Making the Smart Energy Grid Even Smarter

Randal E. Bryant   **Randy H. Katz**   **Erwin P. Gianchandani**
Carnegie Mellon University   **U of California-Berkeley**   **Computing Research Assoc.**

Our nation's infrastructure for generating, transmitting, and distributing electricity — "The Grid" — is a relic based in many respects on century-old technology. It consists of expensive, centralized generation via large plants, and a massive, centrally-controlled transmission and distribution system. It strives to deliver high-quality power to all subscribers simultaneously — no matter what their demand — and must therefore be sized to the peak aggregate demand at each distribution point. Power is transmitted via high voltage lines over long distances, with associated inefficiencies, power losses, and right-of-way costs. Local distribution, via step-down transformers, is expensive in cost and efficiency, and is a single point of failure for an entire neighborhood. The system demands end-to-end synchronization, and it lacks a mechanism for storing ("buffering") energy, thus complicating sharing among grids or independent operation during an "upstream" outage. Recent blackouts demonstrate the existing grid's problems — failures are rare, but spectacular. Average demand per consumer is a small fraction of the peak — a 25 kWhr/day home draws on average less than 5% of its 100 amp service. Consumption correlations, e.g., air conditioners on a hot day, drive demand beyond estimated aggregates, which can result in huge spikes in supply cost and may trigger blackouts. Moreover, the structure cannot accommodate the highly variable nature of renewable energy sources such as solar (generating power only during the day) and wind (generating power only when the wind is strong enough). Meanwhile, consumers are provided little information about their energy usage (just a monthly total) and even fewer opportunities or incentives to adapt their usage to better align their demands to the capabilities of the utility companies.

Many people are pinning their hopes on the "smart grid" — i.e., a more distributed, adaptive, and market-based infrastructure for the generation, distribution, and consumption of electrical energy. This new approach is designed to yield greater efficiency and resilience, while reducing environmental impact, compared to the existing electricity distribution system. Already, the U.S. government is investing billions of dollars in deploying aspects of smart grid technology, primarily "smart" meters and associated communications technology. These meters can reflect real-time prices to consumers, to motivate them to reduce their consumption at times of high demand. But trepidation about trusting such critical infrastructure to unproven technology may limit just how "smart" the smart grid can become in the near-term. In addition, little attention is being paid to more radical approaches, ones that would involve fundamentally new, more decentralized structures for the grid.

Initial plans for the smart grid suggest it will make extensive use of existing information technology (IT). However, *advances* in computer science — and at the intersection of computing and energy — have the potential to greatly enhance the smart grid and, ultimately,

---

greatly amplify its impact. There must be a far closer partnership between computer science researchers and the engineers, scientists, and policymakers already working on smart grid technology in order to fully realize its potential. This partnership would be mutually beneficial: these parties have complementary expertise in building complex distributed systems that are scalable, flexible, and user-friendly in environments that have stringent requirements for reliability, resiliency, efficiency, and security. Basic, collaborative, and interdisciplinary research with a strong IT component is required to achieve the best possible smart grid.

In this document, we consider two different time horizons for visions of a future electricity grid. Our "10-year" horizon takes current concepts for smart grid technology and enhances them to include more aggressive use of information technology. Our "20-year" horizon imagines that we are willing to fundamentally restructure the entire system to use a more decentralized "packet-based" architecture, with intelligence pushed to the network's edges, even to the level of individual loads and sources (i.e., move increasingly toward the loads and supplies, as opposed to top-down from supplies to loads). This vision is inspired by the way in which the Internet has replaced the historic, centralized architecture of the telephone system.

