

Smart-Grid Demand Management Requires Robust Control Systems

ENGINEERS CREATING THE SMART GRID FACE SIGNIFICANT CONTROL-SYSTEM CHALLENGES IN ELECTRICAL GENERATION AND TRANSMISSION LOAD MANAGEMENT.



The Smart Grid can best be seen as a big control system with negative feedback imposed to achieve demand control (see “What Is The Smart Grid Really?” at electronicdesign.com). The objective is demand management to deal with the difficulty of siting new, large-scale electrical generation sites and transmission lines (Fig. 1). The communications technology and software for the application of feedback to the energy consumer has already been developed by companies such as Battelle and Silver Spring Networks (see “Demand Response Today” at electronicdesign.com).

The biggest challenges that remain lie in creating the mechanism that will generate the negative feedback. In North America, the buying and trading of electricity along the path from generator to transcontinental transmission lines to more local distributors to the local electricity retailer is supposed to produce real-time changes in pricing at the consumer level, where presets in, for example, clothes dryers and air conditioners will determine when those devices will turn on or what thermostat settings to follow.

The current control model is a forecast-and-dispatch model. Based on each utility’s own forecast together with its long-term agreements and such, the model determines whether it’s economically advantageous to sell some of what the utility has available or to buy more on the open market. To that extent, this is a completely open loop model. There’s no feedback, merely hunches, based on practical experience.

Making this work during the course of the day is the responsibility of a balancing authority (see “Balancing Authorities,” p. 33). To support the more than 3000 utilities in the United States, there are 131 balancing authorities that monitor power at transmission-station-level meters. The information obtained from this monitoring is called “bulk consumption.”



1. Bonneville Dam is roughly 200 feet high and half a mile long as it spans the Columbia River. Its 20 turbine-powered generators have a combined capacity of 1.189 GW. They literally “don’t build them like that anymore.” There isn’t room. Load management addresses that limitation, but engineers also are dealing with an accumulated legacy of systems and processes that present daunting prospects for creating feedback systems that will work at this scale. (courtesy of the Bonneville Power Administration)

CLASSIC LOAD BALANCING

Operational changes driven by the bulk consumption information are applied to a special class of generators called load-following generators. Turned on, they let the balancing authorities “gen up” a little more capacity. Also, if some unforeseen event occurs, they can provide the resources to bring online at the last minute. Conversely, if consumption is unexpectedly less than forecast demand, they can easily shut down to remove supply from the grid.

This is useful, but there is a very small tolerance within which the grid has enough resilience to simply absorb demand variations. One catastrophic outcome might occur if the balancing authority lost too much of its load. The utilities would be putting out too much power, and equipment might be destroyed.

Today, the source of real-time information for the balancing authorities comprises a series of real-time monitoring systems, called supervisory control and data acquisition (SCADA) systems. The control generally involves a “person in the loop.” Control engineers at the balancing authorities literally pick up the phone and call people, asking their sources for a gen-up for more capacity, or asking managers at factories that have negotiated lower rates to shut down certain operations.

In the U.S., most utilities have standardized on a set of communications protocols called Distributed Network Protocol 3 (DNP3).¹ Three classes of communications equipment use DNP3 for SCADA: master stations, remote terminal units, and intelligent electronic devices (Fig. 2). The Inter-Control Center Communications Protocol (ICCP) is used for inter-master station

communications. ICCP is a part of IEC 60870-6, an international standard.² In fact, IEC 61850 is the standard for substation control in Europe. It is still being assessed for North America.

The Smart Grid won't consist of experts communicating over telephone lines to balance demand, though. Instead, home appliances would turn on or off depending on pennies per kilowatt hour. Adding those kinds of "price responsive loads" would present some problems, according to Rick Geiger, Cisco's executive director of utilities and Smart Grid. When the retail price hits a certain target and "500 MW of load suddenly drops off the grid," he said, the existing control model would be destabilized significantly.

"Something more is called for," he said. "You've got to be able to close that control loop."

An industry team spearheaded by Jeff Taft, chief architect of Cisco's Connected Energy Networks Business Unit, put together "Ultra Large Scale Control Systems for the Grid" last October.³ The ideas in this 25-page document covering the fine-grain detail of generating, transmitting, and distributing electricity, Geiger said, are gaining traction in the industry.

ADVANCED METERING

The "final mile," the two-way link to the energy-user's electric meter, may appear to be the easiest part of the puzzle at first. The Advanced Metering Interface (AMI) is already part of millions of meters. What happens to that information is the part that gets interesting. And it isn't going to be a smooth road ahead.

The effects of government deregulation must be addressed. Also, the utility that delivers the power might not be the company that has the customer relationship. Plus, any of those entities can buy power from multiple power marketers. Just like mortgage derivatives, then, the power marketers will aggregate those retail contracts and make power purchase agreements with generator and distributor operations. There are aggregators in generation and delivery as well.

