141,000	Control of VAVs, fans, Lighting w/ a BAS	Internet/wired	Meter/submeter
Small Commercial 1	San Jose	5,900	Control of Tstats (Split system)	Cellular/WiFi	Meter/submeter
Small Commercial 2	San Jose	6,250	Control of Tstats	Cellular/WiFi	Meter/submeter
Small Commercial 3	San Jose	14,400	Control of Tstats	Cellular/WiFi	Meter/submeter
Small Commercial 4	San Jose	9,500	Control of Tstats	Cellular/WiFi	Meter/submeter
Residential 1	Stockton	2,076	VRV limiting - mini split system	Internet/Zigbee	Smart Meter
Residential 2	Stockton	1,690	compressor off/on	Internet/Zigbee	Smart Meter
Residential 3	Stockton	1,140	compressor off/on	Internet/Zigbee	Smart Meter
Residential 4	Oakland	1,400	Direct fan VFD	Internet/Zigbee	Smart Meter

Results

The server, B-Box and the meter are time synchronized and the communication between the server to B-Box or the device is measured as well as the time it took for the response to be visible from the telemetry equipment. The time it took for the server to send a signal to a controller is called “control-to-response” latency and the time it took for the device to respond to a control signal is called “transition” latency. In addition to capturing this data through repeated tests, we also calculated load shed based on an outside air temperature regression baseline (Piette

et al. 2007). Table 3 summarizes the results from the field tests. The results include the average savings in Watts, Watts per square feet and percentage of whole building loads. We also summarize the results from the latency measurements from various communications implementations. In addition, we calculated the cost of telemetry including installation costs.

Table 2. Summary of load impact and latency measurements

Site	Ave. Load Impact			Latency (Control-to-Response)	Latency (Transition)	Cost of Telemetry Enablement
	kW	W/ft ²	%			
Large Commercial	110	0.8	19%	<5 min.	<5 min	\$5000
Small Commercial 1	5.0	0.8	17%	<4 sec.	<1 min.	\$1000
Small Commercial 2	9.4	1.5	52%	<4 sec.	<1 min.	\$1000
Small Commercial 3	14.8	1.0	37%	<4 sec.	<1 min.	\$1000
Small Commercial 4	3.9	0.4	16%	<4sec.	<1 min.	\$1000
Residential 1	5.0	2.4	85%	<4 sec.	<1 min.	\$100
Residential 2	3.0	2.6	80%	<4 sec.	<1 min.	\$100
Residential 3	3.0	2.1	80%	<4sec.	<1 min.	\$100
Residential 4	0.5	0.3	25%	<4 sec.	<8s	\$100

Figure 5 shows the aggregated load profile during a coordinated test event that was executed at the small commercial sites. A variety of loads were dispatched to shed load in a staggered manner to provide a stair step type decrease in aggregated loads. The last load was only off for about 20 minutes while the longest off load was off approximately 45 minutes. Providing coordination across several sites in this way is critical for deploying fast DR at sites with small load. In aggregate these sites can be combined to provide complex load shapes as needed to fill ancillary or other grid service needs.




Figure 5. Coordinated test event across small commercial sites.

Conclusion and Next Steps

In this project, we aimed to answer the following questions:

- Can we reliably control loads simultaneously and with low latency across multiple sites within a specified response time;
- How reliable is using the Internet or 4G cellular network as the mode of control and metering; and
- Can existing Smart Meters provide data for real-time telemetry services?

The field studies with various building types, sizes, equipment and with various connectivity and control options, we can reliably control loads simultaneously and with low latency across multiple sites within a specified response time using the Internet and 4G network for control and metering. In addition, we showed that Smart Meters can provide secure data for real-time telemetry services over the Zigbee radios.

The next steps for the research include evaluating the economics including analyzing value relative to the load shed capability, developing hardware price-point targets and defining target markets for technology transfer. On the technical side, we plan to further evaluate the cost and reliability of end-to-end cellular systems as well as thermostat reliability. Finally, we recommend evaluating flexibility of end use systems and DR control options to reduce response time and increase participation of all available flexible loads year round.

Acknowledgements

This work was sponsored by ARPA-E through the Highly Dispatchable and Distributed Demand Response for the Integration of Distributed Generation project and in part by the Demand Response Research Center, which is funded by the California Energy Commission (Energy Commission), Public Interest Energy Research (PIER) Program, under Work for Others Contract No. 500-03-026 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Authors would like to thank the City of San Jose and the owners of the residences for participating in the field tests.

References

- Aikat, J., Kaur, J., Smith, F., and Jeffay, K. "Variability in TCP round-trip times," In Proceedings of the ACM SIGCOMM Internet Measurement Conference, Miami, FL, October 2003.
- CAISO, "ISO Generation Monitoring and Control requirements for AGC/Regulation Units V4.6",
<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.194.1264&rep=rep1&type=pdf>
accessed April 2012.
- Connected Nation. "2012 Residential Survey" <http://www.connectednation.org/survey-results/residential>. accessed February 2014.

- Keynote 2003. Keynote Warns That Millions of Wireless Text Messages Are Lost. First Ever Study Shows Critical Interoperability Problems; Lost Revenue Potential is Staggering, http://www.keynote.com/company/press_room/releases_2003/01.15.03b.html, accessed April 2012.
- Kiliccote, S., P. Sporborg, I. Sheikh, E. Huffaker, and M. A. Piette, "Integrating Renewable Resources in California and the Role of Automated Demand Response," *Lawrence Berkeley National Laboratory, Berkeley, CA*, 2010. LBNL 4189E. November 2010.
- Kiliccote, S., M. A. Piette, J. Mathieu, and K. Parrish, "Findings from seven years of field performance data for automated demand response in commercial buildings," *Proc. of ACEEE summer study on energy efficiency in buildings. Pacific Grove, CA*, 2010.
- Lanzisera, S., B. Nordman and R. Brown, "Data network equipment energy use and savings potential in buildings," *Energy Efficiency*, Vol. 5, No. 2, pp 149-162, May 2012.
- MacDonald, J., P. Cappers, D. S. Callaway, and S. Kiliccote, "Demand Response Providing Ancillary Services A Comparison of Opportunities and Challenges in the US Wholesale Markets", *Grid-Interop 2012*, Irving, TX, December 2012.
- Piette, M. A., D. S. Watson, N. Motegi, and S. Kiliccote, *Automated Critical Peak Pricing Field Tests: 2006 Pilot Program Description and Results*, LBNL-62218. August 2008.
- Xu, Q., A. Gerber, Z. M. Mao, and J. Pang. 2011. AccuLoc: practical localization of performance measurements in 3G networks. In *Proceedings of the 9th international conference on Mobile systems, applications, and services (MobiSys '11)*. ACM, New York, NY, USA, 183-196.