A 10-Year Smart Grid Vision

The first major leap in migrating intelligence beyond smart meters at the edge of buildings is to extend the smart grid into factories, offices, and homes. To illustrate an ambitious vision for the smart grid 10 years from now, Smart Grid 2.0 if you will, consider the following residential energy scenario:

> *The Jones family, of Phoenix, Arizona, lives in a house with state-of-the art sensing and control capabilities. Their home management computer has "learned" the habits and preferences of the family and has an interface directly to the electric utility's computer system. On August 15, 2020, the following interchange takes place between these two systems:*
>
> *The utility company sees that the day's temperature will exceed 110°F. It needs to reduce the peak demand later in the day. It contacts the Jones's home management computer (JC) and requests that it keep the load below 8KW between 4pm and 9pm. JC is aware that the family will be out until 8pm, and it knows (a) it can safely let the household temperature rise to 90°F without harming the family cat or tropical fish, and (b) it takes 30 minutes to bring the temperature back down to the Jones family 's preferred reading of 75°F. Consequently, JC offers to reduce usage to 4KW from 4pm to 7:30pm, then to 12KW from 7:30pm to 8pm and to 6KW from 8pm to 9pm. JC can maintain the latter limit because it anticipates that the Joneses will be returning with their electric car still partially charged; JC will be able to tap into the energy remaining in the car's batteries. It bundles this offer with a requirement that the utility company supply enough energy to fully recharge the car by 7am the next day, and it offers to supply energy to the grid from the household solar panels for the rest of the day. The two systems negotiate an appropriate price for the total energy package.*

While the ideas of using pricing incentives and household automation have been proposed as part of the system for achieving more balanced loads and better efficiencies in the smart grid, the above scenario includes aspects that go much further than other existing plans for smart grid technology:

- The home management computer serves as the "energy czar" for the home. It uses sensors to "learn" the usage patterns and characteristics of the household appliances and lighting (e.g., power draw, ramp-up time), as well as the occupants' habits, needs, and preferences beyond default settings. It can selectively control devices, such as the household thermostat. It can diagnose anomalies (e.g., defective fluorescent lighting ballasts or refrigerator doors left ajar) and report these to the homeowners. It has access to the homeowners' calendar programs and is able to determine when they will be away. Unlike some smart grid proposals that require humans to continually monitor and adjust their energy usage, this example illustrates how the home management computer could function much like a good butler — it would automatically learn the occupants' preferences and make the right choices in the background.
- The home management computer negotiates with the utility company, engaging in complex transactions involving bundled combinations of current and future quotas and prices, as both a consumer and a provider of electricity.
- The utility company's system must perform these negotiations with hundreds of thousands of households, as well as with hundreds of other entities such as power generators, transmission line operators, other utility companies, etc. The system has at its disposal data about weather and seasonal patterns, but it must also consider possible statistical variations and unexpected demands and outages.
- Although we described the above scenario as involving a single negotiation per day, it is likely that utility companies and home computer systems will be engaging in these exchanges more frequently (e.g., hourly).

Much of the technology underlying such a system requires advances in computer science research:

- Modern automobiles contain 20 or more embedded microprocessors and far more embedded sensors, optimizing fuel usage and emissions, maintaining safety systems, and supplying the driver with routing information. Research in the field of *cyber-physical systems* seeks means to implement highly adaptive, reliable, safe, and efficient sensor-based systems. For the case of the smart grid, these cyber-physical systems must scale to millions of loosely-coupled autonomous agents, designed in such a way that their collective behavior achieves very high degrees of reliability and efficiency.
- *Machine learning/data mining* can readily detect usage patterns and preferences automatically. For example, researchers at Georgia Tech and the University of Washington have shown that instrumenting a home with just three sensors — for electric power, water, and gas — makes it possible to determine the resource usage of individual appliances, lights controlled by individual switches, and the HVAC and plumbing systems — simply by analyzing patterns in the data. The sensor and control system can apply machine learning to continually improve energy efficiency, reliability, and comfort by monitoring operations and algorithmically tuning parameters and behaviors, largely

eliminating the need for users to manually set configuration parameters. Realizing this potential requires that sensors and the control system work reliably in all environments and at sufficiently low cost for a consumer marketplace.