How is it possible to balance all of those transactions? How is it possible to make sure that what consumer A actually consumed matches what the power marketer issued as a purchase agreement on the generation company? And how do you ensure that the generation company delivered that much power?

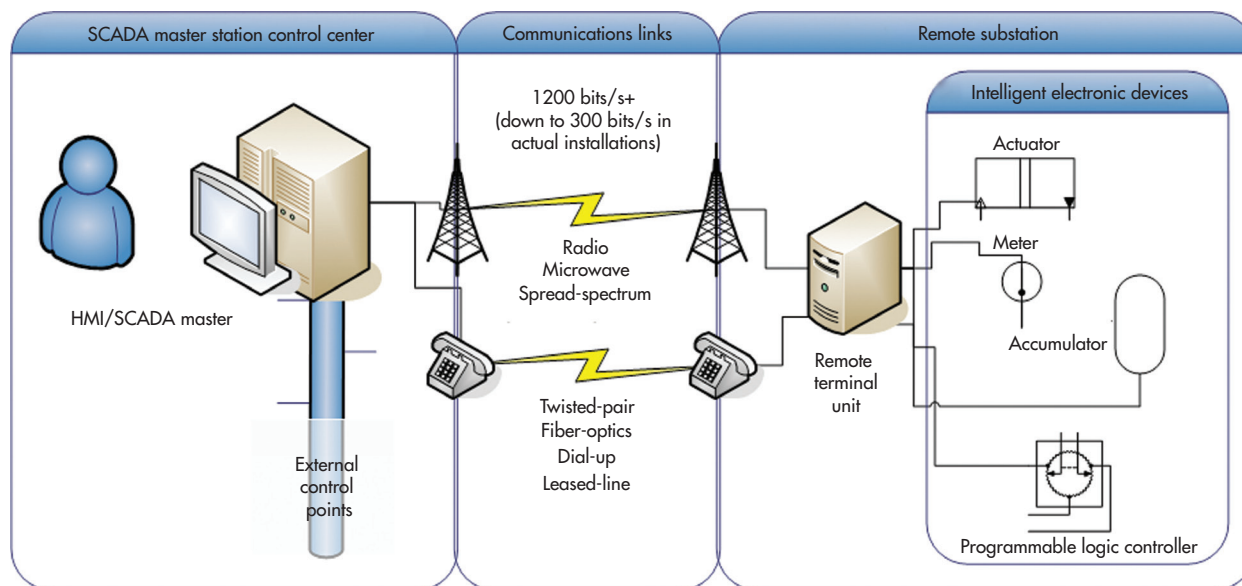
Previously, profiles managed the complexities. A profile is a standard spreadsheet that estimates what a customer's consumption profile looks like during the course of the day. Now that pricing will be used for demand management, it's necessary to break the profile down into more granular data based on data that the utility gets via the AMI backhaul.

The initial people who would be able to take advantage of the variable pricing would seem to be the industrial customers, who are already shedding load in response to a phone call, rather than a consumer with a major appliance. Actually, smart chargers for electric vehicles (EVs) are leading the way. They represent the first domestic load that asks whether it's the right time to charge. Otherwise, industrial electricity users will do the trail-breaking.

Industrial customers are billed differently than residential customers. Depending on what state they live in, residential customers may have either a flat tariff or a graduated tariff. In California, it is a graduated tariff where a certain base number of kilowatt hours per month is at one rate, and there is a higher tiered rate for the next number of kilowatt hours up to, in some places, a fifth tier. (Implemented in parts of the Fresno, Calif., area, that fifth tier created a huge backlash against smart meters in 2009.)⁴ There also are time-of-use rates, which are one of the incentives for businesses and homeowners to install solar generating systems.

As the Fresno protests suggest, ordinary consumers are not enthusiastic supporters of rate complexity. That's why the aim is to pre-load the criteria, as in the case of EV chargers, or at least make setting trip points as simple as possible. That home-user aversion to complexity is why the pioneers of demand control by means of variable pricing are most likely to be industrial users.

Industries purchase power with reference to a "base load," which is a taker payload that they guarantee to the utility. So if they don't actually consume that much, they pay anyway. On top of that, there is a rate for variable consumption and a peak demand rate. If they exceed their contractual peak demand, industries pay very expensive rates for electric power because it stresses the utility's supply and balancing algorithm for that day.



2. Three classes of communications equipment—master stations, remote terminal units, and intelligent electronic devices—use Distributed Network Protocol 3 (DNP3) for supervisory control and data acquisition (SCADA). (courtesy of Wiki Commons)

To avoid those overcharges, industrial companies spend a lot of money on management systems to make sure that they never exceed their demand rate. Eventually, entrepreneurs, working with home appliance makers, may be able to use the same kinds of algorithms to maximize home consumers' rates of return on smart appliances. ("A Sub Zero refrigerator? Sure, there's an app for that!")