- *Agent-based (auction) systems* can negotiate complex contracts on a massive scale. For example, Google performs automated auctions millions of times per day to place paid advertisements within its search result pages. But it's one thing to sell space on Web pages; it's a much more serious proposition to operate a market-based system at this scale that controls critical infrastructure.
- The field of *human-computer interaction (HCI)* has studied how to create computer systems that provide appropriate levels of information and control to users in ways that are understandable and that minimize the chances of error. Many skeptics about the smart grid state that consumers are unlikely to make the effort to play an active role in energy management. We believe that a combination of automation combined with providing people with the right information in understandable form along with the right level of control will be essential to achieving major improvements in efficiency.
- *Advanced optimization* can guide the adoption of renewable energy sources, such as wind and solar, based on projected macro-scale demand, grid capacity with anticipated upgrades, and consideration of the inherent intermittency of renewable power sources. For instance, the optimal location in terms of wind-energy density may not be as desirable as a slightly suboptimal location where projections indicate maximal need; two smaller wind farms on opposite sides of a geographical barrier (e.g., a mountain range) may prove most efficient due to offsetting intermittency, reasonable grid access, and consistency with planned grid upgrades.

A 20-Year Smart Grid Vision

Taking a longer-term and more transformative view, the Internet suggests alternative organizing principles for a 21st Century decentralized peer-to-peer smart grid. The Internet succeeded by pushing intelligence to the network's edges while hiding the diversity of underlying technologies through well-defined interfaces. In other words, any device can be a source or sink of routable traffic, but it's the "intelligent" endpoints that adapt their behavior to what the infrastructure can deliver in accordance with localized utility functions.

Imagine an electric power system built on *packetized energy*: energy is stored where it is generated and conceptually "routed" to where it is needed. The existing infrastructure is generally unable to store energy for later use, yielding a centralized system with crude mechanisms to adapt to changing loads (e.g., regional exchanges, peaker plants, curtailment) and provisioned so that users lack the means to present an easier-to-manage load to the infrastructure. New environmentally friendly energy storage technologies are needed, with capacity/cost metrics suitable for deployment in homes, buildings, and throughout the transmission and generation system, including intermittent energy sources. Combining intelligent communication protocols with energy transmission in a common architecture makes possible distributed control and demand response to pricing signals. Such an infrastructure design would permit a shift from designing for the peak/worst case to the average case with sufficient headroom, analogous to statistical multiplexing in packet networks. Standardized intelligent "interfaces," at the level of homes or even individual appliances, would allow independently powered operation, distributed

generation, and energy exchange. The architecture should allow aggregation to plug into the regional grid, the neighborhood peer-to-peer grid, or the facility grid to use localized storage and control to smooth load, adapt demand, and engage in exchange.

Consider the following scenario for another version of the smart grid, Smart Grid 3.0 if you will, roughly 20 years from now:

> *It is now August 15, 2030, and the Jones, Ortiz, and Garcia families of Phoenix, Arizona are neighbors. Their community has entered into an agreement that permits its members to resolve their energy demands cooperatively, optionally using the resources of the utility's grid only when necessary to make up deficits arising when local demand exceeds local supply. The families have deployed solar panels on their roofs, but their co-op "grid" is augmented with access to an independently operated wind farm on the local mesa with associated energy storage. The latter allows some economies of scale, for larger wind turbines and high-capacity energy storage technologies.*
>
> *During a sunny day, the solar deployments generate sufficient energy to meet the immediate demands of each household. Since the Garcias are on vacation, their generated energy is deposited in the community storage "bank," and they receive a credit to be used at a later time. By contrast, the Ortizs are home, and they are also having a big family reunion with many out-of-town guests. They may draw extra energy from the local storage bank, debiting the credits they have built up in the past.*
>
> *The nighttime breezes allow energy to be extracted from the wind and locally stored. This stored energy can be made available to the community on demand, thus displacing aggregate demand from the utility grid and allowing the utility grid's resources to be deployed elsewhere. When the community is in a situation of energy surplus with respect to demand, it can sell its excess resources to the utility.*
>
> *Finer-grained control and scheduling of loads allows intermittent energy sources like wind and solar to be harnessed even without storage. Consider the Jones' home air conditioner system. On a partly cloudy day, it can opportunistically overcool their home slightly when the sun is shining, exploiting the thermal mass of the home and the time it takes to warm up (i.e., to "schedule" the load). This is a form of "brilliant" load, i.e., load that follows available supply, intelligently making use of inexpensive energy when it is available, even if for short intervals, so as to defer or eliminate demand at a later time when energy could be scarce and expensive. Such decentralized management is only possible with rich awareness of energy availability and local demand, married to the necessary computational modeling to drive such decision-making.*

Radical proposals to replace existing infrastructures, given their wide deployment, high capital costs, and well-understood technologies, are unlikely to succeed. Here, too, the Internet offers a model — of infrastructural *co-existence* and *service displacement*. The early Internet was deployed on top of the telephone network. It provided a more resilient set of organizing principles, became its own infrastructure, and eventually the roles reversed; services such Voice over IP (VoIP) telephony are recent additions, having been added over time. The same approach

can yield a new architecture for local energy generation and distribution that leverages the existing energy grid, but achieves new levels of efficiency and robustness, similar to how the Internet has improved the phone network.

Some Challenges

Fully realizing a future in which millions, perhaps billions, of computerized agents manage our energy production, distribution and consumption requires many capabilities well beyond those of any current system. Moving to a future with packetized energy presents even more fundamental challenges:

- *The systems must adapt to unexpected events.* What if, in our first scenario, the Jones' return home early, or their car returns with less charge in its batteries than was anticipated? Many other events can undo the careful planning made by the utility company: new usage patterns, unexpected weather conditions, failures of components or subsystems, etc. The systems must operate with sufficient capacity margins to avoid failures. The agents must be able to dynamically renegotiate contracts, with appropriate pricing mechanisms to avoid abuse.
- *The system components must be able to cooperate with one another.* For example, shouldn't the Jones' home computer be able to remotely query the Jones' car during the course of the day to assess the precise level of charge anticipated upon the Jones' return, and could the Jones' home computer in turn use this information to renegotiate contracts in "real-time"? This high degree of connectivity and coordination could make the system more reliable; however, if poorly designed, the system could also be more vulnerable to cascading failures leading to large-scale blackouts.
- *The system must guarantee sufficient privacy.* The Jones family might not want the utility company (or a malicious eavesdropper) to know things like when the house is vacant or when the teenage daughter is home alone. Unlike scenarios where utility companies are provided direct control over household appliances, we envision that the home management computer will serve as an "information firewall" to the outside world. It will act on behalf of the homeowners while restricting the flow of information to the outside world. It may even choose to obfuscate externally visible usage patterns, e.g., by having some form of energy storage within the house that can be charged or utilized at different times of day. (As with many other real-world systems, there may be a benefit-cost tradeoff between privacy and efficiency.) Ultimately, these technologies will require computer scientists interfacing with policymakers directly.
- *The system must be resilient to abuse or attack.* Experience with the California energy market in 2000 demonstrated the possibility for companies to "game the system," creating havoc while reaping huge monetary benefits by exploiting flaws in the computerized marketplace. Given the rise in the amount and sophistication of Internet-/cybercrime, there are justifiable fears that malicious agents will target any network-based smart grid both for monetary gain and to disrupt the U.S. economy. An effectively designed, agent-based system can potentially be less vulnerable to manipulation or attack than a centralized, monolithic one, but a bad design could yield just the opposite effect.
- *The system must learn and improve over time.* As new electrical devices become available, how should they be incorporated into the power optimization equation most

efficiently? As usage patterns evolve with households becoming more energy conscious, or as families evolve (e.g., children are born, or children go off to college) and so do their energy consumption patterns, the underlying machine learning must track and adapt, both to long-term lifestyle changes and to transient ones (e.g., a family goes on a two-week vacation, or workmen remodel a home and their power tools draw substantial electricity for a short time).

- *Efficient energy storage and conversion are required to realize packetized energy.* Current technologies for storing useful quantities of electrical energy require converting it to a different form, e.g., chemical (in batteries), compressed air, or gravitational (pumping water into raised tanks), and then converting it back into electricity. All existing technologies have poor efficiency and high cost. Indeed, this lack of efficient storage is the fundamental reason our electric power system must operate without any buffering, and there are many research efforts to find better alternatives. Sending a packet of energy over a network styled after the Internet requires that the packet can be efficiently stored and then retransmitted along each hop in the network.