Between the home owner and the industrial plant, there is a large middle class of commercial electricity users: office buildings, multi-tenant facilities, retail spaces, and hotels and restaurants that represent a large untapped area because of the variability of their tenants and the inaccessibility of their energy and building control systems. They don't have the control that industrial companies have because of their largely multi-tenant nature.

Privacy also opens opportunities for new designs. Typically, ordinary meters are read once per billing cycle. But smart meters designed to monitor demand in real time are read often enough to reveal a daily pattern of electricity use. Burglars who gain access to those patterns could identify homes that are likely targets for felonious visits when the owner is away—homes where the oven doesn't warm up until after 4 p.m., for example.⁵

In 2011, Toshiba suggested masking fine-grained meter readings with a battery/charger arrangement like a perpetually connected EV battery. While that seemed far-fetched at the time, an organization spearheaded in Canada called Privacy by Design has set forth a model for addressing confidentiality and privacy in metering.

Also, researchers from Pennsylvania State University and the University of British Columbia have developed a battery-based approach they call non-intrusive load leveling (NILL).⁶ The approach is conceptually similar to queue delay perturbation countermeasures that prevent networking timing analysis. By de-correlating the timing and amplitude of on/off events in the load profile, NILL removes the signal that load monitoring algorithms use to identify behavior.

MANAGING POWER OUTAGES

One cannot assume that power will always be available. In fact, the more sophisticated parts of the Smart Grid that have already been implemented use smart meters to report, via radio, where outages are occurring so workers can be dispatched efficiently and full power restoration can be verified before they are dispatched elsewhere.

Consider Silver Spring Networks' Smart Energy Platform, in which smart meters outfitted

BALANCING AUTHORITIES

BALANCING AUTHORITIES CAN get complicated. Take the Balancing Authority of Northern California (BANC), the third largest balancing authority in California, as an example. Its charter calls it a joint powers authority (JPA).

BANC assumed control of the Sacramento Municipal Utility District (SMUD) in 2011. It now embraces SMUD, the Modesto Irrigation District (MID), Roseville Electric, and Redding Electric. It makes decisions con-

cerning the use of about 40% of the Central Valley Project (CVP) hydro facilities, which are owned by the California Bureau of Reclamation.

In addition to generating capacity, BANC contracts for transmission facilities owned by the Western Area Power Administration (WAPA) and the high-voltage (500 kV) California-Oregon Transmission Project (COTP), which is jointly owned by several more parties including WAPA and some BANC members via

the Transmission Agency of Northern California.

BANC cooperates with other balancing authorities, including the California Independent System Operator (CAISO). Standards developed and enforced by the Federal Energy Regulatory Commission (FERC), the North American Electric Reliability Corporation (NERC), and the Western Electricity Coordinating Council (WECC) govern these relationships. ■

with the company's communications modules can communicate real-time outage information back to a utility's operations center (Fig. 3). An outage detection system (ODS) processes the raw, unstructured outage data arriving from the meters and presents it in an easily comprehended fashion. This data can be passed on to a third-party outage management system (OMS), yielding detail and real-time outage status down to individual meters.

The ODS also performs complex event processing to filter the data before it is passed on. For example, it knows when meters are undergoing maintenance or are part of a planned outage, so it won't forward alerts from those meters. It also can filter out momentary outages so crews are only dispatched for sustained customer interruptions.


In monitoring outages, utilities use such information to have a better idea of where the