Our examples and discussion have focused mainly on residential electricity usage. However, we believe many of these ideas can be extended to all forms of energy generation, distribution, and usage, including transportation. We believe computer scientists can play an important role in developing much of the core technology for the smart grid.

Some Lessons from the World of IT

In addition to the specific areas of expertise described above, computer scientists have learned some important lessons over the years that are critical in the design of the smart grid:

- *Create well-defined interfaces based on open standards.* Contrast the Internet, based on an open set of standards, along with an open process for revising and expanding them, to the proprietary interfaces of today's home automation technology, in which a single company or small collection of companies limits the abilities of others to create products compatible with its standards. The Internet has seen explosive growth with vibrant commercial activity. Its protocols are used even within private networks and corporate data centers. By contrast, home automation has failed to deliver significant levels of adoption or functionality.
- *Believe in Moore's Law and its implications.* It might not be cost-effective to instrument individual light bulbs today, but that should not deter us from planning for future systems that have many embedded sensors. The cost of electronics and computing will continue to drop relative to all other expenses. In the interim, much could be gained from a small number of sensors coupled with sophisticated learning algorithms.
- *Plan for scaling beyond anticipated levels.* Design systems with distributed control. Minimize the amount of information about the system that must be stored and maintained at any one location. Again, we can look to the Internet as an example of a scalable design. It has been able to evolve into a network with billions of nodes due largely to a set of fundamental design principles on how information is formatted and routed, and a very weak "best effort" service model.

- *Create systems that automatically configure and adapt to their environments.* Make sure the smart grid is resilient to component failures and anomalies. Minimize the human effort required to deploy new components or subsystems.
- *Devise protocols and algorithms that are guaranteed to work.* In a large, distributed system, especially one interfacing with a world where there are many unpredictable events, it is important to ensure that the system will function fully when possible and gracefully degrade otherwise. This can be done via a combination of careful, mathematically based design, analysis, and verification. Just running simulations of many different scenarios cannot exercise enough "corner cases" to ensure high reliability. For example, Intel took a loss of $475 million when the initial copies of the Pentium microprocessor computed the wrong results for the division instruction (for less than one out of a billion cases,) despite extensive simulations. Nowadays, Intel has a set of formal verification tools and a carefully engineered design flow such that it no longer relies on simulation to validate its floating-point hardware.
- *Embrace the powerful concept of overlays.* IT systems permit the construction of virtualized systems on top of underlying physical systems. Mobile telecommunications has seen the emergence of virtual telecommunication operators, essentially a marketing and billing operation built on top of someone else's mobile phone network. The same concept could apply for virtual utility companies, essentially billing clearinghouses. For example, if the Jones drive their electric car to their friends, the Garcias, and plug it in for a recharge, why should they not be billed for their energy usage like a mobile phone roaming on some other operator's network? Third parties may emerge to collect and analyze energy consumption and billing data, making this information available both to the consumer and to potential suppliers in a useful, organized, well-visualized, and privacy-enhanced form. Google's PowerMeter project suggests how this capability might develop.

Compared to the centrally planned and top-down engineering approach seen in today's utilities, the smart grid will need to provide a more open, collaborative, and entrepreneurial environment. This environment is very familiar to computer scientists, who have built many large-scale systems based on these principles. Most prominently, the Internet is a remarkable instance of a distributed system that has scaled in size, performance, and functionality well beyond what the original developers conceived. It has some flaws that would not be acceptable for the smart grid, but it demonstrates many important principles that could be adapted. On the other hand, computer scientists will need to work closely with power system designers and policymakers to be effective in an environment with very high capital costs, major safety concerns, and a myriad of legal and regulatory requirements. In the end, we believe experts in the energy domain must form deep collaborations with computer scientists to achieve the potential for the smart grid.

Leadership and Research Funding

Our country lacks any central organization, whether in the private sector or within the government, that has either the authority or the motivation to make major changes in our electric grid. Indeed, the utility companies operate as regulated monopolies with limited incentive to innovate. For example, much of the Federal funding for utility companies from the American Recovery and Reinvestment Act of 2009 has been spent on buying smart meters, not on

supporting new research and development efforts for identifying tomorrow's technological innovations today. Other countries, including Denmark, Spain, and Brazil are more advanced in their use of renewable energy sources. Numerous industry-driven smart grid initiatives have been undertaken, yet these have been fairly incremental in nature, and none has yielded a comprehensive system architecture for the future grid. Nevertheless, as the home of much of the information technology innovations of the last half century, the U.S. is well positioned to benefit from information technology enabled energy efficiency. Undoubtedly, a focus on this effort plays to America's technological strengths. **Thus, the Federal government should take the initiative in this area by (a) funding research into and development of proof-of-concept prototypes, (b) organizing testbeds and partnerships between the information technology and energy technology industries, and (c) developing new research communities at the intersection of distributed and energy systems.**

Current research funding for smart grid technology is highly fragmented across the Federal government, to the detriment of the collaborative and interdisciplinary approach that is necessary in order to appropriately address the many challenges that must be overcome in this area. For example, 84% of all Federal funding for basic research in computer science comes from the National Science Foundation's Directorate for Computer and Information Science and Engineering (CISE), but the NSF on its own cannot be expected to deal with the broad set of issues that must be addressed to develop and deploy smart grid technology. NSF/CISE does have plans to fund research in sustainable energy through a NSF-wide Science, Engineering, and Education for Sustainability (SEES) initiative included in the President's FY 11 budget request to Congress (see www.nsf.gov/sees). SEES promises to create interdisciplinary collaborations among researchers across the disciplines supported by the NSF to generate "decision capabilities and technologies aimed at mitigating against and adapting to environmental change that threatens sustainability." However, several factors will limit the impact this program can have on smart grid technology. First, it will span a broad range of topics, such as environmental protection and climate modeling, and hence the portion that addresses energy issues will be just a fraction of the total. Second, the U.S. Department of Energy serves as the locus of activity for energy-related technology and policy; consequently, with NSF as the funding source for SEES, it is not likely that the full complement of engineers, scientists, and policy researchers, across academia, industry, and government, will be at the table to work together on addressing the complexity of creating and implementing smart grid technology.

The natural home for smart grid research, the U.S. DoE, has to date provided only minimal funding to the core IT components. For example, the initial solicitation by the newly-established Advanced Research Projects Agency-Energy (ARPA-E) in 2009 led to 37 funded projects, but only one — a Stanford project on providing consumers with better information about their energy usage — has a significant IT component. More recent solicitations from ARPA-E appear to narrow the range of topics and even further reduce support for IT research. The DoE Office of Science has funded computing research in the past, but this work has been limited to high performance computing, simulation, and modeling. While these are certainly important and relevant research areas, they fail to touch the major IT issues underlying smart grid technology discussed in this report.

We feel it is imperative for the U.S. DoE, including ARPA-E and perhaps in conjunction with NSF/CISE, to provide avenues of support for highly collaborative, multi-disciplinary teams comprising computer scientists that seek to address the challenges of our energy infrastructure. At the very least, we have provided justification here for ARPA-E to include strong IT components in its current and future solicitations. The director of ARPA-E recently wrote, "The nation that successfully grows its economy with more efficient energy use, a clean domestic energy supply, and a smart energy infrastructure will lead the global economy of the 21st century. In many cases, [the U.S.] is lagging behind. We as a nation need to change course with fierce urgency." Achieving a truly smart energy infrastructure—for energy generation, distribution, and consumption—inherently requires basic and advanced computing research, as outlined above. Today, computer scientists are well-equipped to collaborate with other scientists and engineers, enabling the current concepts for the smart grid to be realized and then taking the vision to entirely new levels, yielding fundamental improvements in efficiency, reliability, and security, all the while reducing environmental impact. We can no longer afford to wait for this work to get underway.