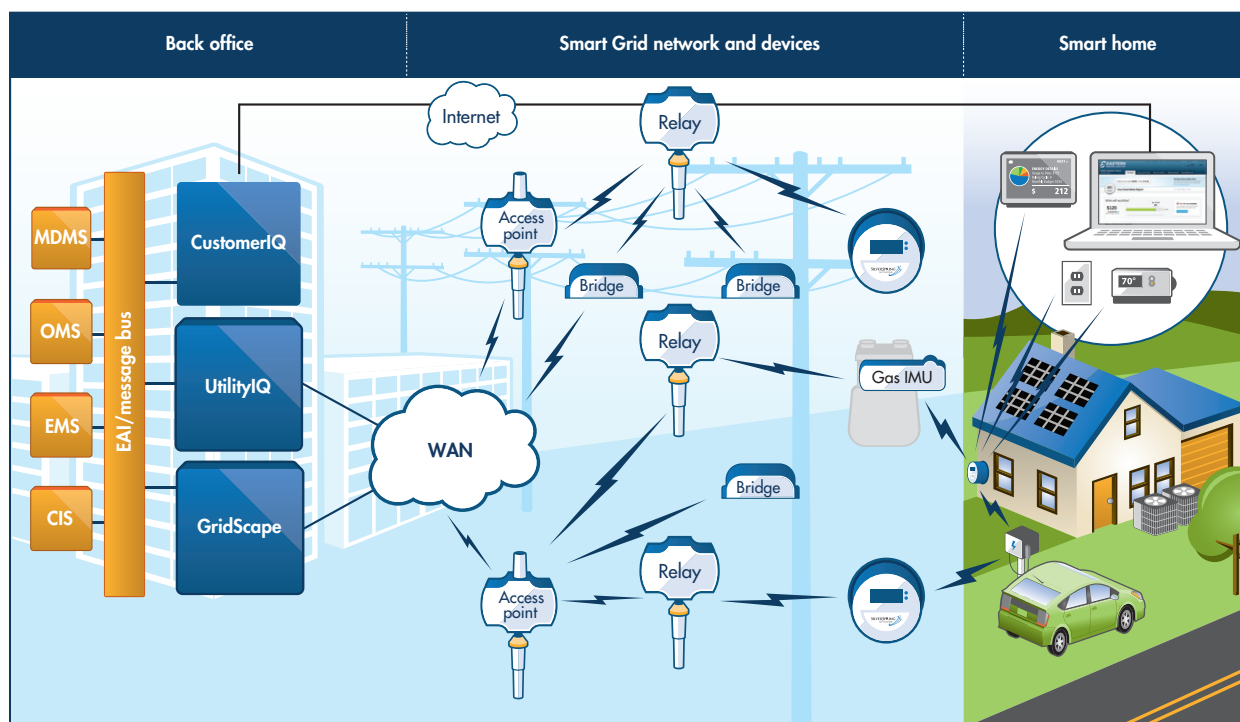
likely common point would be in the distribution system so they know where to look for the fault location. "They might have other sensors that would help further identify not only the location of the fault but also what it is that failed," Geiger said. "Then, based on that, they are able to dispatch a remediation crew and able to reduce the overall duration of the outage based on superior logistics and work-order dispatch."

There also so-called self-healing distribution architectures, which don't depend on smart meters but rather on sensors in the distribution network. "There are numbers of different architectures. There are mesh architectures, loop architectures, but perhaps one of the simplest to explain would be the loop," Geiger said.

"Imagine a power line that is a topologically broken circle fed by the distribution line at both ends. Switches would be spaced around

the loop, and the switch at the top of the loop would normally be open. And if a fault occurred had a fault on either side, it would be possible to open the switches on either side of the fault and open the switch at the top of the loop," he continued. "That way, you could rapidly restore power to as many customers as possible while isolating the fault and making it safe to work on."

Self-healing has acquired something of a misleading reputation, though. It doesn't fix what's broken. It only isolates it. "So if a car hits a power pole, a self-healing distribution grid is not going to replace the pole. You'll still have to dispatch a crew to go do that," Geiger said. "Sometimes you'll see articles that talk about, well, such and such storm came through, and if we'd only had self-healing grids then we'd have been able to restore power more quickly. That's true only in specific instances." 



3. Silver Spring Networks' Smart Energy Platform enables smart meters to communicate real-time outage information back to a utility's operations center. Once there, an outage detection system (ODS) processes the raw, unstructured outage data arriving from the meters and presents it in an easily comprehended fashion. An outage management system (OMS) then can present operations personnel with detailed real-time status information down to the addresses of individual meters.

FURTHER READING

Renewable generation resources also can be added to the Smart Grid, ranging from rooftop solar panels to wind farms. For an overview of today's renewable power generation and distribution issues, Geiger recommends "Integration of Renewable Resources," a 300-page document published in 2007 by the California Independent System Operator (ISO).⁷

References

1. Distributed Network Protocol 3 (DNP3), www.dnp.org/pages/aboutdefault.aspx
2. Inter-Control Center Communications Protocol (ICCP), http://en.wikipedia.org/wiki/IEC_60870-6
3. "Ultra Large Scale Control Systems for the Grid," Cisco's Connected Energy Networks Business Unit, www.cisco.com/web/strategy/docs/energy/control_architecture.pdf
4. "Smart-Meter Rollouts And Standards Stir Controversy," Don Tuite, <http://electronicdesign.com/news/smart-meter-rollouts-and-standards-stir-controversy>
5. "Conferences Highlight Smart Grid Opportunities," Don Tuite, <http://electronicdesign.com/power/conferences-highlight-smart-grid-opportunities>
6. "Protecting Consumer Privacy from Electric Load Monitoring," Stephen McLaughlin, Patrick McDaniel, William Aiello, www.cse.psu.edu/~smclaugh/cse598e-f11/papers/mclaughlin-ccs.pdf
7. "Integration of Renewable Resources," California ISO, www.caiso.com/1ca5/1ca5a7a026270.pdf